ExtremeSwitching™

Troubleshooting VOSS

Release 8.1.5 (VOSS)
9035884-01 Rev AA
April 2020

© 2017-2020, Extreme Networks, Inc.
All Rights Reserved.

Legal Notice

Extreme Networks, Inc. reserves the right to make changes in specifications and other information contained in this document and its website without prior notice. The reader should in all cases consult representatives of Extreme Networks to determine whether any such changes have been made.

The hardware, firmware, software or any specifications described or referred to in this document are subject to change without notice.

Trademarks

Extreme Networks and the Extreme Networks logo are trademarks or registered trademarks of Extreme Networks, Inc. in the United States and/or other countries.

All other names (including any product names) mentioned in this document are the property of their respective owners and may be trademarks or registered trademarks of their respective companies/owners.

For additional information on Extreme Networks trademarks, please see: www.extremenetworks.com/company/legal/trademarks

Open Source Declarations

Some software files have been licensed under certain open source or third-party licenses. End-user license agreements and open source declarations can be found at: www.extremenetworks.com/support/policies/software-licensing

Contents

Chapter 1: About this Document 9
Purpose 9
Conventions 9
Text Conventions 10
Documentation and Training 12
Getting Help 12
Providing Feedback 13
Chapter 2: New in this Document 14
Notice about Feature Support 14
Chapter 3: Troubleshooting Overview 15
Data Collection Required for Technical Support Cases 15
Data Collection for an Outage 15
Data Collection for Non Outage Problems 17
Troubleshooting Planning Fundamentals 18
Proper Installation and Routine Maintenance 18
Network Configuration 19
Normal Behavior on the Network 20
Troubleshooting Fundamentals 21
Connectivity Problems 21
Routing Table Problems 22
Cable Connection Problems 23
Alarm Database 23
LED Indications of Problems 24
Timestamp in show command outputs 24
Chapter 4: Logs and Traps 25
Logs and Traps Fundamentals 26
Overview of Traps and Logs 26
Secure Syslog 29
Simple Network Management Protocol 29
Log Message Format 30
Log Files 33
Log File Transfer 34
Email Notification 35
Log Configuration Using CLI 36
Configuring a UNIX System Log and Syslog Host 37
Configuring Secure Forwarding 39
Installing Root Certificate for Syslog Client 41
Configuring Logging 42
Configuring the Remote Host Address for Log Transfer 43

Configuring System Logging..... 44
Configuring System Message Control..... 45
Extending System Message Control..... 46
Viewing Logs..... 47
Configuring CLI Logging..... 49
Configuring Email Notification..... 51
Log Configuration Using EDM..... 55
Configuring the System Log..... 55
Configuring the System Log Table..... 56
Configuring Email Notification..... 58
SNMP Trap Configuration Using CLI..... 60
Configuring an SNMP Host..... 60
Configuring an SNMP Notify Filter Table..... 61
Configuring SNMP Interfaces..... 62
Enabling SNMP Trap Logging..... 63
SNMP Trap Configuration Using EDM..... 65
Configuring an SNMP Host Target Address..... 65
Configuring Target Table Parameters..... 66
Configuring SNMP Notify Filter Profiles..... 67
Configuring SNMP Notify Filter Profile Table Parameters..... 68
Enabling Authentication Traps..... 69
Viewing the Trap Sender Table..... 69
Chapter 5: Connectivity Fault Management..... 71
CFM Fundamentals..... 71
Maintenance Domain (MD)..... 73
Maintenance Association (MA)..... 74
Maintenance Association Endpoint (MEP)..... 75
Fault Verification..... 75
LBM Message..... 75
L2 Ping..... 76
Fault Isolation..... 76
Link Trace Message..... 76
L2 Traceroute..... 77
L2 Tracetree..... 79
L2 Tracetree-fan..... 79
Maintenance Domain Intermediate Point (MIP)..... 79
Layer 2 Tracemroute..... 79
Nodal MPs..... 80
Configuration Considerations..... 80
CFM Configuration Using CLI..... 82
Autogenerated CFM..... 83
Configuring Explicit Mode CFM..... 87
Triggering a Loopback Test (LBM)..... 94

Triggering Linktrace (LTM)..... 95
Triggering a Layer 2 Ping..... 96
Trigger a Layer 2 Traceroute..... 98
Trigger a Layer 2 Tracetree..... 100
Triggering a Layer 2 Tracetree-fan..... 101
Trigger a Layer 2 Tracemroute..... 102
Using trace CFM to Diagnose Problems..... 104
Using trace SPBM to Diagnose Problems..... 107
CFM Configuration Using EDM..... 110
Autogenerated CFM..... 110
Configuring Explicit CFM..... 114
Configuring Layer 2 Ping..... 118
Initiate a Layer 2 Traceroute..... 120
View Layer 2 Traceroute Results..... 123
Configuring Layer 2 IP Ping..... 124
Viewing Layer 2 IP Ping Results..... 126
Configuring Layer 2 IP Traceroute..... 127
Viewing Layer 2 IP Traceroute Results..... 128
Triggering a Loopback Test..... 130
Triggering Linktrace..... 132
Viewing Linktrace Results..... 135
Configure a Layer 2 Tracetree..... 136
View Layer 2 Tracetree Results..... 139
Configure Layer 2 Trace Multicast Route on a VLAN..... 140
Configure Layer 2 Tracemroute on a VRF..... 143
Viewing Layer 2 Trace Multicast Route Results..... 145
CFM Configuration Example..... 146
CFM Configuration Example..... 146
CFM Sample Output..... 147
Chapter 6: Software Troubleshooting tool configuration..... 152
Troubleshooting Tool Fundamentals..... 152
Troubleshooting Overview..... 152
Debug Files..... 153
Digital Diagnostic Monitoring..... 154
Flight Recorder..... 155
Port Mirroring..... 156
General Diagnostic Tools..... 161
Fabric RSPAN (Mirror to I-SID)..... 163
Software Troubleshooting Tool Configuration Using CLI..... 164
Using CLI for Troubleshooting..... 164
Using Software Record Dumps..... 170
Using Trace to Diagnose Problems..... 170
Using Trace to Diagnose IPv6 Problems..... 173

View and Delete Debug Files..... 175
Configuring Port Mirroring..... 180
Displaying Mirror Resource Usage..... 184
Configuring Global Mirroring Actions with an ACL..... 185
Configure ACE Actions to Mirror..... 185
Clearing ARP Information for an Interface..... 186
Flushing Routing, MAC, and ARP Tables for a Port..... 187
Flushing Routing, MAC, and ARP Tables for a VLAN..... 188
Pinging an IP Device..... 189
Running a Traceroute Test..... 192
Showing SNMP Logs..... 194
Using Trace to Examine IS-IS Control Packets..... 195
Viewing the Metric Type of IS-IS Route in TLVs – Detailed..... 200
Viewing the Metric Type of IS-IS Route in TLVs – Summarized..... 201
Configuring I-SID Monitoring..... 201
Displaying I-SID Monitoring Diagnostics..... 203
Displaying I-SID Mirroring Statistics..... 204
Clearing Fabric RSPAN (Mirror to I-SID) Statistics..... 204
Software Troubleshooting Tool Configuration Using EDM..... 205
Flushing Routing Tables by VLAN..... 205
Flushing Routing Tables by Port..... 206
Configuring Port Mirroring..... 206
Configuring ACLs for Mirroring..... 208
Configuring ACEs for Mirroring..... 209
Running a Ping Test..... 211
Viewing Ping Results..... 213
Viewing Ping Probe History..... 214
Running a Traceroute Test..... 214
Viewing Traceroute Results..... 217
Viewing the Traceroute History..... 218
Configuring I-SID Monitoring..... 219
Viewing and Clearing Fabric RSPAN (Mirror to I-SID) Statistics..... 220
Chapter 7: General Troubleshooting..... 221
Hardware Troubleshooting..... 221
Using Trace to Diagnose Hardware Problems..... 221
Troubleshooting USB Viewing Problems..... 222
Software Troubleshooting..... 224
Failure to Read Failed Configuration File..... 224
No Web Management Interface Access to a Device..... 224
CPU Queues..... 224
Chapter 8: Layer 1 Troubleshooting..... 227
Troubleshooting Fiber Optic Links..... 227
Resetting a QSFP+ or QSFP28 Transceiver..... 228

Chapter 9: Layer 2 and 3 Troubleshooting 230
Troubleshooting BPDU Guard 230
No Packets Received on the Port 230
Troubleshooting IPv6 VRRP 232
VRRP Transitions 233
Enabling Trace Messages for IPv6 VRRP Troubleshooting 234
Risks Associated with Enabling Trace Messages 235
VRRP with Higher Priority Running as Backup 235
Troubleshooting RSMLT 236
RSMLT Peers Not Up 236
Enabling Trace Messages for RSMLT Troubleshooting 237
Troubleshooting IPv6 Connectivity Loss 237
Troubleshooting vIST Failure 238
Troubleshooting Transparent Port UNI 239
Viewing all Configured I-SIDs 239
Viewing C-MACs Learned on T-UNI Ports for an I-SID 243
Multicast Troubleshooting 245
Multicast Feature Troubleshooting 245
Multicast Routing Troubleshooting Using CLI 251
Multicast routing troubleshooting using EDM 270
Troubleshoot MACsec 275
Viewing the MACsec Connectivity Association Details 276
Viewing MACsec Status 277
Troubleshoot MACsec Using EDM 279
View MACsec Connectivity Association Details 279
Troubleshooting Fabric Attach 280
Troubleshooting Fabric Attach using the CLI 281
Troubleshooting Fabric Attach using the EDM 295
Fabric Attach troubleshooting example 307
Troubleshooting FAN Transit 311
Viewing FAN Transit information - Detailed 311
Chapter 10: Upper Layer Troubleshooting 316
Troubleshooting SNMP 316
SNMP Trap not Received 317
Troubleshooting DHCP 318
Troubleshooting DHCP Relay 319
Troubleshooting Client Connection to the DHCP Server 321
Troubleshooting IPv6 DHCP Relay 322
IPv6 DHCP Relay Switch Side Troubleshooting 322
IPv6 DHCP Relay Server Side Troubleshooting 323
IPv6 DHCP Relay Client Side Troubleshooting 323
Enabling Trace Messages for IPv6 DHCP Relay 324
Downgrading or Upgrading from Releases that Support Different Key Sizes 324

Troubleshooting TACACS+ 325
Unable to Log On Using Telnet or Rlogin 325
Unable to Log On Using SSH 330
Unable to Log On by any Means (Telnet, Rlogin, or SSH) 331
Administrator Unable to Obtain Accounting Information from the TACACS+ Server 334
Trap Server Cannot Receive Trap Packets from the Switch 335
Troubleshooting TACACS+ Problems 336
Using BGP Debugging Commands 338
Chapter 11: Traps Reference 341
Proprietary Traps 341
Standard Traps 355
Glossary 359

Chapter 1: About this Document

This section discusses the purpose of this document, the conventions used, ways to provide feedback, additional help, and information regarding other Extreme Networks publications.

