Troubleshooting
Avaya Virtual Services Platform 9000

3.2
NN46250-700, 03.01
February 2012

Notice

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

Documentation disclaimer

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

Link disclaimer

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

Warranty

Avaya provides a limited warranty on its Hardware and Software (“Product(s)”). Refer to your sales agreement to establish the terms of the limited warranty. In addition, Avaya’s standard warranty language, as well as information regarding support for this Product while under warranty is available to Avaya customers and other parties through the Avaya Support Web site: http://support.avaya.com. Please note that if you acquired the Product(s) from an authorized Avaya reseller outside of the United States and Canada, the warranty is provided to you by said Avaya reseller and not by Avaya.

Licenses

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

Copyright

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

Third-party components

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

Trademarks

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

Avaya is a registered trademark of Avaya Inc.

All non-Avaya trademarks are the property of their respective owners, and “Linux” is a registered trademark of Linus Torvalds.

Downloading Documentation

For the most current versions of Documentation, see the Avaya Support Web site: http://support.avaya.com.

Contact Avaya Support

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

Contents

Chapter 1: Purpose of this document

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.

Chapter 2: Safety messages

This section describes the different precautionary notices used in this document. This section also contains precautionary notices that you must read for the 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.

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 or damage to Avaya products.

Caution:

ATTENTION

La mention Attention fournit des informations sur les moyens de prévenir une perturbation possible du service et d'éviter d'endommager les produits Avaya.

Caution:

ACHTUNG

Achtungshinweise bieten Informationen dazu, wie man mögliche Dienstunterbrechungen oder Schäden an Avaya-Produkten verhindert.

Caution:

PRECAUCIÓN

Los avisos de Precaución brindan información acerca de cómo evitar posibles interrupciones del servicio o el daño a los productos Avaya.

Caution:

CUIDADO

Os avisos de cuidado oferecem informações sobre como evitar possíveis interrupções do serviço ou danos aos produtos da Avaya.

Caution:

ATTENZIONE

Le indicazioni di attenzione forniscono informazioni per evitare possibili interruzioni del servizio o danni ai prodotti Avaya.

Chapter 3: New in this release

The following sections detail what is new in *Avaya Virtual Services Platform 9000 Troubleshooting*, NN46250–700, for Release 3.2.

Features

See the following sections for information about feature changes.

ACLI and SNMP log consolidation

Prior to Release 3.2, the system stored the CLI log and SNMP log in two separate files on the external flash: clilog.txt and snmplog.txt. The system did not send the SNMP log and ACLI command logs to the syslog server.

In Release 3.2, the ACLI command and SNMP logs are included in the main system log file, which can be sent to an external syslog server.

The following commands are obsolete:

- **`clilog maxfilesize <64–256000>`**
- **`clilog syslog-host enable`**
- **`snmplog maxfilesize <64–256000>`**

The commands **`show logging file module clilog`** and **`show logging file module snmplog`** replace previous commands to show ACLI and SNMP logs. The following commands only apply to log files generated by releases prior to Release 3.2:

- **`show clilog file`**
- **`save clilog file`**
- **`show snmplog file`**
- **`save snmplog file`**

For more information, see:

- Overview of traps and logs on page 52
- Enabling SNMP trap logging on page 84
- Viewing logs on page 69
- Configuring ACLI logging on page 71

Log message format

In Release 3.2, the format of module identification in log messages is updated. Prior to Release 3.2, the log message identified which CP module logged the message, and the slot number of

the affected module was embedded later. The new log message format identifies the module at the beginning of the message text. For more information, see: Log message format on page 54.

System log support of IPv6

In Release 3.2, 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. For more information, see:

- Overview of traps and logs on page 52
- Configuring a UNIX system log and syslog host on page 59
- Configuring the system log table on page 76

Displaying IGMPv3 information

In Release 3.2, a new parameter is added to the Internet Group Management Protocol version 3 (IGMPv3). IGMP is used in IP multicasting. IGMP establishes host memberships in multicast groups on an IP network. The new parameter, `show ip igmp group [group A.B.C.D. detail]`, displays IGMPv3 specific data. For more information on the new parameter, see Viewing IGMP group information on page 214.

Proprietary traps

Proprietary traps for Secure Shell (SSH) are updated to support IPv4 and IPv6.

Old trap	Replaced by
rcnSshUnauthorizedAccess	rcnaSshUnauthroizedAccess
rcnAuthenticationSuccess	rcnaAuthenticationSuccess
rcnSshSessionLogin	rcnaSshSessionLogin
rcnSshSessionLogout	rcnaSshSessionLogout

Proprietary traps for Virtual Router Redundancy Protocol (VRRP) are updated to support IPv6. The traps rcVrrpExtTrapStateTransition and rcVrrpTmpTrapNewMaster are added.

For more information, see:

- Proprietary traps on page 95

Feature licensing

Troubleshooting licensed routing protocols on page 229 is updated to include license requirements for new features.

Other changes

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

Removing external storage devices from the CP module

To safely remove the USB and external Compact Flash devices in the CP module you must follow a specific procedure. You must perform this procedure to prevent data loss or hardware damage.

For more information, see:

- Removing external storage devices from the CP module on page 113

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 you load a patch application properly. For more information, see:

Displaying current patch information on page 19

ACLI and EDM chapters

To improve document usability, most fundamentals chapters are grouped with applicable related chapters such as ACLI and EDM procedures.

ACLI Commands

Examples for ACLI Commands exist for most commands in the document.

Introduction chapter and navigation

Introduction chapters and navigation are removed.

Purpose of this document

To improve documentation usability, a brief description of the purpose of this document is now the first chapter.

Terminology

Terminology no longer exists in a separate document. Terminology is in a glossary at the end of this document.

Common procedures

Common procedures are incorporated in the document.

Chapter 4: 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

About this task

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

Procedure

1. Capture the current state of the chassis:

   ```
   terminal more disable
   ```

   ```
   show tech
   ```

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>
   ```

   This command executes three separate commands: `flight-recorder snapshot`, `flight-recorder trace`, and `flight-recorder archive`. The generated .tar file includes the following types of files:

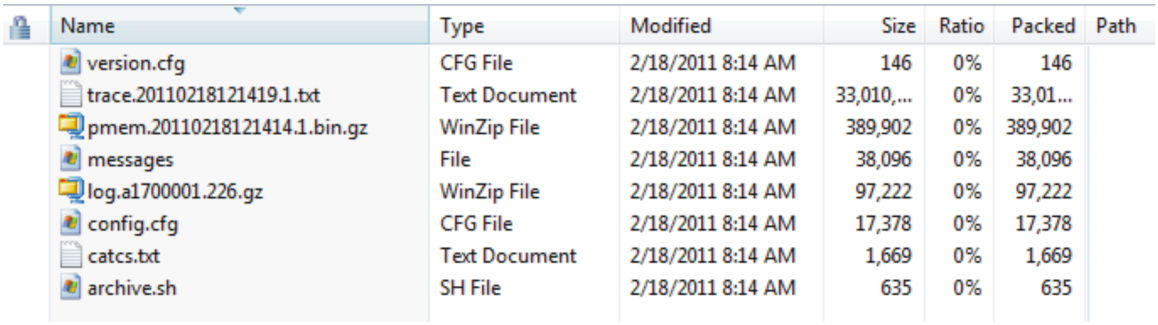

Name	Type	Modified	Size	Ratio	Packed	Path
version.cfg	CFG File	2/18/2011 8:14 AM	146	0%	146	
trace.20110218121419.1.txt	Text Document	2/18/2011 8:14 AM	33,010,...	0%	33,01...	
pmem.20110218121414.1.bin.gz	WinZip File	2/18/2011 8:14 AM	389,902	0%	389,902	
messages	File	2/18/2011 8:14 AM	38,096	0%	38,096	
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. Repeat step1 on page 17 on the IST peer VSP node, and if time permits, on other neighbor nodes to the VSP that exhibits the problem.
4. Reset the chassis:
   `reset –y`
5. Continue with Collecting data after you restart on page 20.

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.
```

The following example shows output of the `flight-recorder all 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/arc
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.2
0111019113952.4.txt.

Processing Flight-recorder archive for slot 4 ....

Flight-recorder archive for slot 4 complete, filename is /intflash/archive/4/arc
hive.20111019113956.4.tar.

Processing Flight-recorder snapshot for 11 ....

Flight-recorder snapshot for slot 11 complete, filename is /intflash/PMEM/11/pme
m.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/a
rchive.20111019114014.11.tar.

Processing Flight-recorder snapshot for SF4 ....

Flight-recorder snapshot for slot SF4 complete, filename is /intflash/PMEM/SF4/p
mem.20111019114030.SF4.bin.gz.

Processing Flight-recorder trace for SF4 ....

Flight-recorder trace for slot SF4 complete, filename is /intflash/flrec/SF4/tra
ce.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.
```

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.

Before you begin

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

Procedure

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)
================================================================================

Patch system information
   Status: idle
   Description: idle

Patch information
  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

About this task

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 17 for example output that shows how to identify the location and filename of the archive files.

2. Show 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/`.

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:

- detailed description of the problem
- date and time when the problem started
- frequency of the problem
- Is this a new installation?
- Have you searched the InSite Knowledge Base? Were related problem solutions found? Is there currently a workaround 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.
- 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.

Chapter 5: 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	*Avaya Virtual Services Platform 9000 Installation – Chassis*, NN46250-304
Control Processor, Switch Fabric, and interface module installation and replacement, cable routing	*Avaya Virtual Services Platform 9000 Installation – Modules*, NN46250-301
Cooling module installation and removal	*Avaya Virtual Services Platform 9000 Installation – Cooling Modules*, NN46250-302
Optical component installation and cleaning	*Avaya Virtual Services Platform 9000 Installation – SFP Hardware Components*, NN46250-305
Power supply installation and removal	*Avaya Virtual Services Platform 9000 Installation – AC Power Supply*, 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 it is fully operational, you can be more effective at troubleshooting 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

```
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 (3.2.0.0) (DEV)
        SysName      : CB-SWA
        SysUpTime    : 10 day(s), 21:10:36
        SysContact   : http://support.avaya.com/
        SysLocation  : 211 Mt. Airy Road,Basking Ridge,NJ 07920

Chassis Info:

        Chassis            : 9012
        Serial#            : SAN1223008S
        H/W Revision       :
        H/W Config         :
        NumSlots           : 12
        NumPorts           : 50
        BaseMacAddr        : 00:24:7f:9f:60:00
        MacAddrCapacity    : 4096

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

Chapter 6: 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 localized to 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

```
VSP-9012:1#show interface vlan

================================================================================
                                  Vlan Basic
================================================================================
VLAN                              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	The switch is experiencing a port	Verify that the cable connections to the link partner are correct. Verify

Symptom	Probable cause	Corrective action
is off or does not blink (and you believe that traffic is present).	connection problem, or the link partner is not auto-negotiating properly.	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 3 minutes.	This LED shows steady amber at module reset. This is normal behavior. 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. 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. On the master CP module, the LED transitions to medium blinking green waiting for communication with all the IO, SF, and standby CP modules. On the IO, SF, and standby CP modules, the LED transitions to steady green after communication with the master CP module is established. On the master CP module, the LED transitions to steady green after communication is established with all other modules and the system	Not applicable.

Symptom	Probable cause	Corrective action
	transitions to the ready state.	
System temperature LED on Master CP module is steady red.	One or more modules exceeds the normal operating temperature.	Identify the module that exceeds the normal operating temperature. The Online LED on the module that exceeds the normal operating temperature will change color from steady green to steady red. Investigate possible cooling module failure. The monitoring logic polls the hardware every 30 seconds, and reports a temperature above 60°C as an over-temperature condition. In this case, the system automatically powers off the module if the condition persists for 15 minutes. The system powers off the module immediately if regular polling indicates a temperature of 70°C or higher. This action protects the offending module and adjacent hardware from the risk of permanent damage. During over-temperature conditions, the system raises an alarm and generates an SNMP trap. You can also use ACLI to obtain current temperature readings.
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 the cables are wired correctly.

1000BASE-T cables

1 Gb/s 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; 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 plus (SFP+), see *Avaya Virtual Services Platform 9000 Installation —SFP Hardware Components*, NN46250-305.

Alarm database

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

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 69.

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 has taken 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 backup CP has been taken 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 will rejoin the system after it boots normally. If the problem persists and the CP once again is taken out of service by the system, call Avaya Support and return the CP module for service.

Chapter 7: 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. The 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.

Certain protocols and tools are tailored 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:
                                                    Interface
Network       Destination        Netmask          Gateway      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 Service (DNS) server, another switch, firewall, or access point can 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. 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.

An interface that supports DDM is a Digital Diagnostic Interface (DDI). These devices provide real-time monitoring of individual DDI 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.

For information about DDM and SFPs, see *Avaya Virtual Services Platform 9000 Installation — SFP Hardware Components*, NN46250-305 and *Avaya Virtual Services Platform 9000 Performance Management*, NN46250-701.

Example