Purpose

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

- Extreme Networks VSP 4000 Series (includes VSP 4450 Series)
- Extreme Networks VSP 4900 Series
- Extreme Networks VSP 7200 Series
- Extreme Networks VSP 7400 Series
- Extreme Networks VSP 8000 Series (includes VSP 8200 Series and VSP 8400 Series)
- Extreme Networks VSP 8600 Series
- Extreme Networks XA1400 Series

Note:

VOSS is licensed on the XA1400 Series as a Fabric Connect VPN (FCVPN) application, which includes a subset of VOSS features. FCVPN transparently extends Fabric Connect services over third-party provider networks.

This troubleshooting document describes common problems and error messages, provides information about traps and command logging, and provides techniques you can use to resolve common problems. This document provides troubleshooting information for VOSS switches. This document also provides information about troubleshooting tools: for example, port mirroring.

Examples and network illustrations in this document may illustrate only one of the supported platforms. Unless otherwise noted, the concept illustrated applies to all supported platforms.

Conventions

This section discusses the conventions used in this guide.

Text Conventions

The following tables list text conventions that can be used throughout this document.

Table 1: Notice Icons

Icon	Alerts you to...
Important:	A situation that can cause serious inconvenience.
Note:	Important features or instructions.
Tip:	Helpful tips and notices for using the product.
Danger:	Situations that will result in severe bodily injury; up to and including death.
Warning:	Risk of severe personal injury or critical loss of data.
Caution:	Risk of personal injury, system damage, or loss of data.

Table 2: Text Conventions

Convention	Description
Angle brackets (< >)	Angle brackets (< >) indicate that you choose the text to enter based on the description inside the brackets. Do not type the brackets when you enter the command. If the command syntax is `cfm maintenance-domain maintenance-level <0-7>` , you can enter `cfm maintenance-domain maintenance-level 4`.
Bold text	Bold text indicates the GUI object name you must act upon. Examples: • Click **OK**. • On the **Tools** menu, choose **Options**.
Braces ({ })	Braces ({ }) indicate required elements in syntax descriptions. Do not type the braces when you enter the command. For example, if the command syntax is `ip address {A.B.C.D}`, you must enter the IP address in dotted, decimal notation.

Table continues…

Convention	Description
Brackets ([ ])	Brackets ([ ]) indicate optional elements in syntax descriptions. Do not type the brackets when you enter the command. For example, if the command syntax is `show clock [detail]`, you can enter either `show clock` or `show clock detail`.
Ellipses (…)	An ellipsis (…) indicates that you repeat the last element of the command as needed. For example, if the command syntax is `ethernet/2/1 [ <parameter> <value> ]...`, you enter `ethernet/2/1` and as many parameter-value pairs as you need.
Italic Text	Italics emphasize a point or denote new terms at the place where they are defined in the text. Italics are also used when referring to publication titles that are not active links.
`Plain Courier Text`	Plain Courier text indicates command names, options, and text that you must enter. Plain Courier text also indicates command syntax and system output, for example, prompts and system messages. Examples: • `show ip route` • `Error: Invalid command syntax [Failed][2013-03-22 13:37:03.303 -04:00]`
Separator (>)	A greater than sign (>) shows separation in menu paths. For example, in the Navigation tree, expand the **Configuration** > **Edit** folders.
Vertical Line (|)	A vertical line (|) separates choices for command keywords and arguments. Enter only one choice. Do not type the vertical line when you enter the command. For example, if the command syntax is `access-policy by-mac action { allow | deny }`, you enter either `access-policy by-mac action allow` or `access-policy by-mac action deny`, but not both.

Documentation and Training

Find Extreme Networks product information at the following locations:

Current Product Documentation
Release Notes
Hardware/software compatibility matrices for Campus and Edge products
Supported transceivers and cables for Data Center products
Other resources, like white papers, data sheets, and case studies

Extreme Networks offers product training courses, both online and in person, as well as specialized certifications. For details, visit www.extremenetworks.com/education/.

Getting Help

If you require assistance, contact Extreme Networks using one of the following methods:

Extreme Portal	Search the GTAC (Global Technical Assistance Center) knowledge base; manage support cases and service contracts; download software; and obtain product licensing, training, and certifications.
The Hub	A forum for Extreme Networks customers to connect with one another, answer questions, and share ideas and feedback. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.
Call GTAC	For immediate support: (800) 998 2408 (toll-free in U.S. and Canada) or 1 (408) 579 2826. For the support phone number in your country, visit: www.extremenetworks.com/support/contact

Before contacting Extreme Networks for technical support, have the following information ready:

- Your Extreme Networks service contract number, or serial numbers for all involved Extreme Networks products
- A description of the failure
- A description of any actions already taken to resolve the problem
- A description of your network environment (such as layout, cable type, other relevant environmental information)
- Network load at the time of trouble (if known)
- The device history (for example, if you have returned the device before, or if this is a recurring problem)
- Any related RMA (Return Material Authorization) numbers

Subscribe to Service Notifications

You can subscribe to email notifications for product and software release announcements, Vulnerability Notices, and Service Notifications.

1. Go to www.extremenetworks.com/support/service-notification-form.
2. Complete the form (all fields are required).
3. Select the products for which you would like to receive notifications.

   Note:
   You can modify your product selections or unsubscribe at any time.
4. Select **Submit**.

Providing Feedback

The Information Development team at Extreme Networks has made every effort to ensure the accuracy and completeness of this document. We are always striving to improve our documentation and help you work better, so we want to hear from you. We welcome all feedback, but we especially want to know about:

- Content errors, or confusing or conflicting information.
- Improvements that would help you find relevant information in the document.
- Broken links or usability issues.

If you would like to provide feedback, you can do so in three ways:

- In a web browser, select the feedback icon and complete the online feedback form.
- Access the feedback form at https://www.extremenetworks.com/documentation-feedback/.
- Email us at documentation@extremenetworks.com.

Provide the publication title, part number, and as much detail as possible, including the topic heading and page number if applicable, as well as your suggestions for improvement.

Chapter 2: New in this Document

The following sections detail what is new in this document.

IQAgent

VOSS integrates with ExtremeCloud IQ using IQAgent. When you enable IQAgent on supported VOSS devices, the devices can register with, and send monitoring and device data, such as port statistics and device health, to ExtremeCloud IQ.

The **`show running-config module`** command is updated to include IQAgent support. For more information, see Using CLI for Troubleshooting on page 164.

Notice about Feature Support

This document includes content for multiple hardware platforms across different software releases. As a result, the content can include features not supported by your hardware in the current software release.

If a documented command, parameter, tab, or field does not appear on your hardware, it is not supported.

For information about physical hardware restrictions, see your hardware documentation.

Chapter 3: Troubleshooting Overview

Use the information in this chapter to learn about general troubleshooting guidelines and helpful tips for common problems.

This chapter includes the following sections:

- Data Collection Required for Technical Support Cases
- Troubleshooting Planning Fundamentals
- Troubleshooting Fundamentals

Data Collection Required for Technical Support Cases

Use the following sections to learn about how to gather information before you contact Technical Support.

Data Collection for an Outage

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

Collecting Data Before You Restart

Perform this procedure before you restart the chassis.

Procedure

1. Capture the current state of the chassis:

   ```
   terminal more disable

   show tech
   ```

2. Capture Flight Recorder trace information.

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

   The **all** command executes three separate commands: `flight-recorder snapshot`, `flight-recorder trace`, and `flight-recorder archive`.

For VSP 8600, you can replace `{slot[-slot][,...]}` with `all` to create a snapshot for all slots.

3. Reset the chassis:
   a. Reset the chassis without creating a core file.

      ```
      reset -y
      ```

   b. For VSP 8600, create a core file of the primary Control Processor (CP):

      ```
      sys action cpu-switch-over -coredump
      ```

   c. For VSP 8600, create a core file for a particular module:

      ```
      slot reset {<1-8> |SF1 | SF2 | SF3} -coredump
      ```

   d. For all other VOSS platforms, create an ssio core file and a cbcp-main.x core file, skip the confirmation question, and then reset the chassis.

      ```
      reset -coredump -y
      ```

   e. Create an ssio core file and a cbcp-main.x core file, prompt for the confirmation question, and then reset the chassis.

      ```
      reset -coredump
      ```

      Note:

      Create a core file only when there is a need to analyze a problem. If you reset the switch for any other reason the command is `reset -y`.

4. Continue with Collecting data after you restart on page 16.

Example

The following example shows output of the `flight-recorder all 1` command.

```
Switch: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.
```

Collecting Data After You Restart

About this task

Perform this procedure after you restart the affected chassis.