```
VSP-9012:1#show pluggable-optical-modules config

================================================================================
                  Pluggable Optical Module Global Configuration
================================================================================
                            ddm-monitor : disabled
                   ddm-monitor-interval : 5
                         ddm-traps-send : enabled
                     ddm-alarm-portdown : disabled
```

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. Use port mirroring in diagnostic sniffing—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. Only one type of mirroring destination is supported 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 per lane
One port to one port	Supported, no restriction per 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

Function	Support information
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 (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 same MGID is used.

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, that is, inPort and inVLAN) and tx (egress, that 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 the ports have to be removed 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 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. Traffic patterns that match the ACL or ACE with an action of permit are forwarded to the destination and also to the mirroring port. Traffic patterns that match an ACE with an action of drop (deny) are not forwarded to the destination, but still reach the mirroring port For example, for an ACL or ACE with a match action of permit and debug mirroring enabled, packets are mirrored 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, then the flow-based mirror takes precedence.

For more information about how to configure ACLs and ACEs, see *Avaya Virtual Services Platform 9000 Configuration — QoS and IP Filtering*, 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 though 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 traffic with a given MAC source address (SA) or MAC destination address (DA) is copied 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, traffic with a given source or destination MAC address is copied 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 it is transmitted on the port. Remote mirrored encapsulated frames are identified when the configured remote mirroring destination MAC address is detected 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, a static FDB entry is added 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 remote mirroring feature is disabled 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 ether type, and monitor control. The original CRC-32 is stripped from a mirrored packet, and a new CRC-32 is computed 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 should 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.

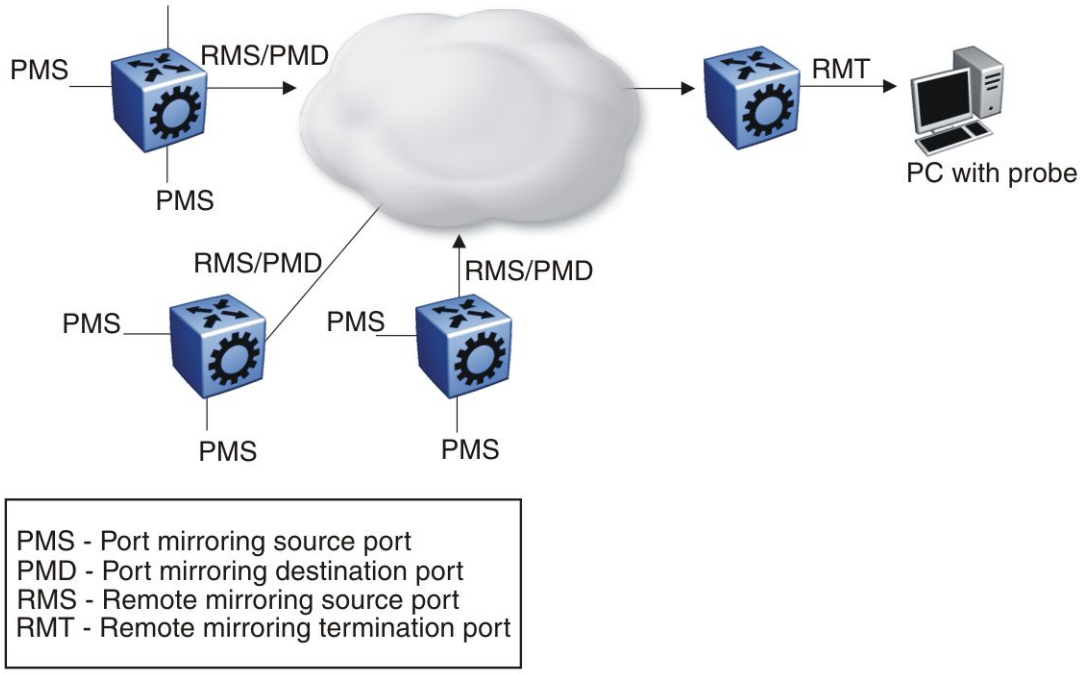

Figure 1: Layer 2 remote mirroring

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 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 OAM ports.

Note the following information:

- If the RMS is a tagged port, the mirrored packet is encapsulated and transmitted with the VLAN ID of the RMS port and forwarded 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 frame is transmitted on the port with the VLAN ID of the mirrored packet (the original packet).
- If port mirroring is disabled, no packets are remote mirrored.
- Packets are captured as long as the RMT is reachable.
- When you enable or disable Layer 2 remote mirroring, a trap is sent 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.

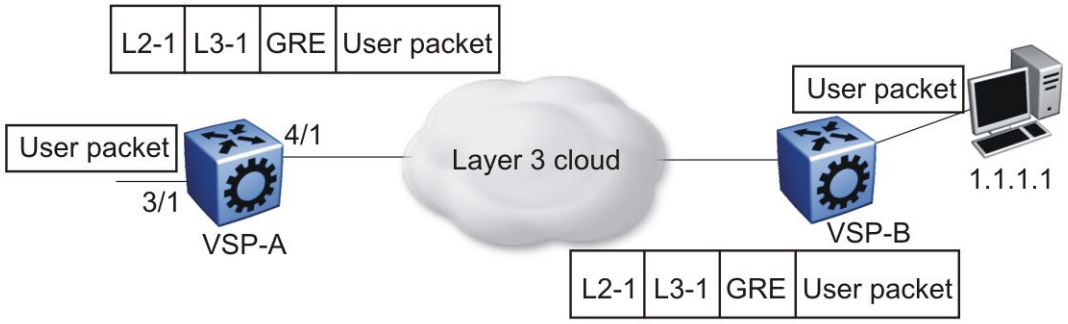

Figure 2: Layer 3 remote mirroring

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 mirroringflow 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 monitor 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 packets are dropped, 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 of 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 IP address creation is restricted based on the mirror configuration.
- You cannot use remote mirroring on OAM ports.

Note the following information:

- If port mirroring is disabled, no packets are remote mirrored.
- 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 *Avaya Virtual Services Platform 9000 Administration*, NN46250-600.

All captured packets are stored in the secondary CP module, used as the PCAP engine. The master CP module maintains protocol handling and is not affected by capture activity.

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.

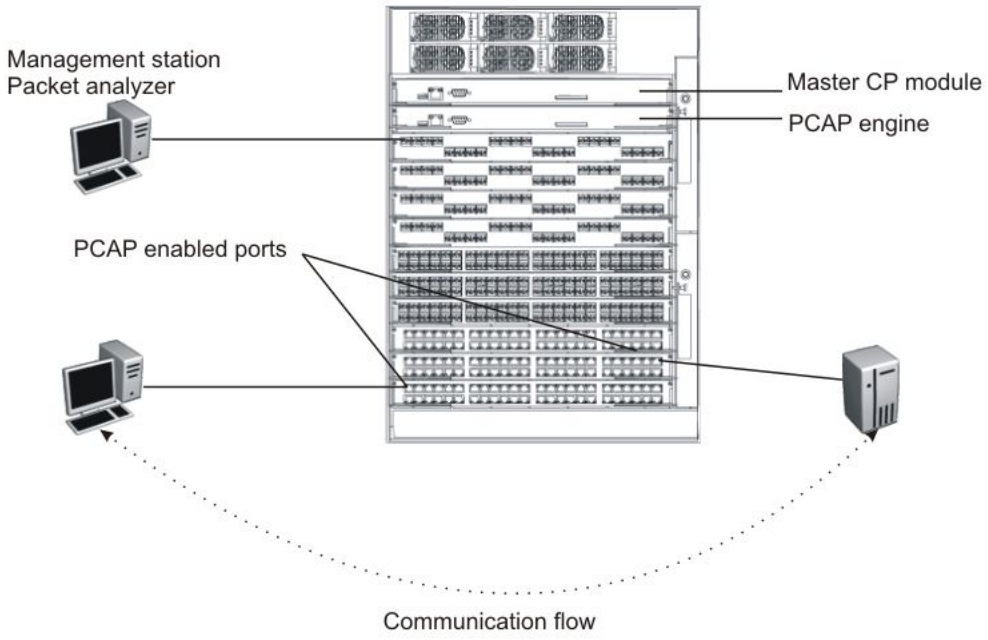

Figure 3: PCAP configuration

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 provides 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 capture filter is disabled 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`
	`pcap capture-filter 1 dstmac 00:00:00:00:00:0a`
		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 engine receives

Example	PCAP configuration	Resulting action

	`pcap capture-filter 1 action trigger-off`
	`pcap capture-filter 1 enable`
	`pcap enable`
		the matching packet, the capture filter and PCAP globally are disabled. The PCAP engine does not capture more packets.

PCAP capture filter with action trigger-on, match dstmac, and interface mode tx.	`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`
		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 capture filter is disabled and the engine captures all packets, starting with the matched one, until you disable PCAP manually.

PCAP capture filter with action capture, match dstmac, and interface mode both.	`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`
		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.

PCAP capture filter with action trigger-on with timer, match dstmac, and interface mode rx.	`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`
		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).

Example	PCAP configuration	Resulting action
	`pcap capture-filter 1 enable` `pcap enable`	A value of 1000 equals 1 second.

PCAP capture filter with action trigger-on with refresh-timer, match dstmac, and interface mode rx.	`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`
		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.
		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.

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 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 using 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 (that is, if enable is 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 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. This data can later be accessed on-demand when debugging systems 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 will be lost. Persistent Memory snapshots are taken 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 central 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 will 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 the ACLI. The following commands are used to make use of this feature:

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

  This command outputs the Persistent Memory information for the specified slot to a file. The user will be notified 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 user will be notified 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 user will be notified 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 user will be notified 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 *Avaya Virtual Services Platform 9000 Performance Management*, 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 could not be forwarded 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 could 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. The trace toolset can be used to trace multiple modules simultaneously and provides 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 trace data is required 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 8: 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.

- the 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

  This message is sent by an SNMP agent 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 send 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 *Avaya Virtual Services Platform 9000 Security*, NN46250–601.

System log messaging

On a UNIX-based management platform, you can use system log (syslog) messaging to manage event messages. The 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 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 CLILOG module to track all ACLI commands executed and for fault management purposes. The ACLI commands are logged to the system log file as CLILOG module.

View logs for SNMPLOG module to track SNMP logs. The SNMP operation log is logged 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. All system messages are tagged 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
```

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. No action is required.

Severity level	Definition
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, 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

UNIX system error codes	System log severity level	Internal VSP 9000 severity level
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 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 external flash. If the external flash does not exist or the system configuration does not log to external flash, the system logs to internal flash instead.

To log to a file on external or 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 internal memory. The system saves internal log messages in a circular list in memory, which overwrite older log messages as the log fills. Unlike the log messages in a log file, the internal log messages in 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 is named as log.xxxxxxxx.sss format. 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 sequence number of the log file is incremented 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, a new log file with the sequence number 000 is created. After a restart, the system finds the newest log file from both external flash and 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 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, a new log file is created 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 log file transfer is unsuccessful, the system keeps the old log files on external flash or 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 .xxxxxxxx.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.

Chapter 9: Log configuration using ACLI

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

Configuring a UNIX system log and syslog host

Before you begin

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

About this task

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

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. Enable the system log:
   ```
   syslog enable
   ```
2. Specify the IP header in syslog packets:
   ```
   syslog ip-header-type <circuitless-ip|default|management-virtual-ip>
   ```
3. Configure the maximum number of syslog hosts:
   ```
   syslog max-hosts <1-10>
   ```
4. Create the syslog host:
   ```
   syslog host <1-10>
   ```
5. Configure the IP address for the syslog host:
   ```
   syslog host <1-10> address WORD <0–46>
   ```
6. Enable the syslog host:
   ```
   syslog host <1-10> enable
   ```
   Configure optional syslog host parameters by using the variables in the following variable definition tables.
7. View the configuration to ensure it is correct:

```
show syslog [host <1–10>]
```

Example

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

VSP-9012:1(config)#syslog host 1 address 47.17.143.52

VSP-9012:1(config)#syslog host 1 enable
```

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

                Id : 1
            IpAddr : 47.17.143.52
           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
```

```
VSP-9012:(config)# syslog host 2 address fe80:0:0:0:22b:4eee:fe5e:
73fd udp-port 515

VSP-9012:(config)# syslog host 2 udp-port 515

VSP-9012:(config)# syslog host 2 enable

VSP-9012:(config)#
```

```
VSP-9012:1(config)#show syslog host 2

                Id : 2
            IpAddr : fe80:0:0:0:22b:4eee:fe5e:73fd
           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.

Table 8: Variable definitions

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. • If the value is default, the IP address of the VLAN is used for syslog packets that are transmitted in-band using input/output (I/O) ports. For syslog packets that are transmitted out-of-band through the management port, the physical IP address of the master CPU is used in the IP header. • 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 are transmitted 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 9: 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

Variable	Value
	x:x: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>] [<info|warning|error|fatal>]	Specifies the severity levels for which to send syslog messages for the specified modules. The default is info.
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

Before you begin

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

About this task

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

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.

Procedure

1. Define which messages to log:
   `logging level <0-4>`
2. Write the log file from memory to a file:
   `logging write WORD<1-1536>`
3. Show logging on the screen:
   `logging screen`

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#logging level 0
VSP-9012:1(config)#logging write log2
VSP-9012:1(config)#logging screen
```

Variable definitions

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

Table 10: Variable definitions

Variable	Value
level <0-4>	Shows and configures the logging level. The level is one of the following values: • 0: Information; all messages are recorded • 1: Warning; only warning and more serious messages are recorded • 2: Error; only error and more serious messages are recorded • 3: Manufacturing; this parameter is not available for customer use • 4: Fatal; only fatal messages are recorded
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

Variable	Value
	external flash. Use the no form of the command to stop logging to external flash and log to internal flash instead: `no logging logToExtFlash`
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: `no logging screen`
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

Before you begin

- The IP address you configure for the remote host must be reachable at the time of configuration.
- You must log on to the Global Configuration mode in ACLI.

About this task

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.

Procedure

Configure the remote host address for log transfer:

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

Example

```
VSP-9012:1>enable

VSP-9012:1#configure terminal

VSP-9012:1(config)#logging transferFile 1 address 172.16.120.10
```

Variable definitions

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

Table 11: Variable definitions

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 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

Before you begin

- You must install a CF card in the CP module before you can log to external storage.
- You must log on to the Global Configuration mode in ACLI.

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

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

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 anytime 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 external flash does not exist, the system raises an alarm, and then logs to internal flash instead.

Procedure

1. Enable system logging to a CF card:
   `boot config flags logging`
2. Configure the logfile parameters:
   `boot config logfile <64-500> <500-16384> <10-90>`

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#boot config logfile 64 600 10
```

Variable definitions

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

Table 12: Variable definitions

Variable	Value
flags logging	Enables or disables logging to a file on external flash. The log file is named using the format log.xxxxxxxx.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. Multiple sequence numbers are generated 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 • *<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

Before you begin

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

About this task

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

Procedure

1. Configure system message control action:
   `sys msg-control action <both|send-trap|suppress-msg>`
2. Configure the maximum number of messages:
   `sys msg-control max-msg-num <2-500>`
3. Configure the interval:
   `sys msg-control control-interval <1-30>`
4. Enable message control:
   `sys msg-control`

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#sys msg-control action suppress-msg
VSP-9012:1(config)#sys msg-control max-msg-num 10
VSP-9012:1(config)#sys msg-control control-interval 15
VSP-9012:1(config)#sys msg-control
```

Variable definitions

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

Table 13: Variable definitions

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

Before you begin

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

About this task

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

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

Configure the force message control option:

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

Example

```
VSP-9012:1> enable

VSP-9012:1# configure terminal
```

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

```
VSP-9012:1(config)# sys force-msg ****
```

Variable definitions

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

Table 14: Variable definitions

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

About this task

View log files by file name, category, severity, or CP module to identify possible problems.

View ACLI command and SNMP trap logs, which are logged as normal log messages and logged to the system log file.

Procedure

Show log information:

```
show logging file [alarm][CPU WORD<0-25>] [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

```
VSP-9012:1>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 14:30:07.036] 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
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.

Table 15: Variable definitions

Variable	Value
alarm	Displays alarm log entries.
CPU *WORD<0-25>*	Filters and lists the logs according to the CP module that generated the message. Specify a string length of 0–25 characters. To specify multiple filters, separate each CP module by the vertical bar (|), for example,`show logging file CPU CP1|CP2|IO1` . Following are some of the available CPU qualifiers: • CP1 • CP2

Variable	Value
	• IO1 • IO2 • SF1 • SF6
event-code *WORD<0–10>*	Specifies a number that precisely identifies the event reported.
module *WORD<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, OSPF|FILTER|QOS.
name-of-file *WORD<1-99>*	Displays the valid logs from this file. For example, /intflash/logcopy.txt. You cannot use this command on the current log file—the file into which the messages are currently logged. Specify a string length of 1–99 characters.
save-to-file *WORD<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: /intflash/<filename>, /extflash/<filename>, or /usb/<filename>.
severity *WORD<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, ERROR|WARNING|FATAL.
tail	Shows the last results first.
vrf *WORD<0–32>*	Specifies the name of a VRF instance to show log messages that only pertain to that VRF.

Configuring ACLI logging

Before you begin

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

About this task

Use ACLI logging to track all ACLI commands executed and for fault management purposes. The ACLI commands are logged 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.

Procedure

1. Enable ACLI logging:
   ```
   clilog enable
   ```
2. Disable ACLI logging:
   ```
   no clilog enable
   ```
3. Ensure that the configuration is correct:
   ```
   show clilog
   ```
4. View the ACLI log:
   ```
   show logging file module clilog
   ```
5. View the ACLI log. The following command only applies to 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 14:30:07.036] 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
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 `clilog` commands.

Table 16: Variable definitions

Variable	Value
enable	Activates ACLI logging. To disable, use the `no clilog enable` command.

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

Note:

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

Table 17: Variable definitions

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

Chapter 10: 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

About this task

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

Procedure

1. In the navigation tree, open 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.

Name	Description
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	Specifies the IP header in syslog packets to circuitlessIP, default, or managementVIP. • If the value is default, the IP address of the VLAN is used for syslog packets that are transmitted in-band using input/output (I/O) ports. For syslog packets that are transmitted out-of-band through the management port, the physical IP address of the master CPU is used in the IP header. • If the value is managementVIP, the virtual management IP address of the device is used in the IP header for syslog packets that are transmitted out-of-band only through the management port. • If the value is circuitlessIP, the circuitless IP address is used in the IP header for all syslog messages (in-band or out-of-band). If you configure multiple circuitless IPs, the first circuitless IP configured is used. The default value is default.

Configuring the system log table

About this task

Use the system log table to customize the mappings between the severity levels and the type of alarms.

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, open 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. 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.
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.
HostFacility	Specifies the syslog host facility used to identify messages (LOCAL0 to LOCAL7). The default is LOCAL7.
Severity	Specifies the message severity for which syslog messages are sent. 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.

Name	Description
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.