Procedure

1. Use FTP to transfer the following information:
   - Configuration files from each chassis: Stored on the internal flash at `/intflash/`.
   - Log files from each chassis: Stored on the internal flash at `/intflash/`.
   - Generated archive files for slot: Stored on the internal flash. For example: `/intflash/archive/<slot>`

   Note:

   For VSP 8600, if the core file is not on the primary control processor (CP), you can use FTP to connect to the internal flash of the primary CP. For example, `cp mnt/intflash/archive/<slot>/<filename> /intflash/<filename>`

   To copy the file to internal flash or usb device in the primary CP, use the following command: `cp mnt/intflash/archive/<slot>/<filename> {/intflash/<filename> | /usb/<filename>}`
2. Show core information:

   ```
   show core—files
   ```

   If the timestamp for an entry in the command output matches the time the outage first occurred, or is later than that time, transfer that core file 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 Technical Support.

You must attempt to resolve the problem using this document. Contact Technical Support 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 Technical Support:

- a detailed description of the problem
- the date and time when the problem started
- the frequency of the problem
- if this is a new installation

- if there is relevant information recorded on the support portal — Were related problem solutions found? Is there currently a work around for this issue? For more information, see support on the Extreme Portal at https://extremeportal.force.com/ExtrSupportHome.
- if the system was recently upgraded — Have you recently changed or upgraded the system, the network, or a custom application? (For example, has configuration or code been changed?) When were these changes made? Provide the date and time. Who made these changes? Were the changes made by a partner or customer? Provide the names of the individuals who made the changes.

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

The following table lists the documents that provide maintenance and installation procedures.

To prevent problems, follow proper maintenance and installation procedures.

Table 3: Maintenance and installation documentation

Subject area	Document
Installation, environmental requirements	Installing the Virtual Services Platform 4450GSX-PWR+ Installing the Virtual Services Platform 4450GTX-HT-PWR+ Installing the Virtual Services Platform 4850GTS Series Installing the Virtual Services Platform 7200 Series VSP 7400 Series Switches: Hardware Installation Guide Installing the Virtual Services Platform 8000 Series Installing the Virtual Services Platform 8600
Transceiver installation and requirements	*Extreme Networks Pluggable Transceivers Installation Guide*

Table continues…

Subject area	Document
	Find VSP Components at *Extreme Hardware/ Software Compatibility and Recommendation Matrices*

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

The switch 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.

Note:

Depending on the hardware platform, the output of **`show khi performance memory`** command can differ.

Example

In the following example, the **`show khi performance memory`** command shows the average memory utilization at various time intervals. The **`show khi performance memory history`** command shows the VMSize utilization values in kilobytes for each process at various time intervals. After 1 hour elapses, the system stores this information in `/intflash/coreFiles/slot/khi_mem_log`.

```
Switch:1#show khi performance memory
    Slot:1Slot:1
         Used: 1609636 (KB)
         Free: 2396068 (KB)
         Current utilization: 40 %
         5-minute average utilization: 40 %
         5-minute high water mark: 40 (%)
         10-minute average utilization: 39 %
         10-minute high water mark: 39 (%)
         1-Hour average utilization: 37 %
         1-Day average utilization: 0 %
         1-Month average utilization: 0 %
         1-Year average utilization: 0 %

Switch:1#show khi performance memory history
    Slot:1
```

```
Values indicate VMSize in KB

Pid    Pname                5-Min    10-Min   1-Hour   1-Day    1-Month  1-Year

 4762  logger               1        1        1        --       --       --
 4779  namServer            20       20       20       --       --       --
 4780  sockserv             4        4        4        --       --       --
 4782  oom95                214      214      214      --       --       --
 4784  oom90                214      214      214      --       --       --
 4786  imgsync.x            19       19       19       --       --       --
 4860  logServer            24       24       24       --       --       --
 4861  trcServer            18       18       18       --       --       --

--More-- (q = quit)

Switch:1#show tech

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

General Info :

        SysDescr     : Switch (4.5.0.1_B008) (PRIVATE)
        SysName      : Switch
        SysUpTime    : 0 day(s), 00:49:06
        SysContact   : support@extremenetworks.com
        SysLocation  :

Chassis Info:

        Chassis            : 8608
        Serial#            : SDNI86CWD018
        H/W Revision       : R0D
        H/W Config         :
        Part Number        : EC9402001-E6
        NumSlots           : 8
        NumPorts           : 80
        BaseMacAddr        : f8:73:a2:03:80:00

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

Troubleshooting Fundamentals

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

Connectivity Problems

Use the following general tasks to isolate connectivity problems:

- Check physical connectivity. Verify if an alarm for link or port down exists.
- Check the link state by viewing the **`show interface {gigabitEthernet|loopback|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

```
Switch:1#show interface vlan

================================================================================================
                                          Vlan Basic
================================================================================================
VLAN                                MSTP
ID    NAME             TYPE         INST_ID PROTOCOLID   SUBNETADDR      SUBNETMASK      VRFID
------------------------------------------------------------------------------------------------
1     Default          byPort       0       none         N/A             N/A             0
2     VLAN-2           byPort       0       none         N/A             N/A             0
50    mcast_smlt_3     byPort       1       none         N/A             N/A             0
60    mcast_smlt_4     byPort       1       none         N/A             N/A             0

All 4 out of 4 Total Num of Vlans displayed


================================================================================
        Vlan Port
================================================================================
VLAN PORT               ACTIVE             STATIC             NOT_ALLOW
ID   MEMBER             MEMBER             MEMBER             MEMBER
--------------------------------------------------------------------------------
1    1/5-1/8,1/10-1/48, 1/5-1/8,1/10-1/48,
     2/2-2/6            2/2-2/6

2

50   1/1,2/1/3-2/1/4    1/1,2/1/3-2/1/4

60   1/3-1/4            1/3-1/4



All 4 out of 4 Total Num of Port Entries displayed


================================================================================
                              VLAN VRF Association
================================================================================
VLAN  VRF
ID    NAME
--------------------------------------------------------------------------------
1     GlobalRouter
2     GlobalRouter
50    GlobalRouter
60    GlobalRouter
4086  GlobalRouter
4087  GlobalRouter

--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 to MAC assignment
- problems with administrative distance or other parameters

Important:

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.

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).

Pluggable optic cables

Cables for the optical transceivers vary depending on the specific device type.

For more information about the cable requirements for optical transceivers, see *Extreme Networks Pluggable Transceivers Installation Guide*.

Alarm Database

The switch 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 CLI. Local alarms are an automatic mechanism run by the system that do not require additional configuration.

The fabric drivers on the I/O and control (IOC) modules or switch fabric (SF) modules can trigger an alarm if the module is not operational. You can view the alarm from the primary control processor (CP) console or from the **`show alarm database`** output. The message indicates the slot and module that is not operational, and the remote slot and SFI to which it is connected. The alarm clears when the SFI state returns to operational.

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 47.

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 shows the following alarm text:

```
CP1 [01/01/70 00:03:06.796] 0x00010844 00000000 Global Router HW WARNING
USB found in slot 1 has VendorId 05dc ProductId a01a and Manufacturer
Lexar and did not match supported devices
```

This alarm means that you have tried to insert an unsupported USB device into the USB slot. Only the USB device provided with your system can be inserted into the USB slot.

LED Indications of Problems

For information on LEDs on the chassis, see the hardware installation documentation.

Timestamp in show command outputs

The output for all CLI show commands includes a timestamp header to indicate when the command output was generated. This information can be helpful when communicating with Support.

The following command output shows a timestamp example.

```
Switch:1#show alarm statistics
********************************************************************************
                Command Execution Time: Wed Nov 07 19:55:15 2018 UTC
********************************************************************************

================================================================================================
                                ALARM STATISTICS
================================================================================================
PERSISTENT PERSISTENT PERSISTENT  PERSISTENT  DYNAMIC   DYNAMIC  DYNAMIC DYNAMIC
  ALARM      ACTIVE    CLEARED      WRPRD      ALARM    ACTIVE   CLEARED  WRPRD
    0          0          0          0         11        8         3       0
```

Chapter 4: Logs and Traps

Table 4: Logs and traps product support

Feature	Product	Release introduced
For configuration details, see Troubleshooting VOSS.		
Logging to a file and syslog (IPv4)	VSP 4450 Series	VSP 4000 4.0
	VSP 4900 Series	VOSS 8.1
	VSP 7200 Series	VOSS 4.2.1
	VSP 7400 Series	VOSS 8.0
	VSP 8200 Series	VSP 8200 4.0
	VSP 8400 Series	VOSS 4.2
	VSP 8600 Series	VSP 8600 4.5
	XA1400 Series	VOSS 8.0.50
Logging to a file and syslog (IPv6)	VSP 4450 Series	VOSS 4.1
	VSP 4900 Series	VOSS 8.1
	VSP 7200 Series	VOSS 4.2.1
	VSP 7400 Series	VOSS 8.0
	VSP 8200 Series	VOSS 4.1
	VSP 8400 Series	VOSS 4.2
	VSP 8600 Series	VSP 8600 6.2
	XA1400 Series	Not Supported
Simple Mail Transfer Protocol (SMTP) for log notification	VSP 4450 Series	VOSS 6.0
	VSP 4900 Series	VOSS 8.1
	VSP 7200 Series	VOSS 6.0
	VSP 7400 Series	VOSS 8.0
	VSP 8200 Series	VOSS 6.0
	VSP 8400 Series	VOSS 6.0
	VSP 8600 Series	VSP 8600 6.1
	XA1400 Series	VOSS 8.0.50
System Logging compliance with RFC 5424 and RFC 3339	VSP 4450 Series	VOSS 6.1.2
	VSP 4900 Series	VOSS 8.1

Table continues…

Feature	Product	Release introduced
	VSP 7200 Series	VOSS 6.1.2
	VSP 7400 Series	VOSS 8.0
	VSP 8200 Series	VOSS 6.1.2
	VSP 8400 Series	VOSS 6.1.2
	VSP 8600 Series	VSP 8600 6.1
	XA1400 Series	VOSS 8.0.50
TLS client for secure syslog **Note:** VOSS Releases 6.0 and 6.0.1 do not support this feature.	VSP 4450 Series	VOSS 5.1.2
	VSP 4900 Series	VOSS 8.1
	VSP 7200 Series	VOSS 5.1.2
	VSP 7400 Series	VOSS 8.0
	VSP 8200 Series	VOSS 5.1.2
	VSP 8400 Series	VOSS 5.1.2
	VSP 8600 Series	Not Supported
	XA1400 Series	VOSS 8.0.50

Logs and Traps Fundamentals

This section details SNMP traps and log files, available as part of the switch System Messaging Platform.