Chapter 11: 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

Before you begin

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

About this task

Configure an SNMP host so that the system can forward SNMP traps to a host for monitoring. You can use SNMPv1, SNMPv2c, or SNMPv3. You configure the target table parameters (security name and model) as part of the host configuration.

Procedure

1. Configure an SNMPv1 host:
   ```
   snmp-server host WORD<1-256> [port <1-65535>] v1 WORD<1-32>
   [filter WORD<1-32>]
   ```
2. 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>]
   ```
3. 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>]
   ```
4. Ensure that the configuration is correct:

```
show snmp-server host
```

Example

```
VSP-9012:1> enable

VSP-9012:1#configure terminal
```

Configure the target table entry:

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

Variable definitions

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

Table 18: Variable definitions

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: 1. timeout <0-2147483647> specifies the timeout value in seconds with a range of 0–214748364. 2. retries <0-255> specifies the retry count value with a range of 0–255. 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 host server port number.
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.

Configuring an SNMP notify filter table

Before you begin

- You must log on to the Global Configuration mode in ACLI.
- For more information about the notify filter table, see RFC3413.

About this task

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. Create a new notify filter table:
   `snmp-server notify-filter` *`WORD<1-32>`* *`WORD<1-32>`*
2. Ensure that the configuration is correct:
   `show snmp-server notify-filter`

Example

```
VSP-9012:1> enable

VSP-9012:1# configure terminal

VSP-9012:1(config)# snmp-server notify-filter profile1
+99.3.4.1.4.3.1.1.4.1

VSP-9012:1(config)# 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           0x7f
```

Variable definitions

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

Table 19: Variable definitions

Variable	Value
WORD<1-32> WORD<1-32>	Creates a notify filter table. The first instance of *WORD<1-32>* specifies the name of the filter profile with a string length of 1–32. The second instance of *WORD<1-32>* identifies the filter subtree OID with a string length of 1–32. 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. 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.

Configuring SNMP interfaces

Before you begin

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

About this task

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

Procedure

1. Configure the destination and source IP addresses for SNMP traps:
   ```
   snmp-server sender-ip {A.B.C.D} {A.B.C.D}
   ```
2. If required, send the source address (sender IP) as the sender network in the notification message:
   ```
   snmp-server force-trap-sender enable
   ```
3. If required, force the SNMP and IP sender flag to use the same value:
   ```
   snmp-server force-iphdr-sender enable
   ```

Example

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

```
VSP-9012:1(config)# snmp-server sender-id 44.16.13.12 42.13.15.10
VSP-9012:1(config)# snmp-server force-trap-sender enable
VSP-9012:1(config)# snmp-server force-iphdr-sender enable
```

Enable the generation of authentication traps:

```
VSP-9012:1(config)# snmp-server authentication-trap enable
```

Variable definitions

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

Table 20: Variable definitions

Variable	Value
agent-conformance enable	Enables the agent conformance mode. Conforms to MIB standards if disabled. If you activate this option, feature configuration is stricter and error handling less informative. Avaya recommends that you do not activate this option; it is not a normally supported mode of operation.
authentication-trap enable	Activates the generation of authentication traps.
force-iphdr-sender enable	Automatically configures the SNMP and IP sender to the same value. The default is no force-iphdr-sender enable .
force-trap-sender enable	Sends the configured source address (sender IP) as the sender network in the notification message.
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 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.

Enabling SNMP trap logging

Before you begin

- You must log on to the Global Configuration mode in ACLI.
- You must configure and enable the syslog server.

About this task

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

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.

Procedure

1. Enable SNMP trap logging:
   `snmplog enable`
2. Disable SNMP trap logging:
   `no snmplog enable`
3. View the contents of the SNMP log:
   `show logging file module snmplog`
4. View the contents of the SNMP log. The following command only applies to 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
VSP-9012:1(config)# snmplog enable
VSP-9012:1(config)# show logging file module snmplog
```

Variable definitions

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

Table 21: Variable definitions

Variable	Value
enable	Enables the logging of traps. Use the command `no snmplog enable` 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: • grep WORD<1–255> performs a string search in the log file. *WORD<1–255>* is the string, of up to 255 characters in length, to match. • tail shows the last results first.

Chapter 12: 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 in 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

About this task

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, open 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, type the SnmpAdminString.
11. In the **TMask** box, type the mask.
12. In the **MMS** box, type the maximum message size.

13. 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.
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.
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

About this task

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, open 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.

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

About this task

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, open 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	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

Name	Description
	• trap—messages generated contain Unconfirmed-Class Protocol Data Units (PDU) • inform—messages generated contain Confirmed-Class PDUs The default value is trap.

Configuring SNMP notify filter profiles

About this task

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

Procedure

1. In the navigation tree, open 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.x. format.
7. In the **Mask** box, type the mask location in hex string format.
8. From the **Type** options, select **configure the filter flag**.
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

Name	Description
	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.

Configuring SNMP notify filter profile table parameters

Before you begin

- The notify filter profile exists.

About this task

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

Procedure

1. In the navigation tree, open 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.

Name	Description
NotifyFilterProfileName	Specifies the name of the filter profile to use to generate notifications.

Enabling authentication traps

About this task

Enable the SNMP agent process to generate authentication-failure traps.

Procedure

1. In the navigation tree, open 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.
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

Chapter 13: Traps reference

The Virtual Services Platform 9000 generates alarms, traps, and logs. For more information about specific log messages, see the *Virtual Services Platform 9000 Logs Reference*, NN46250-702. This section provides information about traps.

Proprietary traps

The following tables describe Avaya proprietary traps for the Virtual Services Platform 9000. All these traps have a status of current.

Table 22: 1.3.6.1.4.1.2272.1.21.0.xx series

OID	Notification type	Objects	Description
1.3.6.1.4.1.2272.1.21.0.1	rcnCardDown	rcCardIndex rcCardAdminStatus rcCardOperStatus	A rcCardDown trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the rcCardOperStatus object for one of its cards is about to transition into the down state.
1.3.6.1.4.1.2272.1.21.0.2	rcnCardUp	rcCardIndex rcCardAdminStatus rcCardOperStatus	A rcCardUp trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the rcCardOperStatus object for one of its cards is about to transition into the up state.
1.3.6.1.4.1.2272.1.21.0.3	rcnErrorNotification	rcErrorLevel rcErrorCode rcErrorText	A rcErrorNotification trap signifies that the SNMPv2 entity, acting in an agent role, has detected that an error condition has occurred.
1.3.6.1.4.1.2272.1.21.0.4	rcnStpNewRoot	rcStgId	A rcStpNewRoot trap signifies that the SNMPv2 entity, acting in an agent role, has detected the Spanning Tree Protocol has declared the device to be the new root of the spanning tree.
1.3.6.1.4.1.2272.1.21.0.5	rcnStpTopologyChange	rcStgId rcPortIndex	A rcStpTopologyChange trap signifies that the SNMPv2 entity, acting in an agent role, has

OID	Notification type	Objects	Description
			detected the Spanning Tree Protocol has gone due a topology change event.
1.3.6.1.4.1.2272.1.21.0.6	rcnChasPowerSupplyDown	rcChasPowerSupplyId rcChasPowerSupplyOperStatus	A rcChasPowerSupplyDown trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the rcChasPowerSupplyOperStatus object for one of its power supply unit is about to transition into the down state.
1.3.6.1.4.1.2272.1.21.0.7	rcnChasFanDown	rcChasFanId rcChasFanOperStatus	A rcChasFanDown trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the rcChasFanOperStatus object for one of its power supply unit is about to transition into the down state.
1.3.6.1.4.1.2272.1.21.0.8	rcnLinkOscillation	rcPortIndex	A rcLinkOscillation trap signifies that the SNMPv2 entity, acting in an agent role, has detected an excessive number of link state transitions on the specified port.
1.3.6.1.4.1.2272.1.21.0.9	rcnMacViolation	rcErrorText rcPortIndex	A rcMacViolation trap signifies that the SNMPv2 entity, acting in an agent role, has received a PDU with an invalid source MAC address.
1.3.6.1.4.1.2272.1.21.0.11	rcn2kCardDown	rc2kCardIndex rc2kCardFrontAdminStatus rc2kCardFrontOperStatus	A rcCardDown trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the rcCardOperStatus object for one of its cards is about to transition into the down state.
1.3.6.1.4.1.2272.1.21.0.12	rcn2kCardUp	rc2kCardIndex rc2kCardFrontAdminStatus rc2kCardFrontOperStatus	A rcCardUp trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the rcCardOperStatus object for one of its cards is about to transition into the up state.
1.3.6.1.4.1.2272.1.21.0.13	rcn2kTemperature	rc2kChassisTemperature	A rc2kTemperature trap signifies that the SNMPv2 entity, acting in an agent role, has detected the chassis is overheating.

OID	Notification type	Objects	Description
1.3.6.1.4.1.2272.1.21.0.14	rcnChasPowerSupplyUp	rcChasPowerSupplyId rcChasPowerSupplyOperStatus	A rcChasPowerSupplyUp trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the rcChasPowerSupplyOperStatus object for one of its power supply unit is about to transition into the up state.
1.3.6.1.4.1.2272.1.21.0.16	rcnStpTCN	rcStgId rcPortIndex rcStgBridgeAddress	A rcStpTopologyChange trap signifies that the SNMPv2 entity, acting in an agent role, has detected the Spanning Tree Protocol has gone due to a topology change event.
1.3.6.1.4.1.2272.1.21.0.17	rcnSmltIstLinkUp	—	A rcSmltIstLinkUp trap signifies that the split MLT link is from down to up.
1.3.6.1.4.1.2272.1.21.0.18	rcnSmltIstLinkDown	—	A rcSmltIstLinkDown trap signifies that the split MLT link is from up to down.
1.3.6.1.4.1.2272.1.21.0.19	rcnSmltLinkUp	rcMltSmltId	A rcMltSmltId trap signifies that the split SMLT link is up.
1.3.6.1.4.1.2272.1.21.0.20	rcnSmltLinkDown	rcMltSmltId	A rcMltSmltId trap signifies that the split SMLT link is down.
1.3.6.1.4.1.2272.1.21.0.21	rcnChasFanUp	rcChasFanId rcChasFanOperStatus	A rcChasFanUp trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the rcChasFanOperStatus object for one of its power supply unit is about to transition into the up state.
1.3.6.1.4.1.2272.1.21.0.22	rcnPasswordChange	rcCliPasswordChange rcCliPassChangeResult	A rcPasswordChange trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the one of the cli passwords is changed.
1.3.6.1.4.1.2272.1.21.0.23	rcnEmError	rc2kCardIndex rcChasEmModeError	A rcEmError trap signifies that the SNMPv2 entity, acting in an agent role, has detected Em error.
1.3.6.1.4.1.2272.1.21.0.26	rcnSmartCpldTimerFired	rc2kCardIndex	A rcSmartCpldTimerFired trap signifies that the cpld timer fired.

OID	Notification type	Objects	Description
1.3.6.1.4.1.2272.1.21.0.27	rcnCardCpldNotUpDate	rc2kCardIndex	A rcCardCpldNotUpDate trap signifies that the cpld is not up to date.
1.3.6.1.4.1.2272.1.21.0.28	rcnIgapLogFileFull	—	A rcIgapLogFileFull trap signifies that the Igap accounting time-out Log File reach the maximum.
1.3.6.1.4.1.2272.1.21.0.29	rcnCpLimitShutDown	rcPortIndex ifAdminStatus ifOperStatus rcPortCpLimitShutDown	A rcCpLimitShutDown trap signifies that the cp limit for the port is shutting down.
1.3.6.1.4.1.2272.1.21.0.30	rcnSshServerEnabled	rcSshGlobalPort	A rcSshServerEnabled trap signifies that the SSH server is enabled.
1.3.6.1.4.1.2272.1.21.0.31	rcnSshServerDisabled	rcSshGlobalPort	A rcSshServerDisabled trap signifies that the SSH server is disabled.
1.3.6.1.4.1.2272.1.21.0.35	rcnHaCpuState	rc2kCardIndex rcL2RedundancyHaCpuState	A rcHaCpuState trap signifies that the state of the HA-CPU.
1.3.6.1.4.1.2272.1.21.0.36	rcnInsufficientMemory	rc2kCardIndex	A rcInsufficientMemory trap indicates insufficient memory on CPU blade for proper operation.
1.3.6.1.4.1.2272.1.21.0.37	rcnSaveConfigAction	rcSysActionL1	A rcSaveConfigAction trap indicates the switch run time or boot configuration is being saved.
1.3.6.1.4.1.2272.1.21.0.38	rcnLoopDetectOnPort	rcVlanId rcPortIndex	A rcLoopDetectOnPort trap indicates that a loop has been detected on a port. The VLAN on that port will be disabled.
1.3.6.1.4.1.2272.1.21.0.39	rcnbgpEstablished	rcIpBgpPeerIpAddress rcIpBgpPeerLastError rcIpBgpPeerState	The BGP Established event is generated when the BGP FSM enters the established state.
1.3.6.1.4.1.2272.1.21.0.40	rcnbgpBackwardTransition	rcIpBgpPeerIpAddress rcIpBgpPeerLastError rcIpBgpPeerState	The BGPBackwardTransition Event is generated when the BGP FSM moves from a higher numbered state to a lower numbered state.
1.3.6.1.4.1.2272.1.21.0.41	rcnAggLinkUp	rcMltId	A rcAggLinkUp trap is generated when the operational state of the

OID	Notification type	Objects	Description
			aggregator changes from down to up.
1.3.6.1.4.1.2272.1.21.0.42	rcnAggLinkDown	rcMltId	A rcAggLinkDown trap is generated when the operational state of the aggregator changes from up to down.
1.3.6.1.4.1.2272.1.21.0.43	rcnIgmpNewGroupMember	rcIgmpGroupIfIndex rcIgmpGroupIpAddress rcIgmpGroupInPort rcIgmpGroupMember	An IgmpNewGroupMember trap signifies that a new member has come on an interface.
1.3.6.1.4.1.2272.1.21.0.44	rcnIgmpLossGroupMember	rcIgmpGroupMembers rcIgmpGroupIpAddress rcIgmpGroupInPort rcIgmpGroupIfIndex	An IgmpLossGroupMember trap signifies that a group member has been lost on an interface.
1.3.6.1.4.1.2272.1.21.0.45	rcnIgmpNewQuerier	igmpInterfaceIfIndex igmpInterfaceQuerier	An igmpNewQuerier trap signifies that a new querier has come up on an interface.
1.3.6.1.4.1.2272.1.21.0.46	rcnIgmpQuerierChange	igmpInterfaceIfIndex rcIgmpInterfaceExtnNewQuerier igmpInterfaceQuerier	An rcIgmpQuerierChange trap signifies that the querier has changed.
1.3.6.1.4.1.2272.1.21.0.59	rcnFdbProtectViolation	rcPortIndex rcVlanId	The rcFdbProtectViolation trap signifies that the has violated the user configured limit for total number of fdb-entries learned on that port.
1.3.6.1.4.1.2272.1.21.0.60	rcnLogMsgControl	rcSysMsgLogFrequency rcSysMsgLogText	A rcMsgControl trap signifies whether the number of times of repetition of the particular Log message has exceeded the particular frequency/count or not.
1.3.6.1.4.1.2272.1.21.0.61	rcnSaveConfigFile	rcSysActionL1 rcSysConfigFileName	A rcSaveConfig trap signifies that either the runtime config or the

OID	Notification type	Objects	Description
			boot config has been saved on the switch.
1.3.6.1.4.1.2272.1.21.0.62	rcnDNSRequestResponse	rcSysDnsServerListIpAddr rcSysDnsRequestType	A rcDnsRequestResponse trap signifies that the switch had sent a query to the DNS server or it had received a successful response from the DNS Server.
1.3.6.1.4.1.2272.1.21.0.63	rcnDuplicateIpAddress	ipNetToMediaNetAddress ipNetToMediaPhysAddress	A rcDuplicateIpAddress trap signifies that a duplicate IP address is detected on the subnet.
1.3.6.1.4.1.2272.1.21.0.64	rcnLoopDetectPortDown	rcPortIndex ifAdminStatus ifOperStatus	A rcLoopDetectPortDown trap signifies that a loop has been detected on a port and the port is going to shut down.
1.3.6.1.4.1.2272.1.21.0.67	rcnLoopDetectMacDiscard	rcPortIndex rcSysMacFlapLimitTime rcSysMacFlapLimitCount	A rcLoopDetectMacDiscard trap signifies that a loop has been Detected on a port and the MAC address will be discarded on all ports in that VLAN.
1.3.6.1.4.1.2272.1.21.0.68	rcnAutoRecoverPort	rcPortIndex	A rcnAutoRecoverPort trap signifies that autorecovery has reenabled a port disabled by link flap or CP Limit.
1.3.6.1.4.1.2272.1.21.0.69	rcnAutoRecoverLoopDetectedPort	rcVlanNewLoopDetectedAction	A rcnAutoRecoverPort trap signifies that autorecovery has cleared the action taken on a port by loop detect.
1.3.6.1.4.1.2272.1.21.0.80	rcnVlacpPortDown	rcPortIndex	A rcnVlacpPortDown trap signifies that VLACP is down on the port specified.
1.3.6.1.4.1.2272.1.21.0.81	rcnVlacpPortUp	rcPortIndex	A rcnVlacpPortUp trap signifies that VLACP is up on the port specified.
1.3.6.1.4.1.2272.1.21.0.83	rcnEapMacIntrusion	rcSysIpAddr rcRadiusPaePortNumber rcRadiusEapLastAuthMac rcRadiusEapLastRejMac	A rcnEapMacIntrusion trap signifies that an EAP MAC intrusion has occurred on this port.

OID	Notification type	Objects	Description
1.3.6.1.4.1.2272.1.21.0.110	rcnMaxRouteWarnClear	rcVrfName	A rcnMaxRouteWarnClear trap signifies that the number of routes in the routing table of the virtual router has dropped below the warning threshold.
1.3.6.1.4.1.2272.1.21.0.111	rcnMaxRouteWarnSet	rcVrfName	A rcnMaxRouteWarnSet trap signifies that the virtual router routing table is reaching its maximum size. Take action to prevent this.
1.3.6.1.4.1.2272.1.21.0.112	rcnMaxRouteDropClear	rcVrfName	A rcnMaxRouteDropClear trap signifies that the virtual router routing table is no longer dropping new routes as it is below the maximum size.
1.3.6.1.4.1.2272.1.21.0.113	rcnMaxRouteDropSet	rcVrfName	A rcnMaxRouteDropSet trap signifies that the virtual router routing table has reached the maximum size, and is now dropping all new nonstatic routes.
1.3.6.1.4.1.2272.1.21.0.117	rcnMstpNewCistRoot	rcStgBridgeAddress	A rcMstpNewCistRoot trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the Multiple Spanning Tree Protocol has declared the device to be the new root of the common internal spanning tree.
1.3.6.1.4.1.2272.1.21.0.118	rcnMstpNewMstiRoot	rcStgBridgeAddress rcStgId	A rcMstpNewMstiRoot trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the Multiple Spanning Tree Protocol has declared the device to be the new root of the spanning tree instance.
1.3.6.1.4.1.2272.1.21.0.119	rcnMstpNewCistRegionalRoot	rcStgBridgeAddress	A rcMstpNewCistRegionalRoot trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the Multiple Spanning Tree Protocol has declared the device to be the new regional root of the common internal spanning tree.

OID	Notification type	Objects	Description
1.3.6.1.4.1.2272.1.21.0.120	rcnRstpNewRoot	rcStgBridgeAddress	A rcRstpNewRoot trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the Rapid Spanning Tree Protocol has declared the device to be the new root of the spanning tree.
1.3.6.1.4.1.2272.1.21.0.124	rcnRsmltEdgePeerModified	rcVlanId	A rcnRsmltEdgePeerModified trap signifies that the RSMLT peer address is different from that of the stored address. You must save the configuration if EdgeSupport has to use this information on the next restart.
1.3.6.1.4.1.2272.1.21.0.167	rcnChasPowerSupplyNoRedundancy	—	A rcnChasPowerSupplyNoRedundancy trap signifies that the chassis is running on power supply without redundancy.
1.3.6.1.4.1.2272.1.21.0.168	rcnChasPowerSupplyRedundancy	—	A rcnChasPowerSupplyRedundancy trap signifies that the chassis is running on power supply with redundancy.
1.3.6.1.4.1.2272.1.21.0.171	rcnLicenseTrialPeriodExpired	—	A rcnLicenseTrialPeriodExpired trap signifies that the Trial Period License has expired.
1.3.6.1.4.1.2272.1.21.0.172	rcnLicenseTrialPeriodExpiry	rcSysLicenseTrialDaysLeft	A rcnLicenseTrialPeriodExpiry trap signifies the time remaining, in days, before the License Trial Period expires.
1.3.6.1.4.1.2272.1.21.0.173	rcnVrfUp	rcVrfName rcVrfOperStatus	This notification is generated when the operational status of the specified VRF is toggled from down to up.
1.3.6.1.4.1.2272.1.21.0.174	rcnVrfDown	rcVrfName rcVrfOperStatus	This notification is generated when the operational status of the specified VRF is toggled from up to down.
1.3.6.1.4.1.2272.1.21.0.175	rcnMrouteIngressThresholdExceeded	rcIpResourceUsageGlobalIngressRecInUse rcIpResourceUsag	This notification is generated when the number of mroute ingress records exceeds the ingress threshold.

OID	Notification type	Objects	Description
		eGlobalIngressThreshold	
1.3.6.1.4.1.2272.1.21.0.176	rcnMrouteEgressThresholdExceeded	rcIpResourceUsageGlobalEgressRecInUse rcIpResourceUsageGlobalEgressThreshold	This notification is generated when the number of mroute egress records exceeds the egress threshold.
1.3.6.1.4.1.2272.1.21.0.177	rcnRemoteMirroringStatus	rcPortRemoteMirroringIndex rcPortRemoteMirroringEnable rcPortRemoteMirroringMode	A rcRemoteMirroringStatus trap signifies whether the remote mirroring is enabled or disabled on a particular port.
1.3.6.1.4.1.2272.1.21.0.185	rcnChasPowerSupplyRunningLow	—	A rcnChasPowerSupplyRunningLow trap signifies that the chassis is running on low power supply.
1.3.6.1.4.1.2272.1.21.0.196	rcnChasFanCoolingLow	rcChasFanOperStatus rcChasFanType rcErrorLevel rcErrorText	An rcnaChasFanCoolingLow trap signifies that the chassis is running on low fan cooling.
1.3.6.1.4.1.2272.1.21.0.285	rcnaSshSessionLogout	rcSshGlobalHostIpAddr	An rcnaSshSessionLogout trap signifies that there is an SSH session logout.
1.3.6.1.4.1.2272.1.21.0.286	rcnaSshUnauthorizedAccess	rcSshGlobalHostIpAddr	An rcnaSshUnauthorizedAccess trap signifies that an unauthorized access has occurred. It is deprecated by rcnaSshUnauthorizedAccess.
1.3.6.1.4.1.2272.1.21.0.287	rcnaAuthenticationSuccess	rcLoginUserName, rcLoginHostIpAddress	An rcnaAuthenticationSuccess trap signifies that a login is successful. The Trap includes the login username and the host IP address. It is deprecated by rcnaAuthenticationSuccess.
1.3.6.1.4.1.2272.1.21.0.288	rcnaSshSessionLogin	rcSshGlobalHostIpAddr	An rcnaSshSessionLogin trap signifies that there is an SSH session login.

Table 23: 1.3.6.1.4.1.2272.1.64.1.x.xx series

OID	Notification type	Objects	Description
1.3.6.1.4.1.2272.1.64.1.0.1	rcnSlppPortDownEvent	rcSlppPortSlppEnable rcSlppVlanSlppEnable rcSlppIncomingVlanId rcSlppSrcMacAddress	A port down event that has occurred due to SLPP. The user is notified of the expected VLAN ID along with the VLAN ID and source MAC address of the packet coming in on the port identified. The first two objects can be used to lookup instance info for port ID and VLAN ID.

Table 24: 1.3.6.1.4.1.2272.1.206.x.x.x series

OID	Notification type	Objects	Description
1.3.6.1.4.1.2272.1.206.1.0.1	rcVrrpTmpTrapNewMaster	rcVrrpTmpOperationsMasterIpAddr rcVrrpTmpNewMasterReason	This notification is generated when VRRP transitions to the master.
1.3.6.1.4.1.2272.1.206.2.2.1	rcVrrpExtTrapStateTransition	ifIndex rcVrrpExtTrapStateTransitionType rcVrrpExtTrapStateTransitionCause rcVrrpExtOperationsVrId rcVrrpTmpOperationsPrimaryIpAddr rcVrrpTmpOperationsMasterIpAddr	This notification is generated when a transition happens in the state of VRRP, for instance, a transition from master to backup when shutdown is received.

Standard traps

The following table describes standard traps that the Virtual Services Platform 9000 can generate.

Table 25: Standard traps

OID	Notification type	Objects	Description
1.3.6.1.2.1.14 .16.2.1	ospfVirtIfState Change	ospfRouterId ospfVirtIfAreaId ospfVirtIfNeighbor ospfVirtIfState	An ospfIfStateChange trap signifies that there has been a change in the state of an OSPF virtual interface. This trap is generated after the interface state regresses, for example, goes from Point-to-Point to Down, or progresses to a terminal state, for example, Point-to-Point.
1.3.6.1.2.1.14 .16.2.2	ospfNbrStateC hange	ospfRouterId ospfNbrIpAddr ospfNbrAddressLe ssIndex ospfNbrRtrId ospfNbrStat	An ospfNbrStateChange trap signifies a change in the state of a non-virtual OSPF neighbor. This trap is generated after the neighbor state regresses, for example, goes from Attempt or Full to 1-Way or Down. or progresses to a terminal state, for example, 2-Way or Full. When a neighbor transitions from or to Full on non-broadcast multiple access and broadcast networks, the trap is generated by the designated router. A designated router transitioning to Down will be noted by ospfIfStateChange
1.3.6.1.2.1.14 .16.2.3	ospfVirtNbrSta teChange	ospfRouterId ospfVirtNbrArea ospfVirtNbrRtrId ospfVirtNbrState	An ospfIfStateChange trap signifies a change in the state of an OSPF virtual neighbor. This trap is generated after the neighbor state regresses, for example, goes from Attempt or Full to 1-Way or Down, or progresses to a terminal state, for example, Full.
1.3.6.1.2.1.14 .16.2.4	ospfIfConfigEr ror	ospfRouterId ospfIfIpAddress ospfAddressLessIf ospfPacketSrc ospfConfigErrorTyp e ospfPacketType	An ospfIfConfigError trap signifies that a packet has been received on a nonvirtual interface from a router whose configuration parameters conflict with the configuration parameters of this router. Note that the event optionMismatch should cause a trap only if it prevents an adjacency from forming.
1.3.6.1.2.1.14 .16.2.5	ospfVirtIfConfi gError	ospfRouterId ospfVirtIfAreaId	An ospfConfigError trap signifies that a packet has been received on

OID	Notification type	Objects	Description
		ospfVirtIfNeighbor ospfConfigErrorType ospfPacketType	a virtual interface from a router whose configuration parameters conflict with the configuration parameters of this router. Note that the event optionMismatch should cause a trap only if it prevents an adjacency from forming.
1.3.6.1.2.1.14.16.2.6	ospfIfAuthFailure	ospfRouterId ospfIfIpAddress ospfAddressLessIf ospfPacketSrc ospfConfigErrorType authTypeMismatch authFailure ospfPacketType	An ospfIfAuthFailure trap signifies that a packet has been received on a nonvirtual interface from a router whose authentication key or authentication type conflicts with the authentication key or authentication type of this router.
1.3.6.1.2.1.14.16.2.7	ospfVirtIfAuthFailure	ospfRouterId ospfVirtIfAreaId ospfVirtIfNeighbor ospfConfigErrorType authTypeMismatch authFailure ospfPacketType	An ospfVirtIfAuthFailure trap signifies that a packet has been received on a virtual interface from a router whose authentication key or authentication type conflicts with the authentication key or authentication type of this router.
1.3.6.1.2.1.14.16.2.16	ospfIfStateChange	ospfRouterId ospfIfIpAddress ospfAddressLessIf ospfIfState	An ospfIfStateChange trap signifies a change in the state of a nonvirtual OSPF interface. This trap is generated after the interface state regresses, for example, goes from Dr to Down, or progresses to a terminal state, for example, Point-to-Point, DR Other, Dr, or Backup.
1.3.6.1.2.1.16.0.1	risingAlarm	alarmIndex alarmVariable alarmSampleType alarmValue alarmRisingThreshold	The SNMP trap that is generated after an alarm entry crosses the rising threshold and generates an event that is configured to send SNMP traps. TRAP TYPE ENTERPRISE rmon
1.3.6.1.2.1.16.0.2	fallingAlarm	alarmIndex alarmVariable alarmSampleType alarmValue alarmFallingThreshold	The SNMP trap that is generated after an alarm entry crosses the falling threshold and generates an event that is configured to send SNMP traps. TRAP TYPE ENTERPRISE rmon

OID	Notification type	Objects	Description
1.3.6.1.2.1.46 .1.3.0.3	vrrpTrapState Transition	ifIndex vrrpTrapStateTrans itionType vrrpTrapStateTrans itionCause vrrpOperVrId vrrpOperIpAddr ipAdEntAddr	A vrrpTrapStateTransition trap signifies a state transition has occurred on a particular VRRP interface. Implementation of this trap is optional. vrrpOperIpAddr contains the IP address of the VRRP interface while ipAdEntAddr contains the IP address assigned to physical interface.
1.3.6.1.2.1.68 .0.1	vrrpTrapNewM aster	vrrpOperMasterIpA ddr	The newMaster trap indicates that the sending agent has transitioned to Master state.
1.3.6.1.2.1.68 .0.2	vrrpTrapAuthF ailure	vrrpTrapPacketSrc vrrpTrapAuthErrorT ype	A vrrpAuthFailure trap signifies that a packet has been received from a router whose authentication key or authentication type conflicts with the authentication key or authentication type of this router.
1.3.6.1.2.1.80 .0.1	pingProbeFail ed	pingCtlTargetAddre ssType pingCtlTargetAddre ss pingResultsOperSt atus pingResultsIpTarge tAddressType pingResultsIpTarge tAddress pingResultsMinRtt pingResultsMaxRtt pingResultsAverag eRtt pingResultsProbeR esponse pingResultsSentPr obes pingResultsRttSum OfSquares pingResultsLastGo odProbe	This trap is generated after a probe failure is detected when the corresponding pingCtlTrapGeneration object is configured to probeFailure(0) subject to the value of pingCtlTrapProbeFailureFilter. The object pingCtlTrapProbeFailureFilter can specify the number of successive probe failures required before this notification can be generated.
1.3.6.1.2.1.80 .0.2	pingTestFailed	pingCtlTargetAddre ssType pingCtlTargetAddre ss pingResultsOperSt atus	This trap is generated after a ping test fails when the corresponding pingCtlTrapGeneration object is configured to testFailure(1). In this instance pingCtlTrapTestFailureFilter

OID	Notification type	Objects	Description
		pingResultsIpTargetAddressType pingResultsIpTargetAddress pingResultsMinRtt pingResultsMaxRtt pingResultsAverageRtt pingResultsProbeResponses pingResultsSentProbes pingResultsRttSumOfSquares pingResultsLastGoodProbe	specifies the number of probes in a test required to fail to consider the test as failed.
1.3.6.1.2.1.80.0.3	pingTestCompleted	pingCtlTargetAddressType pingCtlTargetAddress pingResultsOperStatus pingResultsIpTargetAddressType pingResultsIpTargetAddress pingResultsMinRtt pingResultsMaxRtt pingResultsAverageRtt pingResultsProbeResponses pingResultsSentProbes pingResultsRttSumOfSquares pingResultsLastGoodProbe	This trap is generated at the completion of a ping test when the corresponding pingCtlTrapGeneration object is configured to testCompletion(4).
1.3.6.1.2.1.81.0.1	traceRoutePathChange	traceRouteCtlTargetAddressType traceRouteCtlTargetAddress traceRouteResultsIpTgtAddrType traceRouteResultsIpTgtAddr	This trap is generated after the path to a target changes.

OID	Notification type	Objects	Description
1.3.6.1.2.1.81.0.2	traceRouteTestFailed	traceRouteCtlTargetAddressType traceRouteCtlTargetAddress traceRouteResultsIpTgtAddrType traceRouteResultsIpTgtAddr	This trap is generated is traceroute cannot determine the path to a target (traceRouteNotifications 2).
1.3.6.1.2.1.81.0.3	traceRouteTestCompleted	traceRouteCtlTargetAddressType traceRouteCtlTargetAddress traceRouteResultsIpTgtAddrType traceRouteResultsIpTgtAddr	This trap is generated after the path to a target is determined.
1.3.6.1.6.3.1.1.5.1	coldStart	—	A coldStart trap signifies that the SNMPv2 entity, acting in an agent role, is reinitializing and that its configuration may have been altered.
1.3.6.1.6.3.1.1.5.2	warmStart	—	A warmStart trap signifies that the SNMPv2 entity, acting in an agent role, is reinitializing such that its configuration is unaltered.
1.3.6.1.6.3.1.1.5.3	linkDown	—	A linkDown trap signifies that the sending protocol entity recognizes a failure in one of the communication links represented in the agent configuration. TRAP-TYPE ENTERPRISE snmp
1.3.6.1.6.3.1.1.5.4	linkUp	—	A linkUp trap signifies that the sending protocol entity recognizes that one of the communication links represented in the agent configuration has come up. TRAP-TYPE ENTERPRISE snmp
1.3.6.1.6.3.1.1.5.5	authenticationFailure	—	—

Chapter 14: Hardware troubleshooting

The following sections provide troubleshooting information for more common hardware problems.

Troubleshooting module failure

About this task

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

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 module type that is different from the module type removed in Step 1, and then wait for the replacement module to initialize.

   Important:

   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 will see one of two recovery menus. The first is illustrated in the following example.