Overview of Traps and Logs

System log messaging

On a UNIX-based management platform, you can use system log (syslog) messaging to manage event messages. The switch 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 the switch that runs in a network accessible to the workstation.

The remote UNIX management workstation performs the following actions:

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

Log consolidation

The switch 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:

- Extensible Authentication Protocol (EAP)
- Remote Authentication Dial-in User Service (RADIUS)
- Remote Monitoring (RMON)
- Web
- 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)
- policy
- Simple Network Management Protocol (SNMP) log

The switch can send information in the system log file, including CLI command log and the SNMP operation log, to a syslog server.

View logs for CLILOG module to track all CLI commands executed and for fault management purposes. The CLI 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 CLI Log and SNMP Log information regardless of the logging level you configure. 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 in EDM, under the System Log Table tab, you must select either IPv4 or IPv6.

Log messages with enhanced secure mode

Enhanced secure mode allows the system to provide role-based access levels, stronger password requirements, and stronger rules on password length, password complexity, password change intervals, password reuse, and password maximum age use. If you enable enhanced secure mode, the system encrypts the entire log file.

With enhanced secure mode enabled, only individuals in the administrator or auditor role can view log files to analyze switch access and configuration activity. However, no access level role can modify the content of the log files, not even the administrator or the auditor access level roles. The administrator has access to the **remove** and **delete** commands.

If you enable enhanced secure mode, you cannot access the following commands for log files at any role-based access level:

- **more**
- **edit**
- **rename**
- **copy**

If someone attempts to access a log file with the preceding commands, an information and warning message displays on the screen.

The following table summarizes log file command access based on role-based access levels.

Table 5: Log commands accessible for various users

Access level role	Commands
Administrator	The **remove** and **delete** commands.
No user at any access level.	The following commands: • **more** • **edit** • **rename** • **copy**
Administrator	All configuration commands can be accessed only by the individual in the administrator role, other than the preceding commands.
Administrator and auditor	All show commands for log files.
All users (Administrator, auditor, security, privilege, operator)	All show commands for log configurations.

With enhanced secure mode enabled, authorized users can use SFTP to transfer files to a remote server with the content encrypted.

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 the switch to send SNMP traps to a server after a port is unplugged or if a power supply fails.

This document only describes SNMP commands related to traps. For more information about how to configure SNMP community strings and related topics, see Configuring Security for VOSS.

Secure Syslog

Syslog is a standard used to send event log messages to devices within a network. The switch sends event messages to a logging server called syslog server. The syslog server stores the log messages and displays them for event reporting. Syslog messages are used for monitoring system activities and troubleshooting.

The secure syslog feature adds security and authenticated access to the plain text event log messages that are communicated between a remote syslog server and a syslog client. The secure syslog feature helps prevent unauthorized access to confidential data transmitted on an unsecured communication channel between a remote syslog server and client.

To implement the security, this feature employs port forwarding using the Transport Layer Security (TLS) to provide encrypted communication between a syslog server and client.

After starting the syslog server, to ensure authentication, you must setup a remote port forwarding connection to connect the switch with a remote TLS Server.

TLS client for secure syslog

The syslog server is installed on a host that serves as a TLS Server. The switch plays the role of a TLS client for secure syslog. A TLS handshake is initiated between the syslog server and the switch. The syslog server transmits a certificate which has a subject common name and an optional subject alternative name (SAN). The subject common name is always present in the certificate but the SAN is optional. The server-cert-name must match the SAN name, if present in the certificate. If the SAN name is not present, it must match the subject common name. Otherwise, TLS negotiation fails and the connection to the server is closed. If the server-cert-name part is not configured, this check is not done.

Once the TLS handshake is successful, the log messages sent from the switch to the syslog server are encrypted. The syslog server decrypts these messages using a private key. The server then stores the messages or forwards them to other servers.

This feature supports the Rsyslog, which is a Linux based open source syslog server for TLS tunneling.

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 SNMP agents and managers use 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:

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

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

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

Log Message Format

The log messages for the switch 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:

- CPU slot number—Indicates the CP slot where the command is logged.
- 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].
- hostname—The Hostname from which the message is generated.
- event code—Precisely identifies the event reported.
- alarm code—Specifies the alarm code.
- 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.
- module name—Identifies the software module or hardware from which the log is generated.
- severity level—Identifies the severity of the message.
- sequence number—Identifies a specific CLI command.
- context—Specifies the type of the session used to connect to the switch. If the session is a remote session, the remote IP address is identified.
- user name—Specifies the user name used to login to the switch.
- CLI command—Specifies the commands typed during the CLI session. The system logs anything type during the CLI session as soon as the user presses the `Enter` key.

The following messages are examples of an informational message for CLILOG:

```
CP1  [07/18/14 13:23:11.253] 0x002c0600 00000000 GlobalRouter CLILOG INFO    13   TELNET:
192.0.2.200 rwa show log file name-of-file log.40300001.1806

CP1  [07/18/14 13:24:19.739] 0x002c0600 00000000 GlobalRouter CLILOG INFO    15 TELNET:
192.0.2.200 rwa term more en

 CP1  [07/18/14 13:24:22.577] 0x002c0600 00000000 GlobalRouter CLILOG INFO    16 TELNET:
192.0.2.200 rwa show log

CP1  [01/12/70 15:13:59.056] 0x002c0600 00000000 GlobalRouter CLILOG INFO     5 TELNET:
198.51.100.108 rwa syslog host 4

CP1  [01/12/70 15:13:35.520] 0x002c0600 00000000 GlobalRouter CLILOG INFO     4 TELNET:
198.51.100.108 rwa syslog host enable

CP1  [01/12/70 15:13:14.576] 0x002c0600 00000000 GlobalRouter CLILOG INFO     3 TELNET:
198.51.100.108 rwa show syslog

CP1  [01/12/70 15:12:44.640] 0x002c0600 00000000 GlobalRouter CLILOG INFO     2 TELNET:
198.51.100.108 rwa show logging file tail
```

The following messages are examples of an informational message for SNMPLOG:

```
CP1  [05/07/14 10:24:05.468] 0x002c4600 00000000 GlobalRouter SNMPLOG INFO     1
ver=v2c  public  rcVlanPortMembers.2 =

CP1  [05/07/14 10:29:58.133] 0x002c4600 00000000 GlobalRouter SNMPLOG INFO     2
ver=v2c  public  rcVlanPortMembers.2 =

CP1  [05/07/14 10:30:20.466] 0x002c4600 00000000 GlobalRouter SNMPLOG INFO    3  ver=v2c
public  rcVlanPortMembers.1 =
```

The following messages are examples of an informational message for system logs:

```
CP1  [07/24/14 18:04:08.304] 0x00000670 00000000 GlobalRouter SW INFO Basic license
supports all features on this device
CP1  [07/24/14 18:04:10.651] 0x00034594 00000000 GlobalRouter SW INFO System boot
CP1  [07/24/14 18:04:10.651] 0x00034595 00000000 GlobalRouter SW INFO VSP-8200 System
Software Release 0.0.0.0_B553
CP1  [07/24/14 18:04:10.779] 0x00010774 00000000 GlobalRouter HW INFO Detected 8 284XSQ
chassis
CP1  [07/24/14 18:04:10.779] 0x0001081c 00400010.2 DYNAMIC SET GlobalRouter HW INFO Slot
2 is initializing.
CP1  [07/24/14 18:04:10.780] 0x0001081c 00400010.1 DYNAMIC SET GlobalRouter HW INFO Slot
1 is initializing.
CP1  [07/24/14 18:04:10.810] 0x00010729 00000000 GlobalRouter HW INFO Detected 8284XSQ
Power Supply in slot PS 1. Adding 800 watts to available power
CP1  [07/24/14 18:04:10.811] 0x00010830 00000000 GlobalRouter HW INFO Detected 8242XSQ
module (Serial#: SDNIV84Q2013) in slot 2
```

The system encrypts AP information before writing it to the log file.

The encrypted information in a log file is for debugging purposes. Only a Customer Service engineer can decrypt the encrypted information in a log file. CLI commands display the logs without the encrypted information. Do not edit the log file.

The following table describes the system message severity levels.

Table 6: Severity levels

Severity level	Definition
EMERGENCY	A panic condition that occurs when the system becomes unusable. A severity level of emergency is usually a condition where multiple applications or servers are affected. You must correct a severity level of emergency immediately.
ALERT	Any condition requiring immediate attention and correction. You must correct a severity level of alert immediately, but this level usually indicates failure of a secondary system, such as an Internet Service Provider connection.
CRITICAL	Any critical conditions, such as a hard drive error.
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. An indication that an error can occur if action is not taken within a given amount of time.
NOTIFICATION	Significant event of a normal nature. An indication that unusual, but not error, conditions have occurred. No immediate action is required.
INFO	Information only. No action is required.
DEBUG	Message containing information useful for debugging.
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 in EDM, under the System Log Table tab, you must select either IPv4 or IPv6.

Internally, the switch has four severity levels for log messages: INFO, WARNING, ERROR, and FATAL. The system log supports eight different severity levels:

- Debug
- Info
- Notice
- Warning

- Critical
- Error
- Alert
- Emergency

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

Table 7: Default and system log severity level mapping

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

Log Files

The log file captures hardware and software log messages, and alarm messages. The switch logs to internal flash.

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 are 01. 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 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. And once the maximum configured size is reached, system

continues to create a new log file with incremental sequence number on the internal flash for logging.

Log File Transfer

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

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

If log file transfer is unsuccessful, the system keeps the old log files on 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.

With enhanced secure mode enabled, authorized users can use SFTP to transfer files to a remote server with the content encrypted.