```
  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 second menu is illustrated in the following example.

```
  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 USB and 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 VSP 9000 stop command will 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 USB or 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.

Procedure

1. 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.
2. 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

```
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 VSP 9000 automatically recognizes them. The devices are accessible within seconds after insertion.

After you insert the external Compact Flash, you should enable logging to the external Compact Flash by using 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

Before you begin

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

About this task

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. Check the file system:
   `ls /usb/`
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 and then reinsert the device.
4. Check the device for errors:
   `dos-chkdsk /usb`
5. If errors are detected, then you can reformat the device:
   `dos-format /usb`

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

Example

```
VSP-9012:> enable

VSP-9012:# 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/

VSP-9012:# usb-stop

It is now safe to remove the USB device.
```

```
VSP-9012:# 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
```

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

Troubleshooting 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. If you cannot write to a USB device, perform this procedure.

Procedure

Verify that the write-protect switch is in the write position.

Troubleshooting USB writing problems

About this task

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

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

Procedure

1. If you try to create file number 257, the device gives the no more space error.
2. You can delete files from the device.

   or
3. Reformat the device. See Troubleshooting USB Viewing Problems for more information: Troubleshooting USB viewing problems on page 115.

Troubleshooting external Compact Flash viewing problems

Before you begin

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

About this task

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. Check the file system:
   ```
   ls /extflash/
   ```
2. 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.
3. Remove and then reinsert the device.
4. Check the device for errors:
   ```
   dos-chkdsk /extflash
   ```
5. If errors are detected, then you can reformat the Compact Flash:
   ```
   dos-format /extflash
   ```

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

Example

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

Using trace to diagnose hardware problems

Before you begin

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

About this task

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

Procedure

1. Begin the trace operation:
   `line-card {<3–12>|SF1|SF2|SF3|SF4|SF5|SF6} trace level [{67–92} {0–4}]`
2. Search the trace for a specific string value:
   `line-card {<3–12>|SF1|SF2|SF3|SF4|SF5|SF6} trace grep [WORD<0–1024>]`

Example

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

Begin the trace operation:

```
VSP-9012:#line-card SF1 trace level 67 1
```

Search the trace for a specific string value:

```
VSP-9012:#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 26: Variable definitions

Variable	Value
{<3–12>|SF1|SF2|SF3|SF4|SF5|SF6}	Specifies the slot number for the interface module or Switch Fabric module.
{67–92} {0–4}	Starts the trace by specifying the module ID and level. <*67-92*> 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.
WORD<0–1024>	Performs a string search in the trace.

Chapter 15: Software troubleshooting

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

Failure to read 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

```
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 *Avaya Virtual Services Platform 9000 Upgrades and Patches*, NN46250-400.

Chapter 16: Software download

This section describes where to download software or documentation.

Downloading the software

Download the new software to upgrade the Avaya Virtual Services Platform 9000.

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.

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 **Downloads**.
5. Click the release name.
6. Click the **Downloads** tab.
7. Download the required software release.
8. Use an FTP client in binary mode to transfer the file to the Virtual Services Platform 9000, or transfer it using an external USB device.

Downloading Avaya Virtual Services Platform 9000 documentation

About this task

Download documentation from the Avaya Web site 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.

Chapter 17: Software troubleshooting tool configuration using the ACLI

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

Using ACLI for troubleshooting

Before you begin

- You must log on to at least Privileged EXEC mode in ACLI to use the **`show running-config`** and **`show interfaces`** commands in this procedure.

About this task

You can use ACLI to provide diagnostic information.

Procedure

1. Disable scrolling of the output display:
   ```
   terminal more disable
   ```
2. View configuration file information:
   ```
   more WORD<1–99>
   ```
3. Capture the output for the following command after you observe a problem with the device:
   ```
   show running-config [verbose] [module <cli|sys|web|rmon|
   vlan|port|qos|mlt|stg|ip|diag|radius|ntp|lacp|naap|cluster|
   boot|filter|ipv6|slpp|nsna|vsptalk>]
   ```
4. Capture the output for the following command after you observe a problem with the device:
   ```
   show tech
   ```
5. Capture the output for the following command after you observe a problem with the device:
   ```
   show interfaces gigabitEthernet statistics <bridging {slot/
   port[-slot/port][,...]}|dhcp-relay {slot/port[-slot/port]
   [,...]} [vrf name 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][,...]}>
```

6. Capture the output for the following command after you observe a problem with the device:

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

Example

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

Capture the output for the following command after you observer a problem with the device:

```
VSP-9012:1#show running-config module cli
Preparing to Display Configuration...
#
# Wed Aug 03 12:03:42 2011 UTC
# box type             : VSP-9012
# software version     : 0.0.0.0_B340 (PRIVATE)
# cli mode             : ACLI
#


#ASIC Info :
#Slot #1:
#       Module: 9080CP
#       OXATE CPLD: 10032310
#       OXIDE FPGA: 10040918
#       CATSKILL FPGA: 10052013
#       QE version: QE2000_A0

#Slot #4:
#       Module: 9048GT
#       K2 FPGA: 11052520
#       IODATEDC CPLD: 09041015
#       IODATEBB CPLD: 09041016
#       PIM48TX CPLD: 09050110
#       LED48TX CPLD0: 09041016

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

Capture the output for the following command after you observe a problem with the device:

```
VSP-9012:1#show tech

Sys Info:
---------------


General Info :

        SysDescr     : VSP-9012 (3.2.0.0) (DEV)
        SysName      : VSP-9012
        SysUpTime    : 2 day(s), 09:30:18
        SysContact   : http://support.avaya.com/
        SysLocation  : 211 Mt. Airy Road,Basking Ridge,NJ 07920

Chassis Info:

        Chassis            : 9012
        Serial#            :
```

```
        H/W Revision       :
        H/W Config         : 01
        NumSlots           : 12
        NumPorts           : 328
        BaseMacAddr        : 00:24:7f:a1:70:00
        MacAddrCapacity    : 4096

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

Capture the output for the following command after you observe a problem with the device:

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

================================================================================
                             Port Stats Interface
================================================================================
PORT     IN                  OUT                  IN                  OUT
NUM    OCTETS               OCTETS               PACKET               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 output for the following command after you observe a problem with the device:

```
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 27: Variable definitions

Variable	Value
WORD<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: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 28: Variable definitions

Variable	Value
module <cli|sys|web|rmon|vlan|port|qos|mlt|stg|ip|diag|radius|ntp|lacp|naap|cluster|boot|filter|ipv6|slpp|nsna|vsptalk>	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 statistics`

Variable	Value
bridging *{slot/port[-slot/port][,...]}*	Displays ports bridging statistics.
dhcp-relay *{slot/port[-slot/port][,...]}* [vrf name *WORD,0–16*>][vrfids *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 29: Variable definitions

Variable	Value
collision	Displays port collision error information.
ospf	Displays port Open Shortest Path First (OSPF) error information.
verbose	Dsiplays 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

Dump software record information:

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

Example

```
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 30: 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.

Variable	Value
<0-3>	Specifies the verbosity from 0–3. Higher numbers specify more verbosity.

Using trace to diagnose 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.

- You must log on to at least the Privileged EXEC mode in ACLI to clear or save trace information.

About this task

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

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

Procedure

1. Clear the trace:
   `clear trace`
2. Identify the module ID for which you want to use the trace tool:
   `show trace modid-list`
3. Begin the trace operation:
   `trace level [<0-178>] [<0-4>]`
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
5. Stop tracing:
   `trace shutdown`
6. View the trace results:

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

7. 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.
8. 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.
9. 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
19  - WDT
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 31: Variable definitions

Variable	Value
grep *[WORD<0-128>]*	Performs a comparison of trace messages.
level *[<0-178>] [<0-4>]*	Starts the trace by specifying the module ID and level. • <0-178> 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 32: Variable definitions

Variable	Value
file *WORD<1–99>*	Specifies the file name in one of the following formats: • a.b.c.d:<file> • /intflash/<file> • /extflash/<file> • /usb/<file>

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.

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

About this task

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

Procedure

1. Activate or deactivate the trace for the IPv6 base:
   `trace ipv6 base <disable|enable> <all|debug|error|icmp|info|ipclient|nbr|pkt|warn>`
2. Activate or deactivate the trace for IPv6 forwarding:
   `trace ipv6 forwarding <disable|enable> <all|debug|error|info|pkt|warn>`
3. Activate or deactivate the trace for IPv6 neighbor discovery:
   `trace ipv6 nd <disable|enable> <all|debug|error|info|nbr|pkt|redirect|warn>`
4. 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>`
5. Activate or deactivate the trace for IPv6 transport:
   `trace ipv6 transport <disable|enable> <all|common|tcp|udp>`

Example

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

Activate the trace for all the IPv6 base categories:

```
VSP-9012:1#trace ipv6 base enable all
```

Activate the trace for all the IPv6 forwarding categories:

```
VSP-9012:1#trace ipv6 forwarding enable all
```

Activate the trace for all the IPv6 neighbor discovery categories:

```
VSP-9012:1#trace ipv6 nd enable all
```

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

```
VSP-9012:1#trace ipv6 rtm enable all
```

Activate the trace for all the IPv6 transport caterories:

```
VSP-9012:1#trace ipv6 transport enable all
```

Variable definitions

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

Table 33: 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.
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.

Using autotrace to diagnose problems

About this task

Use autotrace to automatically perform the trace function after a parameter reaches a threshold.

For example, if the CPU fluctuates and you cannot access the device to perform a CP trace, use autotrace to automatically perform this function. Autotrace monitors CPU utilization. After the utilization reaches the threshold and is sustained for the configured amount of time, the device performs a CP trace, and then saves the result to the external flash.

Procedure

1. Configure the module and verbosity:
   `trace auto module add <0-107> <0-4>`
2. Configure the high threshold for a module:
   `trace auto high-percentage <60-100>`
3. Configure the sustained time above the high threshold to trigger a trace:
   `trace auto high-track-duration <3-10>`
4. Configure the low threshold for a module:
   `trace auto low-percentage <50-90>`
5. Configure the sustained time below the low threshold to trigger a trace:
   `trace auto low-track-duration <3-10>`
6. Enable automatic tracing:
   `trace auto enable`

Example

Configure the module and verbosity:

```
VSP-9012:1> trace auto module add 20 2
```

Configure the high threshold for a module:

```
VSP-9012:1> trace auto high-percentage 70
```

Configure the sustained time above the high threshold to trigger a trace:

```
VSP-9012:1> trace auto high-track-duration 5
```

Configure the low threshold for a module:

```
VSP-9012:1> trace auto low-percentage 50
```

Configure the sustained time below the low threshold to trigger a trace:

```
VSP-9012:1> trace auto low-track-duration 3
```

Enable automatic tracing:

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

Variable definitions

Use the data in the following table to use the `trace auto` command.

Table 34: Variable definitions

Variable	Value
disable	Disables the auto-trace function.
enable	Enables the auto-trace function.
high-percentage <60-100>	Specifies the high-percentage threshold for a module. The range is 60–100%.
high-track-duration <3-10>	Specifies, in seconds, the amount of time that the activity must be sustained to trigger the trace. The range is 3–10 seconds.
low-percentage <50-90>	Specifies the low-percentage threshold for a module. The range is 50–90%.
low-track-duration <3-10>	Specifies, in seconds, the amount of time that the activity must be sustained to trigger the trace. The range is 3–10 seconds.
module add <0-107> <0-4>	Configures the trace auto-enable function by specifying the module ID and level. • *<0-107>* specifies the module ID from 0–107. • *<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.
module remove <0-107>	Removes a module ID from the auto-trace instance.

Configuring port mirroring

Before you begin

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

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. 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][,...]}}`
2. Configure the mode:
   `mirror-by-port <1-479> mode <both|rx|tx>`
3. Enable the mirroring instance:
   `mirror-by-port <1-479> enable`
4. 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` .
5. Verify the configuration:
   `show mirror-by-port`

Example

```
VSP-9012:1>enable

VSP-9012:1#configure terminal
```

Create the port mirroring instance. Traffic passing port 5/15 is mirrored 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   DSCP TTL
                                                        VLAN-ID
   --------------------------------------------------------------------------------
   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   DSCP TTL
                                                        VLAN-ID
   --------------------------------------------------------------------------------
   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 35: 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][,...]}}	Creates a new mirror-by-port table entry. • 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][,...]}	Modifies the mirrored port. Before you can modify an existing entry, you must disable the entry: `no mirror-by-port <1-479> enable` .
monitor-ip <*1-479*> {A.B.C.D} [dscp <*0-63*>] [ttl <*2-255*>]	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. 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.
monitor-mlt <1-479> <1-512>	Modifies the monitoring MLT; <*1-512*> specifies the mirroring MLT ID.

Variable	Value
	Before you can modify an existing entry, you must disable the entry: `no mirror-by-port <1-479> enable` .
monitor-port <1-479> {slot/port[-slot/port][,...]}	Modifies the monitoring ports. Before you can modify an existing entry, you must disable the entry: `no mirror-by-port <1-479> enable` .
monitor-vlan <1-479> <1-4084>	Modifies the monitoring VLAN. Before you can modify an existing entry, you must disable the entry: `no mirror-by-port <1-479> enable` .
mode <both|rx|tx>	Configures the mirroring mode. The default is rx. • both mirrors both egress and ingress packets. • rx mirrors ingress packets. • tx mirrors egress packets.
remote-mirror-vlan-id <1-4084>	Configures the remote mirror VLAN ID.

Configuring global mirroring actions with an ACL

Before you begin

- The ACL exists.
- You must log on to Global Configuration mode in ACLI.

About this task

Configure the global action to mirror packets that match an access control list (ACL).

Procedure

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 36: 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: • 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.
- You must log on to Global Configuration mode in ACLI.

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. 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>
```

2. Ensure the configuration is correct:
   ```
   show filter acl action [<1-2048>] [<1-2000>]
   ```

Example

```
VSP-9012:1>enable

VSP-9012:1#configure terminal

VSP-9012:1(config)#filter acl ace action 901 1 permit monitor-dst-ip
10.10.10.1 dscp 10 ttl 155
```

Variable definitions

Use the data in the following table to use the `filter acl ace action` command.

Table 37: Variable definitions

Variable	Value
1-2048	Specifies the ACL ID from 1–2048
1-2000	Specifies the ACE ID from 1–2000.
monitor-dst-ip {A.B.C.D} [dscp <*0–63*>] [ttl <*2–255*>]	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.
monitor-dst-mlt <*1–512*>	Configures mirroring to a destination MLT group.
monitor-dst-ports {slot/port[-slot/port][,...]}	Configures mirroring to a destination port or ports.
monitor-dst-vlan <*1-4084*>	Configures mirroring to a destination VLAN.
{permit|deny}	Configures the action mode for security ACEs. The default value is permit.

Configuring Layer 2 remote mirroring

Before you begin

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

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 136.

Procedure

1. Configure the mode for remote mirroring:
   ```
   remote-mirroring mode <source|termination>
   ```
2. Configure the destination MAC for remote mirroring:
   ```
   remote-mirroring dstMac <0x00:0x00:0x00:0x00:0x00:0x00>
   [ether-type <0x00-0xffff>] [vlan-id <1-4084>]
   ```
3. Configure the source MAC for remote mirroring:
   ```
   remote-mirroring srcMac <0x00:0x00:0x00:0x00:0x00:0x00>
   [ether-type <0x00-0xffff>] [vlan-id <1-4094>]
   ```
4. 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>]
   ```
5. Enable remote mirroring:
   ```
   remote-mirroring enable
   ```
6. 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 38: Variable definitions

Variable	Value
dstMac *<0x00:0x00:0x00:0x00:0x00:0x00>*	Configures the destination MAC address for use in the remote mirroring encapsulation header. The mirrored packet is sent to this MAC address. Only remote mirroring source (RMS) ports use the DstMac.

Variable	Value
	For remote mirroring termination (RMT) ports, one of the unused MAC addresses from the switch port MAC address range is used. This MAC address is saved in the configuration file.
enable	Enables remote mirroring on the port. When remote mirroring is enabled, the following events occur: • A static entry for the DstMac is added 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.
ether-type <*0x00-0xffff*>	Specifies the Ethertype of the remote mirrored packet. The default value is 0x8103.
mode <*source*|*termination*>	Specifies whether the port is an RMT (mode is termination) or an RMS (mode is source).
{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).
srcMac <*0x00:0x00:0x00:0x00:0x00:0x00*>	Configures the source MAC address for use in the remote mirroring encapsulation header. The mirrored packet is sent from the RMS port, and the source MAC parameter in the header is derived from this address. The source MAC address of the encapsulated frame contains the first 45 bits of this MAC address. The three least significant bits are derived 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

- You must log on to at least Privileged EXEC mode in ACLI.
- 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. Access Privileged EXEC mode:
   ```
   enable
   ```
3. Access the secondary CPU:
   ```
   peer telnet
   ```

Example

```
VSP-9012:1> enable
VSP-9012:1# peer telnet
```

Configuring PCAP global parameters

Before you begin

- The secondary CP module is active.
- If you save to external flash, a Compact Flash (CF) card is present.
- You must log on to Global Configuration mode in ACLI.

About this task

Configure PCAP globally to define how PCAP operates on the Avaya Virtual Services Platform 9000.

Procedure

1. Enable PCAP:
   `pcap enable`
2. Configure optional parameters as required.
3. 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 39: Variable definitions

Variable	Value
auto-save [file-name WORD<1-40>] [network {A.B.C.D}|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. file-name WORD<1-40> is the name of the file where captured frames are saved. network configures the save device to network. <A.B.C.D> 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. A log message is generated after the buffer wraps.
enable	Enables PCAP globally. The default is disabled. To disable PCAP, use the `no pcap enable` 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. If this variable is false, the present file is not overwritten on the Compact Flash card.

Enabling PCAP on a port

Before you begin

- If required, IP filters exist.
- If required, ACLs with a global action of mirror exist.
- You must log on to Interface Configuration mode in ACLI.

About this task

Configure PCAP on a port to capture packets on that port.

Procedure

1. Enable PCAP on ports:
   `pcap enable [mode {both|rx|tx}]`
2. Verify the PCAP configuration:
   `show pcap port`

Example

```
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 40: 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 only packets which match the filter criteria are captured.

Configuring PCAP capture filters

Before you begin

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

About this task

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.

Procedure

1. Create a capture filter:
   ```
   pcap capture-filter <1-1000>
   ```
2. Configure the filter action:
   ```
   pcap capture-filter <1-1000> action <capture|drop|trigger-off|trigger-on>
   ```
3. Define the match parameters.
   Use the following variable definitions table to configure match parameters.
4. Enable the filter:
   ```
   pcap capture-filter <1-1000> enable
   ```
5. Ensure the configuration is correct:
   ```
   show pcap capture-filter [<1-1000>]
   ```

Example

```
VSP-9012:1> enable
VSP-9012:1# configure terminal
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 41: Variable definitions

Variable	Value
<1-1000>	Specifies the capture filter ID.
action <capture|drop|trigger-off|trigger-on>	Determines the action taken by the filter: • 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. The default is capture. **Important:** 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.
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 configurematch-zero , 0 is considered a valid value; otherwise, 0 is considered a disable value.

Variable	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 only the first few bytes are compared. <*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 Ethernet type 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 specifiednumber of packets. This variable is similar to the refresh-timer option;after it is invoked, the filter is disabled. This option is activeonly if you configure the action parameter to trigger-on. The defaultvalue is 0, which means this option is disabled.
pbits <0-7> [<0-7>] [match-zero]	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> [<0-255>]	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 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 only the first few bytes are compared. The default is 00:00:00:00:00:00, which means this option is disabled.

Variable	Value
	<*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 WORD<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. *WORD<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> WORD<0-50>	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 (*WORD*<*0-50*>). 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

About this task

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.

Procedure

1. Log on to the secondary CPU.
2. Use the following command:

```
show pcap dump
```

Copying captured packets to a remote machine

Before you begin

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

About this task

You can copy packets to a remote machine, or internal or external flash. If you use PCAP with autosave disabled, captured packets are stored in the secondary CPU DRAM buffer. To copy the packets to a file for later viewing, use the copy command. You can use this command in the pimary CPU.

Captured packets stored in the secondary CPU DRAM buffer use the name PCAP00.

Procedure

1. Copy packets:
   `copy WORD<1-99> WORD<1-99>`
2. Copy packets from DRAM:
   `copy PCAP00 WORD<1-99>`
3. 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 42: Variable definitions

Variable	Value
WORD<1-99> WORD<1-99>	Specifies USB, flash, or an IP host by IP address and specifies the PCAP file (.cap). Formats include • a.b.c.d:<file> • /intflash/<file> • /mnt/intflash/<file> • /extflash/<file> • /mnt/extflash/<file> • /usb/<file>

Resetting the PCAP DRAM buffer

Before you begin

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

About this task

Reset the DRAM buffer to clear the PCAP DRAM buffer and the PCAP counters.

Procedure

1. Log on to the secondary CPU.
2. Disable PCAP:
   ```
   no pcap enable
   ```
3. Reset the PCAP engine DRAM buffer:
   ```
   pcap reset-stat
   ```
4. Reenable PCAP:
   ```
   pcap enable
   ```

Clearing ARP information for an interface

Before you begin

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

About this task

Clear the Address Resolution Protocol (ARP) cache as part of ARP problem resolution procedures.

Procedure

Clear ARP information:

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

OR

```
clear ip arp interface vlan <1-4084>
```

Example

```
VSP-9012:> enable
VSP-9012:# clear ip arp interface gigabithethernet 4/12
```

Variable definitions

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

Table 43: Variable definitions

Variable	Value
1–4084	Specifies the VLAN ID
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).

Flushing routing, MAC, and ARP tables for an interface

Before you begin

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

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. Flush IP routing tables by port:
   ```
   action flushIp
   ```
2. Flush the MAC address tables:
   ```
   action flushMacFdb
   ```
3. Flush ARP tables:
   ```
   action flushArp
   ```
4. Flush all tables with one command:
   ```
   action flushAll
   ```
5. Exit to Global Configuration mode:
   ```
   exit
   ```
6. Clear a routing table:
   ```
   clear ip route gigabitethernet {slot/port}
   ```
   OR
   ```
   clear ip route vlan <1-4084>
   ```

Example

```
VSP-9012:1> enable
VSP-9012:1# configure terminal
VSP-9012:1(config)# interface gigabitethernet 4/10
VSP-9012:1(config-if)# action flushAll
VSP-9012:1(config-if)# exit
VSP-9012:1(config)# clear ip route gigabitethernet 4/10
```

Variable definitions

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

Table 44: Variable definitions

Variable	Value
1–4084	Specifies the VLAN ID.
{slot/port}	Specifies a port number.

Pinging an IP device

About this task

Ping a device to test the connection between the Avaya Virtual Services Platform 9000 and another network device. 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 may 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

Ping an IP network connection:

```
ping WORD<0–256> [-d] [-I <1–60>] [-s] [-t <1–120>] [count <1–9999>] [datasize <16–65487>] [interface WORD<1-256>|gigabitEthernet|mgmtEthernet|tunnel|vlan] [scopeid <1–9999>] [source WORD<1-256>] [vrf WORD<0–16>]
```

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.

Table 45: Variable definitions

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 issues (icmp packet too short or wrong icmp packet type).
datasize <16–65487>	Specifies the size of ping data sent in bytes (16–65487). The default is 16.
-I <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: • gigabitEthernet: {slot/port} gigabit ethernet 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 [-I] 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:x:x) address (string length 0–256).
vrf WORD<0–16>	Specifies the virtual router and forwarder (VRF) name from 1–16 characters.

Variable	Value
	Specify the MgmtRouter VRF if you need to run the ping operation through the management interface.

Running a traceroute test

Before you begin

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

About this task

Use traceroute to determine the route packets take through a network to a destination.

Ping and traceroute may 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

Run a traceroute test:

```
traceroute {A.B.C.D} [<1-1464>] [-m <1-255>] [-p <0-65535>] [-q <1-255>] [-v] [-w <1-255>] [source <WORD 1-256>] [vrf <WORD 0-16>]
```

Example

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

Run traceroute test, with a probe packet size of 200 and a max time to live of 60:

```
VSP-9012:1# traceroute 46.11.10.33 200 –m 60
```

Variable definitions

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

Table 46: Variable definitions

Variable	Value
{A.B.C.D}	Specifies the destination IP address.

Variable	Value
-m <*1-255*>	Specifies the maximum time-to-live (TTL) (1–255).
-p <*0-65535*>	Specifies the base UDP port number (0–65535).
-q <*1-255*>	Specifies the number of probes per TTL (1–255).
-v	Specifies verbose mode (detailed output).
-w <*1-255*>	Specifies the wait time for each probe (1–255).
<*1-1464*>	Specifies the size of the probe packet (1–1464).
source <*WORD 1-256*>	Specifies the source IP address.
vrf <*WORD 0-16*>	Specifies the VRF instance by VRF name.

Showing SNMP logs

Before you begin

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

About this task

Show the full SNMP logs. SNMP logs display the alarms and events that have been registered on the device.

Procedure

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 47: Variable definitions

Variable	Value
WORD<1-99>	This variable represents the log file to be opened and displayed. Use one of three formats: • /intflash/<file> • /extflash/<file> • /usb/<file>

Chapter 18: 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

About this task

For administrative and troubleshooting purposes, sometimes you must flush the routing tables. You can use EDM to flush the routing tables by VLAN or flush them by port. Perform this procedure to flush the IP routing table for a VLAN.

Procedure

1. In the navigation tree, open 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, all entries associated with the VLAN are flushed. You can also flush the Address Resolution Protocol (ARP) entries and IP routes for the VLAN.
5. Click **Apply**.

Flushing routing tables by port

About this task

For administrative and troubleshooting purposes, sometimes you must flush the routing tables. You can use EDM to flush the routing tables by VLAN or flush them by port. Perform this procedure to flush the IP routing table for a port.

Procedure

1. Select a port.
2. In the navigation tree, open the following folders: **Configuration** > **Edit** > **Port**.
3. Click **General**.
4. Click the Interface tab.
5. In the **Action** section, select **flushAll**.

   In a port context, all entries associated with the port are flushed. You can flush the ARP entries and IP routes for a port. After you flush a routing table, it is not automatically repopulated. The repopulation time delay depends on the routing protocols in use.
6. 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

1. In the navigation tree, open the following folders: **Configuration** > **Edit** > **Diagnostics**.
2. Click **General**.
3. Click the **Port Mirrors** tab.
4. Click **Insert**.
5. Configure mirroring as required.
6. To enable port mirroring for the instance, select the **Enable** check box.

7. 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).
Enable	Enables or disables this port mirroring instance. The default value is Enable.
Mode	Specifies the traffic direction of the packet being mirrored: • 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 are forwarded). Used to configure the mirroring port.
MirroringVlanId	Specifies the destination VLAN ID.
MirroringMltId	Specifies the destination multilink trunk ID.
RemoteMirrorVlanId	Specifies the virtual local area network 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 164.

Procedure

1. From the Physical Device View, select a port.
2. In the navigation tree, open 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	Enables or disables remote mirroring on the port. When remote mirroring termination (RMT) is enabled, the following things occur:

Name	Description
	• A static entry for the DstMac is added to the FDB. All packets that come with that remote mirroring dstmac are sent 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.
Mode	Specifies whether the port is an RMT or an RMS.
SrcMac	Specifies the source MAC address of the remote mirrored packet. The remote mirroring packet is sent with this source MAC address.
DstMac	Specifies the destination MAC address of the remote mirrored packet. Packets are bridged to this MAC address. Remote mirroring packets are sent to this MAC address.
EtherType	Specifies the Ethertype of the remote mirrored packet. The default value is 0x8103. Packets are sent with this Ethertype.
VlanIdList	If the port is a termination port, represents the filter lists VLAN in which the destination MAC address resides.

Configuring ACLs for mirroring

Before you begin

• The ACL exists.

About this task

Use the access control list (ACL) global action of mirroring to mirror packets for an access control entry (ACE) that matches a particular packet.

Procedure

1. In the navigation tree, open the following folders: **Configuration** > **Security** > **Data Path**.
2. Click **Advanced Filters (ACE/ACLs)**.
3. Click the **ACL** tab.
4. In the **GlobalAction** column, double-click a value, and then choose the desired mirror option.

5. Click **Apply**.

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, open 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.

Name	Description
AclId	Specifies the ACL ID.
AceId	Specifies a unique identifier and priority for the ACE.

Name	Description
Mode	Indicates the operating mode associated with this ACE. Valid options are deny and permit, with deny as the default.
MltIndex	Specifies whether to override the MLT-index picked by the MLT algorithm when a packet is sent out on MLT ports. Valid values range from 0 to 16, with 0 as the default. MLT index is not supported for multicast traffic, but for unicast traffic only.
RedirectNextHop	Redirects matching IP traffic to the next hop.
RedirectUnreach	Configures the desired behavior for redirected traffic when the specified next-hop is not reachable. The default value is deny.
IpfixState	Enables or disables the IPFIX action for the ACE.
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	Specifies the ports to which to mirror traffic.
DstVlanId	Specifies the VLAN to which to mirror traffic.
DstMltId	Specifies the Multilink Trunking (MLT) group to which to mirror traffic.
DstIp	Configures mirroring to a destination IP address for flow-based mirroring.
DstIpDscp	Optionally, configures the DSCP value. The default is 256 (disabled).
DstIpTtl	Optionally, configures the time-to-live value. The default TTL is 64.