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 are 01. 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 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. The switch does not support the minimum size and percentage of free disk space parameters. The internal flash must be less than 75% full for the system to log a file. If the internal flash is more than 75% full, logging to a file stops to prevent exhausting disk space.

Log file transfer using a wildcard filename

File transfers using SFTP require file permissions.

Use the command **`attribute WORD<1-99> [+/-] R`** to change the permissions of a file.

To change permissions for all log files, use the wildcard filename **`log.*`**. Using the command in the wildcard form **`attribute log.* [+/-]R`** changes permissions for log files with names that begin with the characters “log.”.

Important:

You cannot use a wildcard pattern other than **`log.*`** for this command.

Email Notification

The switch can send email notification for failed components or other critical log-event conditions. The switch can also send periodic health status notifications.

Enable and configure a Simple Mail Transfer Protocol (SMTP) client on the switch for one SMTP server by specifying the server hostname or IPv4 address. To use a hostname, you must also configure a Domain Name System (DNS) client on the switch.

You must configure at least one email recipient and can create a maximum of five email recipients.

The switch can periodically send general health status notifications. Status email messages include information about the following items:

- General switch
- Chassis
- Card
- Temperature
- Power supplies
- Fans

- LEDs
- System errors
- Port lock
- Message control
- Operational configuration changes
- Current Uboot
- Port interfaces
- Port statistics

The switch maintains a default list of event IDs for which it generates an email notification. You can add specific event IDs to this list. To see the default list of event IDs, run the **`show smtp event-id`** command.

The following example shows an email that the switch sends for log events.

```
Subject: Logs from LabSwitch - 50712100008
From: <LabSwitch@default.com>
To: <test1@default.com>
CP1  [08/04/15 21:48:04.527:UTC] 0x00004603 00400003.67108870 DYNAMIC CLEAR GlobalRouter
SNMP INFO 2k card up(CardNum=1 AdminStatus=1 OperStatus=1)
CP1  [08/04/15 21:48:04.527:UTC] 0x00004603 00400003.67108870 DYNAMIC CLEAR GlobalRouter
SNMP INFO 2k card up(CardNum=1 AdminStatus=1 OperStatus=1)
CP1  [08/04/15 21:48:04.527:UTC] 0x00004603 00400003.67108870 DYNAMIC CLEAR GlobalRouter
SNMP INFO 2k card up(CardNum=2 AdminStatus=1 OperStatus=1)
CP1  [08/04/15 21:50:03.511:UTC] 0x00088524 00000000 GlobalRouter SW INFO Boot sequence
successful
```

If you enable the SMTP client but the switch cannot reach the SMTP server, the switch generates an alarm. The switch holds log and status information in a queue until the connection with the SMTP server is restored. The message queue holds a maximum of 2,000 messages. If the queue fills, the switch drops new messages.

The following text is an example of the alarm that the switch generates when it cannot connect to the SMTP server.

```
CP1  [06/10/15 19:27:07.901:EST] 0x00398600 0e600000 DYNAMIC SET GlobalRouter SMTP
WARNING SMTP: Unable to establish connection with server: mailhost.usae.company.com, port:
25
```

If the switch cannot establish a connection to the SMTP server, verify that the server IP address or hostname, and the TCP port are correct. If you specify the server hostname, confirm that the IP address for the DNS server is correct. Check for network issues such as unplugged cables.

If the SMTP server rejects the email message, the switch generates a log message.

Log Configuration Using CLI

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

Configuring a UNIX System Log and Syslog Host

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

About this task

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

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enable the system log:

   ```
   syslog enable
   ```

3. Specify the IP header in syslog packets:

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

4. Configure the maximum number of syslog hosts:

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

5. Create the syslog host:

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

6. Configure the IP address for the syslog host:

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

7. Enable the syslog host:

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

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

8. View the configuration to ensure it is correct:

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

Example

```
Switch:1(config)# syslog enable
Switch:1(config)# syslog host 7 address 192.0.2.1
Switch:1(config)# syslog host 7 enable
Switch:1(config)#show syslog host 7
                Id : 7
```

```
            IpAddr : 192.0.2.1
           UdpPort : 514
          Facility : local7
          Severity : info|warning|error|fatal
   MapInfoSeverity : info
MapWarningSeverity : warning
  MapErrorSeverity : error
    MapMfgSeverity : notice
  MapFatalSeverity : emergency
            Enable : true
SecureForwardingMode: none
          Tcp Port : 1025
```

```
Switch:1(config)#show syslog

Enable    : true
Max Hosts : 5
OperState : active
header : default
Total number of configured hosts : 3
Total number of enabled hosts : 1
Configured host : 7  8  9
Enabled host : 7
```

Variable Definitions

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

Variable	Value
enable	Enables the sending of syslog messages on the device. 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>	Specifies the IP header in syslog packets to circuitless-ip or default. • 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. • 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.

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.

Table continues…

Variable	Value
address WORD <*0–46*>	Configures a host location for the syslog host. WORD <*0–46*> is the IPv4 or IPv6 address of the UNIX system syslog host in the format A.B.C.D or x:x:x:x:x:x: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.
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. You can specify up to four severity levels in the same command string. 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 Secure Forwarding

Configuring secure forwarding includes setting the mode for the particular syslog host and setting the TCP port through which the logs are sent to the syslog server.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Create the syslog host:

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

   Use the no operator before this parameter, that is, `no syslog host` to delete a host instance.

3. Configure an IP address for the syslog host:

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

4. Enable the syslog host:

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

5. Enable syslog globally:

   ```
   syslog enable
   ```

6. Set the mode for secure forwarding on the host:

   ```
   syslog host <1-10> secure-forwarding mode <none | tls [server-cert-name WORD<1-64>]>
   ```

7. Set the TCP port:

   ```
   syslog host <1-10> secure-forwarding tcp-port <1025-49151>
   ```

8. Display the secure forwarding configured values:

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

9. **(Optional)** Remove the server certificate name:

   ```
   no syslog host <1-10> secure-forwarding mode tls server-cert-name
   ```

10. **(Optional)** Set secure-forwarding mode to none for a particular host:

    ```
    default syslog host <1-10> secure-forwarding mode
    ```

Next steps

After configuring secure forwarding on the switch, set the syslog server to be able to see the log messages on the interactive syslog viewer.

- For TLS secure syslog, on the rsyslog server, configure the server to use TLS method and install the root certificate on the server in the switch.

Variable Definitions

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

Variable	Value
host <1–10>	Specifies the ID for the syslog host. The range is 1–10.
address WORD<0–46>	Configures a host location for the syslog host. WORD <0–46> is the IPv4 or IPv6 address of the UNIX system syslog host in the format A.B.C.D or x:x:x:x:x:x: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 CLI.
enable	Enables the syslog host. Use the no operator before this parameter, no syslog host enable to disable syslog host. The default is disabled.
secure-forwarding	Adds protected syslog using remote port forwarding for host.

Use the data in the following table to use the **`syslog host secure-forwarding`** command.

Variable	Value
host <1–10>	Creates and configures a host instance. Use the no operator before this parameter, no syslog host to delete a host instance.
mode <none | tls [server-cert-name WORD<1-64>]>	Specifies the mode of secure forwarding of syslog on the host. The default mode is none, that is, tls mode is disabled by default. **Note:** Certificate validation is done only if the server-cert-name is configured.
tcp-port <1025–49151>	Set tcp-port for secure forwarding of syslog for host. The default tcp-port is 1025. To set the TCP port to default value, use command **`default syslog host <1–10> secure-forwarding tcp-port`**. **Important:** The tcp-port 6000 cannot be used, as it is used as an internal port for Internal Spanning Tree (IST).

Installing Root Certificate for Syslog Client

Use the following procedure to install a root certificate for a syslog client.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Install a root certificate on the store:

   ```
   syslog root-cert install-filename <file-name>
   ```

   The certificate is installed in folder: `/intflash/.cert/.syslogrootinstalledcert/`.

   Note:

   The offline root certificate for TLS syslog must be kept in folder: `/intflash/.cert/..syslogofflinerootcert/`.

3. Uninstall a root certificate from the store:

   ```
   no syslog root-cert install-filename <file-name>
   ```

4. To display the installed syslog server root certificate file:

   ```
   show syslog root-cert-file
   ```

Variable Definition

Use the data in the following table to use the **`syslog root-cert`** command.

Variable	Value
install-filename *WORD<1–128>*	Specifies the name of the root certificate to be installed on the store.

Configuring Logging

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

About this task

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 CLI Log and SNMP Log information regardless of the logging level you configure. This is not the case for other INFO messages.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Define which messages to log:

   ```
   logging level <0-4>
   ```

3. Write the log file from memory to a file:

   ```
   logging write WORD<1-1536>
   ```

4. Show logging on the screen:

   ```
   logging screen
   ```

Example

```
Switch:1(config)#logging level 0
Switch:1(config)#logging write log2
Switch:1(config)#logging screen
```

Variable Definitions

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

Variable	Value
level <0-4>	Shows and configures the logging level. The level is one of the following values: • 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
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 or TFTP server. <1–10> specifies the file ID. The address {A.B.C.D} option specifies the IP address. The filename-prefix WORD<0–200> option sets the filename prefix for the log file at the remote host.
write WORD<1-1536>	Writes the log file with the designated string. *WORD<1-1536>* is the string or command that you append to the log file. If the string contains spaces, you must enclose the string in quotation marks (").