Configuring PCAP globally

Before you begin

- The secondary CPU is active.
- If you save to external storage, a Compact Flash (CF) card is installed.

About this task

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.

Procedure

1. In the navigation tree, open 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).
BufferWrap	Enables buffer wrap-around after the buffer is full. When enabled, PCAP continues to capture packets, otherwise, packet capturing stops.
WrapAutoSaveFile	Select the WrapAutoSaveFile 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 WrapAutoSaveFile checkbox.

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

Before you begin

- If required, IP filters exist.
- If required, ACLs with a global action of mirror exist.

About this task

Configure PCAP on a port so that the port supports PCAP.

Procedure

1. In the Device Physical View tab, select a port.
2. In the navigation tree, open 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.
Mode	Configures the PCAP mode (tx, rx, or both). The default is rx mode.

Configuring PCAP filters

About this task

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.

Procedure

1. In the navigation tree, open 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.
Action	Determines the action taken by the filter:

Name	Description
	• **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. **Important:** 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.
SrcMac	Specifies the source MAC address to match.
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.
DstMac	Specifies the destination MAC address.
DstMacMask	Specifies the destination MAC address mask to specify an address range.
IsInverseDstMac	Specifies the destination MAC address inverse. If you select this variable, all MAC addresses other than the one specified are matched.
VlanId	Specifies the VLAN ID of the packet to match.
ToVlanId	Specifies the destination VLAN ID; use to specify a range.
IsInverseVlanId	Specifies the VLAN ID inverse. If you select this variable, all VLAN IDs other than the one specified are matched.
Pbit	Specifies the 802.1p-bit of the packet to match.
ToPbit	Specifies an 802.1p-bit range.

Name	Description
IsInversePbit	Specifies the p-bit inverse. If you select this variable, all p-bits other than the one specified are matched.
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.
EtherType	Specifies the EtherType of the packet to match.
ToEtherType	Specifies an EtherType range.
IsInverseEtherType	Specifies the EtherType inverse. If you select this variable, all EtherTypes other than the one specified are matched.
SrcIp	Specifies the source IP address of the packet to match.
ToSrcIp	Specifies a source IP address range.
IsInverseSrcIp	Specifies the source IP address inverse. If you select this variable, source IP addresses other than the one specified are matched.
DstIp	Specifies the destination IP address of the packet to match.
ToDstIp	Specifies the destination IP address range.
IsInverseDstIp	Specifies the destination IP address inverse. If you select this variable, all addresses other than the one specified are matched.
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.
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.
ProtocolType	Specifies the protocol of the packet to match.
ToProtocolType	Specifies a protocol type range.
IsInverseProtocolType	Specifies the protocol type inverse. If you select this variable, all protocols other than the one specified are matched.

Configuring advanced PCAP filters

About this task

Use advanced filters to match User Datagram Protocol (UDP) and TCP parameters, as well as to specify user-defined parameters.

Procedure

1. In the navigation tree, open 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.
ToUdpPort	Specifies a range of UDP ports.
IsInverseUdpPort	Indicates that all other values other than the specified range of UDP ports are matched.
TcpPort	Specifies the TCP port of the packet to match.
ToTcpPort	Specifies a range of TCP ports.
IsInverseTcpPort	Indicates that all other values other than the specified range of TCP ports are matched.
UserDefinedData	Specifies the user-defined data to match.
UserDefinedDataSize	Specifies the length of user-defined data.
UserDefinedOffset	Specifies the offset from which the match must start.

Name	Description
IsInverseUserDefined	Indicates that all data other than the specified user-defined data is matched.
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, open 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 and 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 16.
TimeOut	Specifies the timeout value, in seconds, for a remote ping operation. The default is 3 s.
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.
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.

Name	Description
StorageType	Specifies the storage type for this row.
TrapGeneration	Specifies when to generate a notification. The options are: • 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.
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.
Descr	Describes the remote ping test.
SourceAddressType	Specifies the type of the source address used at a remote host when performing a ping operation.
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.
ByPassRouteTable	Enables (optionally) the bypassing of the route table.
DSField	Specifies the value to store in the Differentiated Services (DS) field in the IP packet used to encapsulate the ping probe.

Viewing ping results

About this task

View ping results to view performance-related data.

Procedure

1. From the navigation tree, open 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.
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.

Name	Description
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, open 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).

Name	Description
Time	Indicates the timestamp for this probe result.

Running a traceroute test

About this task

Run a traceroute test to determine the route packets take through a network to a destination.

Procedure

1. From the navigation tree, open 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 and 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 to use on the traceroute request at the remote host.

Name	Description
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.
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.
SourceAddressType	Specifies the type of the source address to use at 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.
IfIndex	Directs the traceroute probes to be transmitted over the specified interface
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.
InitialTtl	Specifies the initial TTL value to use. The default is 1.

Name	Description
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.
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 • PathChange—Generate a TraceRoutePathChange notification after the current path varies from a previously determined path. • TestFailure—Generate a TraceRouteTestFailed notification after the full path to a target cannot be determined. • TestCompletion—Generate 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.
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, open 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, open the following folders: **Configuration** > **Edit** > **Diagnostics**.
2. Click **Ping/Trace Route**.
3. Click the **Trace Route Control** tab.
4. Select an entry.

5. Click **Route Probe History**.

Route Probe History field descriptions

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

Name	Description
OwnerIndex	Identifies the Trace Route entry to which a probe result belongs.
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 19: 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 small form factor pluggable (SFP) transceivers and SFP plus (SFP +) transceivers, see *Avaya Virtual Services Platform 9000 Installation — SFP Hardware Components*, NN46250-305.

Procedure

1. Measure the SFP or 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 could 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 that 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.

Chapter 20: 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 3.1 software and the other side runs a later release, for example, Release 3.2.

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 reenable 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 reenable 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 reenable the IST ports or disable and reenable the IST.
9. Check the software version.

   ```
   show software
   ```

   If the two sides run different versions, upgrade the software.

Example

```
VSP-9012:1>show ist mlt

================================================================================
                                Mlt IST Info
================================================================================
MLT   PEER-IP             VLAN     ENABLE   IST
ID    ADDRESS             ID       IST      STATUS
--------------------------------------------------------------------------------


NEGOTIATED                                    MASTER/
DIALECT      IST STATE                        SLAVE
--------------------------------------------------------------------------------

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>show software

================================================================================
                     software releases in /intflash/release/
================================================================================
3.0.0.0.GA
VSP9K.0.0.0.0int336 (Backup Release)
VSP9K.0.0.0.0int340 (Primary Release)
VSP9K.0.0.0.0int328
VSP9K.0.0.0.0int330
```

```
--------------------------------------------------------------------------------
Auto Commit     : enabled
Commit Timeout  : 10 minutes
```

Variable definitions

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

Table 48: Variable definitions

Variable	Value
1-512	Specifies the MLT ID.

Use the data in the following table to use the `show vlacp interface gigabitethernet` command.

Table 49: 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 50: Variable definitions

Variable	Value
{A.B.C.D}	Specifies the IP address for the IST peer.

Chapter 21: 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

```
VSP-9012:# ping 47.16.10.31
```

```
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:# trace level 9 3
```

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

Variable definitions

Use the data in the following table to use the troubleshooting commands in this procedure.

Table 51: 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).

Variable	Value
WORD<0-128>	Specifies the text string to use as the search criterion.
WORD<0-256>	Specifies the IP address.

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 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 Avaya Virtual Services Platform 9000 only participates in forwarding the requests and replies to and from the client and the DHCP server. The Avaya Virtual Services Platform 9000 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 <1–4094>}
   ```
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, 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, refer to the MS webpage for troubleshooting and configuration issues.

Note:

You may receive some Logs messages that indicate the system cannot forward packets. However, certain situations are not DHCP failures.

Example 1: if you receive 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 message `0x00108705` this indicates a problem at the transmission level. Check 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.

Before you begin

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

Procedure

1. To troubleshoot IPv6 DHCP Relay, you can enable rcip6 trace messages using the following command:
   ```
   trace level 66 3
   ```
2. You can also enable IPv6 forwarding trace using the following command:
   ```
   trace ipv6 forwarding enable <all|debug|error|info|pkt|warn>
   ```

Example

```
VSP-9012:1>enable

VSP-9012:1#trace level 66 3

VSP-9012:1#trace ipv6 forwarding enable all
```

Troubleshooting IPv6 VRRP

The following sections describe troubleshooting information for IPv6 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, refer to 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 52: 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 53: 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

Cause value	Cause definition
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 will form 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.

Before you begin

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

Procedure

1. To troubleshoot IPv6 VRRP, you can enable RCIP6 trace messages using the following command:
   ```
   trace level 66 3
   ```
2. 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
   ```
3. When VRRP is enabled on two routing switches, the Master-Backup relationship will form 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 above trace messages show that the VRRP Master is sending the advertisements correctly at the end of advertisement interval for a VRID.

4. 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 above 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 IST peers with the same IPv6 subnets on the SMLT VLANs (same as for IPv4).
- Make sure that the IST 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. The issue may be observed on dual-stack VLANs after multiple delete and re-adds of IPv4 interfaces, or disabling and reenable of IPv6 forwarding or similar configurations.

Before you begin

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

Procedure

1. Display the RSMLT configuration. This command shows whether the peers are up:
   ```
   show ip rsmlt peer
   ```
2. To recover the peers if they are down, disable and reenable RSMLT on both IST peers:
   ```
   no ip rsmlt
   ```
   ```
   ip rsmlt
   ```
3. If the problem persists, boot from a saved configuration.

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#interface vlan 200
VSP-9012:>(config-if)#show ip rsmlt peer
VSP-9012:>(config-if)#no ip rsmlt
VSP-9012:>(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 may slow down the CPU, especially if executed through 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.

Before you begin

- Log on to Global Configuration mode in ACLI.

Procedure

1. 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).
2. 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 that 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).
3. 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).
4. Disabling and re-enabling IPv6 on the VLAN often recovers connectivity.
5. 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 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 22: 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.

Before you begin

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

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

View IGMP interfaces:

```
show ip igmp interface [gigabitethernet {slot/port[-slot/port][,...]}|vlan <1–4084>]
```

Example

```
VSP-9012:1#show ip igmp interface

================================================================================
                                Igmp Interface
================================================================================
       QUERY              OPER                 QUERY   WRONG            LASTMEM
IF     INTVL STATUS VERS. VERS QUERIER         MAXRSPT QUERY JOINS ROBUST QUERY
--------------------------------------------------------------------------------
V10    125   active 2     2    10.10.10.130    100     0     0     2      10
V11    125   active 2     2    11.11.11.130    100     0     0     2      10
V12    125   active 2     2    12.12.12.130    100     0     0     2      10
V13    125   active 2     2    13.13.13.130    100     0     0     2      10
V14    125   active 2     2    14.14.14.130    100     0     0     2      10
V15    125   active 2     2    15.15.15.130    100     0     0     2      10
V16    125   active 2     2    16.16.16.130    100     0     0     2      10
V17    125   active 2     2    17.17.17.130    100     0     0     2      10
V18    125   active 2     2    18.18.18.130    100     0     0     2      10
V19    125   active 2     2    19.19.19.130    100     0     0     2      10
V20    125   active 2     2    20.20.20.130    100     0     154554 2      10
```

```
V21     125    active 2      2     21.21.21.130     100      0      0      2       10
V22     125    active 2      2     22.22.22.130     100      0      0      2       10
V23     125    active 2      2     23.23.23.130     100      0      0      2       10
V24     125    active 2      2     24.24.24.130     100      0      0      2       10
V25     125    active 2      2     25.25.25.130     100      0      0      2       10


16 out of 39 entries displayed
```

Variable definitions

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

Table 54: Variable definitions

Variable	Value
gigabitethernet {slot/port[-slot/port] [,...]}	Optionally, 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). If you do not specify a slot and port, the command output includes all IGMP interfaces.
vlan <*1–4084*>	Optionally, identifies the VLAN ID. If you do not specify a VLAN ID, the command output includes all IGMP interfaces.

Job aid

The following table shows the field descriptions for the command output if you do not use optional parameters.

Table 55: 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.

Field	Description
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 optional parameters.

Table 56: show ip igmp interface command output with optional parameters

Field	Description
VLAN ID or PORT NUM	Identifies the VLAN or port where IGMP is configured.

Field	Description
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

Before you begin

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

About this task

Multicast group trace tracks the data flow path of the multicast streams.

Procedure

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

```
VSP-9012:1>enable
VSP-9012:1#show ip igmp snoop-trace

================================================================================
                                  Snoop Trace
================================================================================
GROUP           SOURCE        IN    IN    OUT     OUT
ADDRESS         ADDRESS       VLAN  PORT  VLAN    PORT
--------------------------------------------------------------------------------
239.255.255.1   42.13.14.10  200   4/16  100     3/5
```

Variable definitions

Use the data in the following table to use the `show ip igmp snoop-trace` command.

Table 57: 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 58: 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.
OUT PORT	Indicates the outgoing port number.

Viewing IGMP group information

View information about IGMP groups to see the current group operation on the switch.

Before you begin

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

Procedure

View IGMP group information:

```
show ip igmp group group <A.B.C.D> [detail] [vlan <1–4084>]
[port {slot/port[-slot/port][,...]}]
```

```
show ip igmp group group <A.B.C.D> [tracked-members] [vlan <1–
4084>] [port {slot/port[-slot/port][,...]}] [source-subnet
<A.B.C.D/X>] {member-subnet <A.B.C.D./X>]
```

Example

```
VSP-9012:1#show ip igmp group

================================================================================
                                  Igmp Group
================================================================================
GRPADDR         INPORT          MEMBER          EXPIRATION TYPE
--------------------------------------------------------------------------------
224.0.7.130     V20-5/19        20.1.2.3        146        Dynamic
224.0.7.130     V20-5/19        20.1.2.4        184        Dynamic

2 out of 2 group Receivers displayed

Total number of unique groups 1VSP-9012:1#show ip igmp group group 224.0.6.130

================================================================================
```

```
                                    Igmp Group
================================================================================
GRPADDR         INPORT          MEMBER          EXPIRATION TYPE
--------------------------------------------------------------------------------
224.0.7.130     V20-5/19        20.1.2.3        138        Dynamic
224.0.7.130     V20-5/19        20.1.2.4        176        Dynamic

2 out of 2 group Receivers displayed

Total number of unique groups 1

VSP-9012:1#show ip igmp group group 224.0.7.130 detail

Interface:                    Vlan20-5/19
IGMPv3 Group:                 224.0.7.130
Interface Group Mode:         EXCLUDE
Interface Compatibility Mode: IGMP_V3
Interface Group Timer:        134
V2 Host Timer:                Not Running
V1 Host Timer:                Not Running
Interface Group Include Source List:
        Source Address  Expires
        10.10.11.10     172
Interface Group Exclude Source List :
        Source Address  Expires
        10.10.11.11     N/A
        10.10.11.12     N/A

VSP-9012:1#show ip igmp group group 224.0.7.130 tracked-members

================================================================================
                            Members of Channels/Groups
================================================================================
INTERFACE       CHANNEL/GROUP                   MEMBER          MEMBER_MODE EXP
--------------------------------------------------------------------------------
Vlan20-5/19     10.10.11.10/224.0.7.130         20.1.1.4        IS_INCLUDE  165
Vlan20-5/19     */224.0.7.130                   20.1.1.3        IS_EXCLUDE  127
```

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
detail	Use the detail parameter to show IGMPv3–specific data.
group <*A.B.C.D*>	Specifies the address of the IGMP group.
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.

Variable	Value
tracked-members	Use the tracked-members parameter to view all the tracked members for a specific group.
vlan <*1–4084*>	Specifies the VLAN ID.

Job aid

The following table shows the field descriptions for the command output.

Table 59: show ip igmp 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 Avaya Virtual Services Platform 9000 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, you are notified by way of a trap on the console, logged message, or both.

If you do not configure the thresholds, the CLI displays only the ingress and egress records that are currently in use.