Configuring the Remote Host Address for Log Transfer

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

Before you begin

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

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure the remote host address for log transfer:

   ```
   logging transferFile {1-10} address {A.B.C.D} [filename—prefix WORD<0-200>]
   ```

Example

```
Switch:1(config)# logging transferFile 1 address 192.0.2.10
```

Variable Definitions

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

Variable	Value
1–10	Specifies the file ID to transfer.
address {A.B.C.D}	Specifies the IP address of the host to which to transfer the log file. The remote host must be reachable or the configuration fails.
filename-prefix *WORD<0–200>*	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

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

About this task

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

Configure logging to a flash file at all times as a best practice.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enable system logging to a PC card file:

   ```
   boot config flags logging
   ```

3. Configure the logfile parameters:

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

Example

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

Variable Definitions

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

Variable	Value
flags logging	Enables or disables logging to a flash file. 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

Table continues…

Variable	Value
	last three bytes of the chassis base MAC address. The next two characters specify the slot number. The last three characters denote the sequence number of the log file.
logfile <64-500> <500-16384> <10-90>	Configures the following logfile parameters: • *<64-500>* specifies the minimum free memory space on the external storage device from 64–500 KB. The switch 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, ranging from 10–90 percent, of space on the external storage device the logfile can use. The switch does not support this parameter.

Configuring System Message Control

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

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure system message control action:

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

3. Configure the maximum number of messages:

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

4. Configure the interval:

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

5. Enable message control:

   ```
   sys msg-control
   ```

Example

```
Switch:1(config)#sys msg-control action suppress-msg
Switch:1(config)#sys msg-control max-msg-num 10
Switch:1(config)#sys msg-control control-interval 15
Switch:1(config)#sys msg-control
```

Variable Definitions

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

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

Extending System Message Control

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

About this task

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

Procedure

1. Enter Global Configuration mode:

   ```
   enable

   configure terminal
   ```

2. Configure the force message control option:

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

Example

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

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

Variable Definitions

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

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

Viewing Logs

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

About this task

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

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Show log information:

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

Example

Display log file information:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#show logging file
CP1  [02/06/15 22:38:20.678:UTC] 0x00270428 00000000 GlobalRouter SW INFO Lifecy
cle: Start
CP1  [02/06/15 22:38:21.770:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s sockserv started, pid:4794
CP1  [02/06/15 22:38:21.771:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s oom95 started, pid:4795
CP1  [02/06/15 22:38:21.771:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s oom90 started, pid:4796
CP1  [02/06/15 22:38:21.772:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s imgsync.x started, pid:4797
CP1  [02/06/15 22:38:22.231:UTC] 0x0026452f 00000000 GlobalRouter SW INFO No pat
ch set.
CP1  [02/06/15 22:38:22.773:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s logServer started, pid:4840
CP1  [02/06/15 22:38:22.774:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s trcServer started, pid:4841
CP1  [02/06/15 22:38:22.774:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s oobServer started, pid:4842
```

```
CP1  [02/06/15 22:38:22.775:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s cbcp-main.x started, pid:4843
CP1  [02/06/15 22:38:22.776:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s rssServer started, pid:4844
CP1  [02/06/15 22:38:22.777:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s dbgServer started, pid:4845
CP1  [02/06/15 22:38:22.777:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s dbgShell started, pid:4846
CP1  [02/06/15 22:38:22.778:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s coreManager.x started, pid:4847
CP1  [02/06/15 22:38:22.779:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s ssio started, pid:4848
CP1  [02/06/15 22:38:22.780:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s hckServer started, pid:4849
CP1  [02/06/15 22:38:22.780:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s remCmdAgent.x started, pid:4850
CP1  [02/06/15 22:38:24.717:UTC] 0x000006cc 00000000 GlobalRouter SW INFO rcStar
t:  FIPS Power Up Self Test SUCCESSFUL - 0
CP1  [02/06/15 22:38:24.718:UTC] 0x000006c2 00000000 GlobalRouter SW INFO rcStar
t:  Security Stack Init SUCCESSFUL - 0
CP1  [02/06/15 22:38:24.718:UTC] 0x000006c3 00000000 GlobalRouter SW INFO rcStar
t:  IPSEC Init SUCCESSFUL
CP1  [02/06/15 22:38:24.718:UTC] 0x000006bf 00000000 GlobalRouter SW INFO rcStar
t: Security Stack Log init SUCCESSFUL - 0
CP1  [02/06/15 22:38:26.111:UTC] 0x000005c0 00000000 GlobalRouter SW INFO Licens
eLoad = ZERO, loading premier license for developer debugging
IO1  [02/06/15 22:38:26.960:UTC] 0x0011054a 00000000 GlobalRouter COP-SW INFO De
tected Master CP in slot 1

--More-- (q = quit)

Switch:1(config)#show logging file module SNMP
CP1  [02/06/15 22:39:58.530:UTC] 0x00004595 00000000 GlobalRouter SNMP INFO Boot
ed with file
CP1  [02/06/15 22:39:59.547:UTC] 0x00004603 00400003.67108870 DYNAMIC CLEAR Glob
alRouter SNMP INFO 2k card up(CardNum=1 AdminStatus=1 OperStatus=1)
CP1  [02/06/15 22:39:59.547:UTC] 0x00004603 00400003.67108870 DYNAMIC CLEAR Glob
alRouter SNMP INFO 2k card up(CardNum=2 AdminStatus=1 OperStatus=1)
CP1  [02/06/15 22:39:59.547:UTC] 0x00004603 00400003.67108870 DYNAMIC CLEAR Glob
alRouter SNMP INFO 2k card up(CardNum=3 AdminStatus=1 OperStatus=1)
CP1  [02/06/15 22:40:45.839:UTC] 0x000045e5 00400005 DYNAMIC SET GlobalRouter SN
MP INFO Sending Cold-Start Trap
```

Variable Definitions

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

Variable	Value
alarm	Displays alarm log entries.
CPU WORD <0-100>	Filters and lists the logs according to the CPU that generated the message. Specify a string length of 0-25 characters. To specify multiple filters, separate each CPU by the vertical bar (|), for example, CPU1|CPU2.
detail	Displays CLI and SNMP logging information.
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, HW, MLT, FILTER, QOS, CLILOG, SW, CPU, IP, VLAN, IPMC, and

Table continues…

Variable	Value
	SNMPLOG. To specify multiple filters, separate each category by the vertical bar (|), for example, |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 to 99 characters. If you enable enhanced secure mode, the system encrypts the entire log file. After you use the **`show log file name-of-file WORD<1-99>`** command, the system takes the encrypted log file name as input, then decrypts it, and prints the output to the screen. You can then redirect the decrypted output to a file that you can store onto the flash. If enhanced secure mode is disabled, the system only encrypts the proprietary portion of the log file.
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.
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 CLI Logging

Use CLI logging to track all CLI commands executed and for fault management purposes. The CLI commands are logged to the system log file as CLILOG module.

About this task

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 CLI Log and SNMP Log information regardless of the logging level you configure. This is not the case for other INFO messages.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enable CLI logging:

   ```
   clilog enable
   ```

3. **(Optional)** Disable CLI logging:

   ```
   no clilog enable
   ```

4. Ensure that the configuration is correct:

   ```
   show clilog
   ```

5. View the CLI log:

   ```
   show logging file module clilog
   ```

Example

Enable CLI logging, and view the CLI log:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#clilog enable
Switch:1(config)#show logging file module clilog
CP1  [02/13/13 17:27:25.956] 0x002c0600 00000000 GlobalRouter CLILOG INFO     1 CONSOLE
rwa show snmp-server host
CP1  [02/13/13 17:28:10.100] 0x002c0600 00000000 GlobalRouter CLILOG INFO     2 CONSOLE
rwa show snmp-server notif
CP1  [02/13/13 17:28:45.732] 0x002c0600 00000000 GlobalRouter CLILOG INFO     3 CONSOLE
rwa snmp-server force-trap
CP1  [02/13/13 17:29:30.628] 0x002c0600 00000000 GlobalRouter CLILOG INFO     4 CONSOLE
rwa show logging file modug
CP1  [02/14/13 19:39:11.648] 0x002c0600 00000000 GlobalRouter CLILOG INFO     5 CONSOLE
rwa ena
CP1  [02/14/13 19:39:13.420] 0x002c0600 00000000 GlobalRouter CLILOG INFO     6 CONSOLE
rwa conf t
CP1  [02/14/13 19:49:21.044] 0x002c0600 00000000 GlobalRouter CLILOG INFO     7 CONSOLE
rwa filter acl 2 enable
CP1  [02/14/13 19:50:08.540] 0x002c0600 00000000 GlobalRouter CLILOG INFO     8 CONSOLE
rwa filter acl 2 type inpo1
CP1  [02/14/13 19:50:38.444] 0x002c0600 00000000 GlobalRouter CLILOG INFO     9 CONSOLE
rwa filter acl 2 type inpoe
CP1  [02/14/13 19:50:52.968] 0x002c0600 00000000 GlobalRouter CLILOG INFO    10 CONSOLE
rwa filter acl enable 2
CP1  [02/14/13 19:51:08.908] 0x002c0600 00000000 GlobalRouter CLILOG INFO    11 CONSOLE
rwa filter acl 2 enable
CP1  [02/15/13 06:50:25.972] 0x002c0600 00000000 GlobalRouter CLILOG INFO    14 CONSOLE
rwa ena
CP1  [02/15/13 06:50:30.288] 0x002c0600 00000000 GlobalRouter CLILOG INFO    15 CONSOLE
rwa conf t
CP1  [02/15/13 06:50:39.412] 0x002c0600 00000000 GlobalRouter CLILOG INFO    16 CONSOLE
rwa show vlan basic
CP1  [02/15/13 06:51:09.488] 0x002c0600 00000000 GlobalRouter CLILOG INFO    17 CONSOLE
rwa show isis spbm
CP1  [02/15/13 06:56:00.992] 0x002c0600 00000000 GlobalRouter CLILOG INFO    19 CONSOLE
rwa spbm 23 b-vid 2 primar1
CP1  [02/15/13 06:56:59.092] 0x002c0600 00000000 GlobalRouter CLILOG INFO    20 CONSOLE
rwa show isis
CP1  [02/15/13 07:10:54.928] 0x002c0600 00000000 GlobalRouter CLILOG INFO    21 CONSOLE
rwa show isis interface
CP1  [02/15/13 07:12:33.404] 0x002c0600 00000000 GlobalRouter CLILOG INFO    22 CONSOLE
rwa show isis spbm
CP1  [02/15/13 07:45:28.596] 0x002c0600 00000000 GlobalRouter CLILOG INFO    23 CONSOLE
rwa ena
CP1  [02/15/13 07:45:30.236] 0x002c0600 00000000 GlobalRouter CLILOG INFO    24 CONSOLE
rwa conf t
CP1  [02/15/13 07:46:29.456] 0x002c0600 00000000 GlobalRouter CLILOG INFO    25 CONSOLE
rwa interface gigabitEther0
CP1  [02/15/13 07:47:28.476] 0x002c0600 00000000 GlobalRouter CLILOG INFO    26 CONSOLE
```

```
rwa encapsulation dot1q

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

Variable Definitions

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

Variable	Value
enable	Activates CLI logging. To disable, use the **no clilog enable** command.

Configuring Email Notification

Configure the SMTP feature to generate email notifications for component failures, critical conditions, or general system health status.

About this task

The SMTP feature is disabled by default.

Before you begin

- To identify the SMTP server by hostname, you must first configure a DNS client on the switch. For more information about how to configure a DNS client, see Administering VOSS.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure the TCP port the client uses to open a connection with the SMTP server:

   ```
   smtp port <1-65535>
   ```

   Note:

   The port you specify must match the port that the SMTP server uses.

3. Configure email recipients:

   ```
   smtp receiver-email add WORD<3-1274>
   smtp receiver-email remove WORD<3-1274>
   ```

   Note:

   You must configure at least one recipient.

4. Configure the SMTP server hostname or IPv4 address:

   ```
   smtp server WORD<1-256>
   ```

5. **(Optional)** Configure a sender email address:

```
smtp sender-email WORD<3-254>
```

6. **(Optional)** Add or remove log events to the default list that generate email notification:

   ```
   smtp event-id add WORD<1-1100>
   ```

   ```
   smtp event-id remove WORD<1-1100>
   ```

7. **(Optional)** Configure the status update interval:

   ```
   smtp status-send-timer <0 | 30-43200>
   ```

8. Enable the SMTP client:

   ```
   smtp enable
   ```

9. Configure an SMTP domain name:

   ```
   smtp domain-name WORD<1-254>
   ```

10. Verify the configuration:

    ```
    show smtp [event-id]
    ```

Example

Configure the SMTP client to use TCP port 26 to communicate with an SMTP server that is using port 26. Add two receiver email addresses, configure the server information using an IPv4 address, and enable the SMTP feature. Finally, configure an SMTP domain name, and then verify the configuration.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch:1(config)#smtp port 26
Switch:1(config)#smtp receiver-email add test1@default.com,test2@default.com
Switch:1(config)#smtp server 192.0.2.1
Switch:1(config)#smtp enable
Switch:1(config)#smtp domain-name test mailer
Switch:1(config)#show smtp
================================================================================
                                SMTP Information
================================================================================
        SMTP Status:  Enabled
     Server Address:  192.0.2.1
        Server Port:  26
  Status send Timer:  30 (seconds)
       Sender Email:  LabSwitch@default.com
        Domain Name:  test mailer
    Receiver Emails:  test1@default.com
                      test2@default.com
```

Add an event ID to the list for which the switch sends email notification on a log event. Verify the configuration.

```
Switch:1(config)#smtp event-id add 0x0000c5ec
Switch:1(config)#show smtp event-id
================================================================================
                            SMTP Event IDs Information
================================================================================
Log Event IDs: (total: 51)
```

```
            0x000045e3,0x00004602,0x00004603,0x0000c5ec,0x000106ce,0x000106cf
            0x000106d0,0x000106d1,0x000106d2,0x000106d4,0x000106d8,0x000106d9
            0x000106da,0x000106f8,0x000106f9,0x000106fb,0x00010775,0x00010776
            0x000107f5,0x000107f6,0x000305c8,0x000305ca,0x000305f1,0x00030637
            0x00040506,0x00040507,0x00040508,0x00040509,0x000646da,0x000646db
            0x00088524,0x000d8580,0x000d8586,0x000d8589,0x000e4600,0x000e4601
            0x000e4602,0x000e4603,0x000e4604,0x000e4605,0x000e4606,0x000e4607
            0x000e4608,0x000e4609,0x001985a0,0x00210587,0x00210588,0x00210595
            0x00210596,0x0027458a,0x0027458d

Default Event IDs: (total: 50)
            0x000045e3,0x00004602,0x00004603,0x000106ce,0x000106cf,0x000106d0
            0x000106d1,0x000106d2,0x000106d4,0x000106d8,0x000106d9,0x000106da
            0x000106f8,0x000106f9,0x000106fb,0x00010775,0x00010776,0x000107f5
            0x000107f6,0x000305c8,0x000305ca,0x000305f1,0x00030637,0x00040506
            0x00040507,0x00040508,0x00040509,0x000646da,0x000646db,0x00088524
            0x000d8580,0x000d8586,0x000d8589,0x000e4600,0x000e4601,0x000e4602
            0x000e4603,0x000e4604,0x000e4605,0x000e4606,0x000e4607,0x000e4608

            0x000e4609,0x001985a0,0x00210587,0x00210588,0x00210595,0x00210596
            0x0027458a,0x0027458d

Remove From Default: (total: 0)

Add List: (total: 1)
            0x0000c5ec
```

Variable Definitions

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

Variable	Value
<1–65535>	Specifies the TCP port on the switch that the SMTP client uses to communicate with the SMTP server. The default value is 25. **Note:** You must disable the SMTP feature before you can change an existing SMTP port configuration. The port you specify must match the port that the SMTP server uses.

Use the data in the following table to use the **`smtp receiver-email`** command.

Variable	Value
add *WORD<3-1274>*	Adds an email address to the recipient list. The recipients receive the email notification generated by the switch. You must configure at least one email recipient and can create a maximum of five email recipients. You can specify multiple addresses in a single command by separating them with a comma.

Table continues…

Variable	Value
	You cannot use quotation marks (“) or commas (,) in email addresses. Other restrictions for the format of the email address follow RFC 5321. The maximum length for the address is 254 characters.
remove *WORD<3-1274>*	Removes an email address from the recipient list. The recipients receive the email notification generated by the switch. You can specify multiple addresses in a single command by separating them with a comma. You cannot use quotation marks (“) or commas (,) in email addresses. Other restrictions for the format of the email address follow RFC 5321. The maximum length for the address is 254 characters.

Use the data in the following table to use the **`smtp server`** command.

Variable	Value
WORD<1-256>	Specifies the SMTP server address. You can use either a hostname or IPv4 address. If you use a hostname, you must configure the DNS client on the switch.

Use the data in the following table to use the **`smtp sender-email`** command.

Variable	Value
WORD<3-254>	Specifies the email address that appears in the From field of the message that the switch generates. By default, the switch uses <*SystemName*>@default.com.

Use the data in the following table to use the **`smtp event-id`** command.

Variable	Value
add *WORD<1-1100>*	Adds a log event to the list of events that generate email notification. You can specify multiple event IDs in a single command by separating them with a comma. The event ID can be up to 10 digits in hexadecimal format.
remove *WORD<1-1100>*	Removes a log event from the list of events that generate email notification. You can specify multiple event IDs in a single command by separating them with a comma.

Table continues…

Variable	Value
	The event ID can be up to 10 digits in hexadecimal format.

Use the data in the following table to use the **`smtp status-send-timer`** command.

Variable	Value
<0 | 30-43200>	Specifies the interval, in seconds, at which the switch sends status information. The default is 30 seconds. A value of 0 means the switch does not send status information.

Use the data in the following table to use the **`smtp domain-name`** command.

Variable	Value
WORD<1-254>	Specifies the SMTP host name or IPv4 address (string length 1–254).

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

Variable	Value
event-id	Shows a list of active event IDs for which the switch generates email notification. The command output includes the default list of IDs and IDs you specifically add or remove.

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 pane, expand the **Configuration** > **Edit** > **Diagnostics** folders.
2. Click **System Log**.
3. In the **System Log** tab, select **Enable**.
4. Configure the maximum number of syslog hosts.

5. Configure the IP header type for the syslog packet.
6. Click **Apply**.

System Log Field Descriptions

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

Name	Description
Enable	Enables or disables the syslog feature. If you select this variable, this feature sends a message to a server on a network that is configured to receive and store diagnostic messages from this device. You can configure the type of messages sent. The default is enabled.
MaxHosts	Specifies the maximum number of remote hosts considered active and can receive messages from the syslog service. The range is 0–10 and the default is 5.
OperState	Specifies the operational state of the syslog service. The default is active.
Header	Specifies the IP header in syslog packets to circuitlessIP or default. • 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. • 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 pane, expand the **Configuration** > **Edit** > **Diagnostics** folders.
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 IPv6 address.
Address	Specifies the IP address of the syslog host. You can log system log messages to external system log hosts with both IPv4 and IPv6 addresses.
UdpPort	Specifies the UDP port to use to send messages to the syslog host (514–530). The default is 514.
Enable	Enables or disables the sending of messages to the syslog host. The default is disabled.
HostFacility	Specifies the syslog host facility used to identify messages (local0 to local7). The default is local7.
Severity	Specifies the message severity for which syslog messages are sent. The default is info|warning|error|fatal.
MapInfoSeverity	Specifies the syslog severity to use for INFO messages. The default is info.
MapWarningSeverity	Specifies the syslog severity to use for WARNING messages. The default is warning.
MapErrorSeverity	Specifies the syslog severity to use for ERROR messages. The default is error.
MapFatalSeverity	Specifies the syslog severity to use for FATAL messages. The default is emergency.
MapMfgSeverity	Specifies the syslog severity to use for Accelar manufacturing messages. The default is notice.
SecureForwardingTcpPort	Specifies the TCP port to use for secure forwarding for a particular host. The default is 1025.
SecureForwardingMode	Enables or disables secure forwarding of syslog over remote port forwarding. The supported values are tls and none. The default is none, which means that secure forwarding is disabled.
SecureForwardingServerCertName	Specifies the server certificate name. Certificate validation is done only if the server certificate name is configured.