Procedure

Show the hardware resource usage:

```
show ip mroute hw-resource-usage
```

Example

```
VSP-9012:1>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 60: 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

Before you begin

You must log on to Global Configuration mode in ACLI.

About this task

Use Protocol Independent Multicast (PIM) traces to aid in PIM troubleshooting.

Procedure

1. Start debug trace message output:
   ```
   debug ip pim pimdbgtrace
   ```
2. Stop debug trace message output:
   ```
   no debug ip pim pimdbgtrace
   ```
3. Configure the system to display trace messages forwarded by the device:
   ```
   debug ip pim send-dbg-trace
   ```
4. Configure the system to display trace messages received by the device:
   ```
   debug ip pim rcv-dbg-trace
   ```
5. Configure the system to display hello messages forwarded or received by the device:
   ```
   debug ip pim hello
   ```
6. Configure the system to display and log debug trace messages:
   ```
   debug ip pim pimdbglog
   ```
7. Configure the system to display register messages forwarded or received by the device:
   ```
   debug ip pim register
   ```
8. 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}
   ```

Example

```
VSP-9012:> enable
VSP-9012:# configure terminal
VSP-9012:(config)# debug ip pim pimdbgtrace
VSP-9012:(config)# debug ip pim send-dbg-trace
VSP-9012:(config)# debug ip pim rcv-dbg-trace
VSP-9012:(config)# debug ip pim hello
VSP-9012:(config)# debug ip pim pimdbglog
VSP-9012:(config)# debug ip pim register
VSP-9012:(config)# debug ip pim source 239.255.255.33
```

Variable definitions

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

Table 61: 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).
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).

Chapter 23: 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.

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, open 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
IfIndex	Shows the interface where IGMP is enabled.
QueryInterval	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.
Status	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.

Name	Description
Version	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.
OperVersion	Shows the version of IGMP that currently runs on this interface.
Querier	Shows the address of the IGMP querier on the IP subnet to which this interface attaches.
QueryMaxResponseTime	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. 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.) **Important:** You must configure this value lower than the QueryInterval.
WrongVersionQueries	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 queries with the wrong version, this value indicates a version mismatch.
Joins	Shows the number of times this interface added a group membership, that is, the number of times an entry for this interface is added to the cache table. This number gives an indication of the amount of IGMP activity over time.
Robustness	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. 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.
LastMembQueryIntvl	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. 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

Name	Description
	values greater than 10. (The value 3 is equal to 0.3 seconds and 10 is equal to 1 second.)
OtherQuerierPresent Timeout	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.
FlushAction	Configures the flush action to one of the following: • none • flushGrpMem • flushMrouter • flushSender
RouterAlertEnable	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. **Important:** To maximize network performance, Avaya recommends that you configure this parameter according to the version of IGMP currently in use. • IGMPv1—Disable • IGMPv2—Enable • IGMPv3—Enable
SsmSnoopEnable	Enables SSM snoop.
ExplicitHostTrackingEnable	Enables or disables IGMPv3 to track hosts per channel or group. The default is disabled. You must select this field if you want to use fast leave for IGMPv3.

Viewing group trace information for IGMP snoop

About this task

View the multicast group trace to track the data flow path of multicast streams.

Procedure

1. In the navigation tree, open 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
GrpAddr	Displays the IP multicast address of the group traversing the router.
SrcAddr	Displays the IP source address of the multicast group.
OutVlan	Displays the egress VLAN ID for the multicast group.
InPort	Displays the ingress port for the multicast group.
InVlan	Displays the ingress VLAN ID for the multicast group.
OutPort	Displays the egress port of the multicast group.

Viewing IGMP group information

About this task

View information about IGMP groups to see the current group operation on the switch.

Note:

The following procedure displays the dynamically learned IGMP groups. **IP > Groups > 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 the statically configured groups.

Procedure

1. In the navigation tree, open 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.

Chapter 24: Unicast routing troubleshooting

Use this section to troubleshoot Layer 3 unicast routing problems.

Using BGP debugging commands

Before you begin

You must log on to BGP Router Configuration mode in ACLI.

About this task

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. Show specific debug messages for the global BGP configuration:
   `global-debug mask WORD<1-100>`
2. Display specific debug messages for the global BGP neighbors:
   `neighbor-debug-all mask WORD<1-100>`
3. Display specific debug messages for BGP peers or peer groups:
   `neighbor <nbr_ipaddr|peer-group-name> neighbor-debug-mask WORD<1-100>`
4. Display debug messages on the console:
   `debug-screen <on|off>`

Example

```
VSP-9012:> enable
VSP-9012:# configure terminal
VSP-9012:(config)# router bgp
VSP-9012:(router-bgp)#
```

Display the global debug messages for error and packet:

```
VSP-9012:(router-bgp)#global-debug mask error,packet
```

End (disable) the display of global debug messages:

```
VSP-9012:(router-bgp)#global-debug mask none
```

Display specific debug messages for the global BGP neighbors:

```
VSP-9012:(router-bgp)#neighbor-debug-all mask packet,event
```

Display specific debug messages for BGP peers or peer groups:

```
VSP-9012:(router-bgp)#neighbor 45.17.10.23 neighbor-debug-mask
event,trace
```

Display debug messages on the console:

```
VSP-9012:(router-bgp)#debug-screen on
```

Variable definitions

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

Table 62: Variable definitions

Variable	Value
<nbr_ipaddr|peer-group-name>	Specifies the IP address or the group name of the peer.
WORD<1-100>	Specifies one or more mask choices that you enter, separated by commas with no space between choices. For example: [<mask>,<mask>,<mask>...]. 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 63: Mask categories and messages

Mask category	Message
none	None disables the display of all debug messages.

Mask category	Message
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 *Avaya Virtual Services Platform 9000 Administration,* NN46250-600.

Procedure

1. Verify that the license is the correct type.
2. Verify that you installed the license properly.

Example

```
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

The Base License activates the features not included in either the Advanced or Premier Licenses.

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)
- 1.5 million IP routes, 500 000 IP FIB entries
- 256 BGP peers
- Lossless Ethernet

The Premier License activates all licensed features on the Virtual Services Platform 9000.

Viewing OSPF errors

Before you begin

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

About this task

Check Open Shortest Path First (OSPF) errors for administrative and troubleshooting purposes.

Procedure

Display information about OSPF errors:

```
show ip ospf port-error [port {slot/port[-slot/port][,...]}]
[vrf WORD<0–16>] [vrfids WORD<0–512>]
```

Variable definitions

Use the data in the following table to use the `show ip ospf port-error` command.

Table 64: 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).
vrf WORD<0–16>	Specifies a VRF by name.
vrfids WORD<0–512>	Specifies a range of VRF IDs.

Job aid

The following table explains the fields in the OSPF error command output.

Table 65: 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.
AUTHTYPEMISMATCH	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.
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

Before you begin

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

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

View the current state of all OSPF neighbors and their current state of adjacency:

```
show ip ospf neighbor
```

Example

```
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 66: 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.

Step	State	Description
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

Before you begin

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

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. Configure the trace level for the OSPF module to terse:
   ```
   trace level 6 2
   ```
2. View the trace information on screen:
   ```
   trace screen enable
   ```
3. Verify if the path is not reachable due to access lists implemented on the device.

4. 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

```
VSP-9012:> enable

VSP-9012:# trace level 6 2

VSP-9012:# trace screen enable
```

Troubleshooting OSPF ExStart/Exchange problems

Before you begin

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

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 the 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. The Virtual Services Platform 9000 automatically checks for OSPF MTU mismatches.

Procedure

1. View the OSPF packets:
   ```
   trace level 6 2
   ```
2. Ensure that the MTU size value for both devices match.
3. Configure the interface to accept OSPF DD packets with a different MTU size:
   ```
   ip ospf mtu-ignore enable
   ```

Example

```
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 25: Customer service

Visit the Avaya Web site to access the complete range of services and support that Avaya provides. Go to www.avaya.com or go to one of the pages listed in the following sections.

Getting technical documentation

To download and print selected technical publications and release notes directly from the Internet, go to www.avaya.com/support.

Getting product training

Ongoing product training is available. For more information or to register, you can access the Web site at www.avaya.com/support. From this Web site, you can locate the Training contacts link on the left-hand navigation pane.

Getting help from a distributor or reseller

If you purchased a service contract for your Avaya product from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance.

Getting technical support from the Avaya Web site

The easiest and most effective way to get technical support for Avaya products is from the Avaya Technical Support Web site at www.avaya.com/support.

Glossary

attenuation	The decrease in signal strength in an optical fiber caused by absorption and scattering.
Control Processor (CP) module	The Control Processor module is responsible for running all high level protocols (BGP, OSPF) and distributing the results (routing updates) to the rest of the system, managing and configuring the IO and Switch Fabric modules, and maintaining and monitoring the health of the chassis.
cyclic redundancy check (CRC)	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.
Dynamic Random Access Memory (DRAM)	A read-write random-access memory, in which the digital information is represented by charges stored on the capacitors and must be repeatedly replenished in order to retain the information.
Electrostatic Discharge (ESD)	The discharge of stored static electricity that can damage electronic equipment and impair electrical circuitry that results in complete or intermittent failures.
forwarding database (FDB)	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.
Generalized Regular Expression Parser (grep)	A Unix command used to search files for lines that match a given regular expression (RE).
Interface (IO) cooling module	The IO cooling module is a hot swappable fan tray used to cool the interface and CP modules.
Interface module	An interface module is a module that provides network connectivity for various media (sometimes called Layer 0) and protocol types. Interface modules are also called Ethernet modules.
Internet Assigned Numbers Authority (IANA)	The central registry for various assigned numbers, for example, Internet protocol parameters (such as port, protocol, and enterprise numbers), options, codes, and types.
Internet Control Message Protocol (ICMP)	A collection of error conditions and control messages exchanged by IP modules in both hosts and gateways.

Internet Group Management Protocol (IGMP)	A host membership protocol used to arbitrate membership in multicast services.
Internet Protocol multicast (IPMC)	The technology foundation for audio and video streaming, push applications, software distribution, multipoint conferencing, and proxy and caching solutions.
interswitch trunking (IST)	A feature that uses one or more parallel point-to-point links to connect two aggregation switches. The two aggregation switches use this channel to share information and operate as a single logical switch. Only one interswitch trunk can exist on each Split Multilink Trunking (SMLT) aggregation switch.
Layer 1	The Physical Layer of the Open System Interconnection (OSI) model. Layer 1 interfaces with the MAC sublayer of Layer 2, and performs character encoding, transmission, reception, and character decoding.
Layer 2	The Data Link Layer of the OSI model. Examples of Layer 2 protocols are Ethernet and Frame Relay.
Layer 3	The Network Layer of the OSI model. Example of a Layer 3 protocol is Internet Protocol (IP).
Layer 4	The Transport Layer of the OSI model. An example of a Layer 4 protocol is Transfer Control Protocol (TCP).
light emitting diode (LED)	A semiconductor diode that emits light when a current passes through.
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.
management information base (MIB)	Defines system operations and parameters used for the Simple Network Management Protocol (SNMP).
mask	A bit string that is used along with an IP address to indicate the number of leading bits in the address that correspond with the network part.
maximum transmission unit (MTU)	The largest number of bytes in a packet—the maximum transmission unit of the port.
Media Access Control (MAC)	Arbitrates access to and from a shared medium.
mirrored port	The port to mirror. The port is also called the source port.

mirroring multilink trunk	The multilink trunk to which the traffic is mirrored.
mirroring port	The port to which all traffic is mirrored, also referred to as the destination port.
mirroring VLAN	The virtual Local Area Network (VLAN) to which the traffic is mirrored.
multicast group ID (MGID)	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 data is directed 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.
MultiLink Trunking (MLT)	A method of link aggregation that uses multiple Ethernet trunks aggregated to provide a single logical trunk. A multilink trunk provides the combined bandwidth of multiple links and the physical layer protection against the failure of a single link.
next hop	The next hop to which a packet can be sent to advance the packet to the destination.
nonbroadcast multiaccess (NBMA)	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.
Open Systems Interconnection (OSI)	A suite of communication protocols, network architectures, and network management standards produced by the International Organization for Standardization (ISO). OSI-compliant systems can communicate with other OSI-compliant systems for a meaningful exchange of information.
Packet Capture Tool (PCAP)	A data packet capture tool that captures ingress and egress (on Ethernet modules only) packets on selected ports. You can analyze captured packets for troubleshooting purposes.
port mirroring	A feature that sends received or transmitted traffic to a second destination.
Protocol Data Units (PDUs)	A unit of data that is specified in a protocol of a given layer and that consists of protocol-control information of the given layer and possibly user data of that layer.
quality of service (QoS)	Use QoS features to reserve resources in a congested network. For example, you can configure a higher priority to IP deskphones, which need a fixed bit rate, and, split the remaining bandwidth between data connections if calls in the network are important than the file transfers.
remote mirroring	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.
remote mirror source (RMS)	The port that generates the mirrored encapsulated traffic.

remote mirror target (RMT)	The port that decapsulates the remote mirror traffic and transmits it out of the device.
remote monitoring (RMON)	A remote monitoring standard for Simple Network Management Protocol (SNMP)-based management information bases (MIB). The Internetwork Engineering Task Force (IETF) proposed the RMON standard to provide guidelines for remote monitoring of individual LAN segments.
route table manager (RTM)	Determines the best route to a destination based on reachability, route preference, and cost.
Routing Information Protocol (RIP)	A distance vector protocol in the IP suite, used by IP network-layer protocol, that enables routers in the same AS to exchange routing information by means of periodic updates. The RIP is most often used as a very simple IGP within small networks.
Secure Shell (SSH)	Used for secure remote logons and data transfer over the Internet. SSH uses encryption to provide security.
Secure Sockets Layer (SSL)	An Internet security encryption and authentication protocol for secure point-to-point connections over the Internet and intranets, especially between clients and servers.
Simple Loop Prevention Protocol (SLPP)	Simple Hello Protocol that prevents loops in a Layer 2 network (VLAN).
Simple Network Management Protocol (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 gigabit per second (Gb/s).
spanning tree	A simple, fully-connected active topology formed from the arbitrary physical topology of connected bridged Local Area Network components by relaying frames through selected bridge ports. The protocol parameters and states that are used and exchanged to facilitate the calculation of the active topology and to control the bridge relay function.
Spanning Tree Group (STG)	A collection of ports in one spanning tree instance.

time-to-live (TTL)	The field in a packet used to determine the valid duration for the packet; the TTL determines the packet lifetime. The system discards a packet with a TTL of zero.
Trivial File Transfer Protocol (TFTP)	A protocol that governs transferring files between nodes without protection against packet loss.
trunk	A logical group of ports that behaves like a single large port.
unshielded twisted pair (UTP)	A cable with one or more pairs of twisted insulated copper conductors bound in a single plastic sheath.
user-based security model (USM)	A security model that uses a defined set of user identities for authorized users on a particular Simple Network Management Protocol (SNMP) engine.
User Datagram Protocol (UDP)	In TCP/IP, a packet-level protocol built directly on the Internet Protocol layer. TCP/IP host systems use UDP for application-to-application programs.
view-based access control model (VACM)	Provides context, group access, and group security levels based on a predefined subset of management information base (MIB) objects.
Virtual Link Aggregation Control Protocol (VLACP)	Virtual Link Aggregation Control Protocol (VLACP) is a Layer 2 handshaking protocol that can detect end-to-end failure between two physical Ethernet interfaces.
virtual router (VR)	An abstract object managed by the Virtual Router Redundancy Protocol (VRRP) that acts as a default router for hosts on a shared LAN.
virtual router forwarding (VRF)	Provides traffic isolation between customers operating over the same node. Each virtual router emulates the behavior of a dedicated hardware router by providing separate routing functionality, and the network treats each VRF as a separate physical router.
Virtual Router Redundancy Protocol (VRRP)	A protocol used in static routing configurations, typically at the edge of the network. This protocol operates on multiple routers on an IP subnet and elects a primary gateway router. When the primary router fails, a backup router is quickly available to take its place.