Configuring Email Notification

Configure the SMTP feature to generate email notifications for component failures, critical conditions, or general system health status.

About this task

The SMTP feature is disabled by default.

Before you begin

- To identify the SMTP server by hostname, you must first configure a DNS client on the switch. For more information about how to configure a DNS client, see Administering VOSS.

Procedure

1. In the navigation pane, expand the **Configuration** > **Edit** folders.
2. Click **SMTP**.
3. Click the **Globals** tab.
4. In the **ServerAddress** field, configure the SMTP server address.
5. In the **ReceiverEmailsList** field, add email recipients.

   Note:

   You must configure at least one recipient.
6. **(Optional)** In the **SenderEmail** field, configure a sender email address to use an address other than the default.
7. In the **DomainName** field, configure an SMTP domain name.
8. In the **Port** field, configure the TCP port that the client uses to open a connection with the SMTP server.
9. **(Optional)** In the **SystemStatusSendTimer** field, configure the status update interval.
10. Click **enable** to enable the SMTP client.
11. **(Optional)** In the **LogEventIds** field, add or remove log events to the default list that generates an email notification.
12. Click **Apply**.

Globals Field Descriptions

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

Name	Description
ServerAddressType	Specifies the type of server address as either an IPv4 address or a hostname. If you use a hostname, you must configure the DNS client on the switch.

Table continues…

Name	Description
ServerAddress	Specifies the SMTP server address. You can use either a hostname or an IPv4 address. If you use a hostname, you must configure the DNS client on the switch.
ReceiverEmailsList	Specifies the recipient list. The recipients receive the email notification generated by the switch. You must configure at least one email recipient and can create a maximum of five email recipients. You can specify multiple addresses in a single command by separating them with a comma. You cannot use quotation marks (“) or commas (,) in email addresses. Other restrictions for the format of the email address follow RFC5321. The maximum length for the address is 254 characters.
NumOfEmails	Shows the total number of addresses in **ReceiverEmailsList**.
SenderEmail	Specifies the email address that appears in the From field of the message that the switch generates. By default, the switch uses *SystemName*@default.com.
DomainName	Specifies the SMTP domain name. The maximum length is 254 characters.
Port	Specifies the TCP port on the switch that the SMTP client uses to communicate with the SMTP server. The default value is 25. **Note:** You must disable the SMTP feature before you can change an existing SMTP port configuration. The port you specify must match the port that the SMTP server uses.
SystemStatusSendTimer	Specifies the interval, in seconds, at which the switch sends status information. The default is 30 seconds. A value of 0 means the switch does not send status information.
Enable	Enables or disables the SMTP feature. By default, SMTP is disabled.
LogEventIds	Specifies the list of events that generate email notification. You can specify multiple event IDs in a single command by separating them with a comma.

Table continues…

Name	Description
	The event ID can be up to 10 digits in hexadecimal format.
NumOfEventIds	Shows the total number of IDs in **LogEventIds**.
DefaultLogEventIds	Shows the default list of event IDs that generate email notification.
NumOfDefaultEventIds	Shows the total number of IDs in **DefaultLogEventIds**.

SNMP Trap Configuration Using CLI

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 Configuring Security for VOSS.

Configuring an SNMP Host

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. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure an SNMPv1 host:

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

3. Configure an SNMPv2c host:

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

4. Configure an SNMPv3 host:

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

5. Ensure that the configuration is correct:

   ```
   show snmp-server host
   ```

Example

Configure the target table entry. Configure an SNMPv3 host.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#snmp-server host 192.0.2.207 port 162 v2c ReadView inform timeout 1500
retries 3 mms 484
Switch:1(config)#snmp-server host 192.0.2.207 port 163 v3 authPriv Lab3 inform timeout
1500 retries 3
```

Variable Definitions

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

Variable	Value
inform [timeout <1-2147483647>] [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 <1-2147483647> specifies the timeout value in seconds with a range of 1–214748364. 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

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

Before you begin

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

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   ```

```
configure terminal
```

2. Create a new notify filter table:

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

3. Ensure that the configuration is correct:

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

Example

```
Switch:1(config)#snmp-server notify-filter profile3 99.3.6.1.6.3.1.1.4.1
Switch:1(config)#show snmp-server notify-filter

================================================================================
                          Notify Filter Configuration
================================================================================
Profile Name                    Subtree                           Mask
--------------------------------------------------------------------------------
profile1                        +99.3.6.1.6.3.1.1.4.1             0x7f
profile2                        +99.3.6.1.6.3.1.1.4.1             0x7f
profile3                        +99.3.6.1.6.3.1.1.4.1             0x7f
```

Variable Definitions

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

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

Configure an interface to send SNMP traps. If the switch has multiple interfaces, configure the IP interface from which the SNMP traps originate.

Procedure

1. Enter Global Configuration mode:

```
enable

configure terminal
```

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

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

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

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

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

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

Example

```
Switch:1(config)#snmp-server sender-ip 192.0.2.2 192.0.2.5
Switch:1(config)#no snmp-server force-iphdr-sender enable
```

Variable Definitions

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

Variable	Value
authentication-trap enable	Activates the generation of authentication traps.
force-iphdr-sender enable	Automatically configures the SNMP and IP sender to the same value. The default is disabled.
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

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

Before you begin

- You must configure and enable the syslog server.

About this task

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 CLI Log and SNMP Log information regardless of the logging level you set. This is not the case for other INFO messages.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enable SNMP trap logging:

   ```
   snmplog enable
   ```

3. **(Optional)** Disable SNMP trap logging:

   ```
   no snmplog enable
   ```

4. View the contents of the SNMP log:

   ```
   show logging file module snmplog
   ```

Example

Enable SNMP trap logging and view the contents of the SNMP log:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#snmplog enable
Switch:1(config-app)#show logging file module snmp
CP1  [02/06/15 22:39:58.530:UTC] 0x00004595 00000000 GlobalRouter SNMP INFO Boot
ed with file
CP1  [02/06/15 22:39:59.547:UTC] 0x00004603 00400003.67108870 DYNAMIC CLEAR Glob
alRouter SNMP INFO 2k card up(CardNum=1 AdminStatus=1 OperStatus=1)
CP1  [02/06/15 22:39:59.547:UTC] 0x00004603 00400003.67108870 DYNAMIC CLEAR Glob
alRouter SNMP INFO 2k card up(CardNum=2 AdminStatus=1 OperStatus=1)
CP1  [02/06/15 22:39:59.547:UTC] 0x00004603 00400003.67108870 DYNAMIC CLEAR Glob
alRouter SNMP INFO 2k card up(CardNum=3 AdminStatus=1 OperStatus=1)
CP1  [02/06/15 22:40:45.839:UTC] 0x000045e5 00400005 DYNAMIC SET GlobalRouter SN
MP INFO Sending Cold-Start Trap
```

Variable Definitions

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

Variable	Value
enable	Enables the logging of traps. Use the command `no snmplog enable` to disable the logging of traps.

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 Configuring Security for VOSS.

Configuring an SNMP Host Target Address

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

Procedure

1. In the navigation pane, expand the **Configuration** > **Edit** > **SnmpV3** folders.
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.

Table continues…

Name	Description
	ipv6Tdomain specifies the transport type of address is IPv6. The default is ipv4Tdomain.
TAddress	Specifies the transport address in xx.xx.xx.xx:port format, for example: 192.1.2.12:162, where 162 is the trap listening port on the system 192.1.2.12.
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 message size 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 pane, expand the **Configuration** > **Edit** > **SnmpV3** folders.
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.
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 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 pane, expand the **Configuration** > **Edit** > **SnmpV3** folders.
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 **included** or **excluded**.
9. Click **Insert**.

Notify Filter Table Field Descriptions

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

Name	Description
NotifyFilterProfileName	Specifies the name of the filter profile used to generate notifications.
Subtree	Specifies the MIB subtree that, if you combine it with the mask, defines a family of subtrees, which are included in or excluded from the filter profile. For more information, see RFC 2573.
Mask	Specifies the bit mask (in hexadecimal format) that, in combination with the subtree, defines a family of subtrees, which are included in or excluded from the filter profile.
Type	Indicates whether the family of filter subtrees are included in or excluded from a filter. The default is included.

Configuring SNMP Notify Filter Profile Table Parameters

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 pane, expand the **Configuration** > **Edit** > **SnmpV3** folders.
2. Click **Notify Table**.
3. Click the **Notify Filter Profile Table** tab.
4. Click **Insert**.
5. In the **TargetParamsName** box, type a name for the target parameters.
6. In the **NotifyFilterProfileName** box, type a name for the notify filter profile.
7. Click **Insert**.

Notify Filter Profile Table Field Descriptions

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

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

Enabling Authentication Traps

About this task

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

Procedure

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

Error Field Descriptions

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

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

Viewing the Trap Sender Table

About this task

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

Procedure

1. In the navigation pane, expand the **Configuration** > **Edit** folders.
2. Click **Chassis**.
3. Click the **Trap Sender Table** tab.

Trap Sender Table Field Descriptions

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

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

Chapter 5: Connectivity Fault Management

Table 8: Connectivity Fault Management product support

Feature	Product	Release introduced
For configuration details, see Troubleshooting VOSS.		
IEEE 802.1ag Connectivity Fault Management (CFM): • Layer 2 Ping • TraceRoute • TraceTree	VSP 4450 Series	VSP 4000 4.0
	VSP 4900 Series	VOSS 8.1
	VSP 7200 Series	VOSS 4.2.1
	VSP 7400 Series	VOSS 8.0
	VSP 8200 Series	VSP 8200 4.0
	VSP 8400 Series	VOSS 4.2
	VSP 8600 Series	VSP 8600 4.5
	XA1400 Series	VOSS 8.0.50
CFM configuration on C-VLANs	VSP 4450 Series	VSP 4000 4.0
	VSP 4900 Series	Not Supported
	VSP 7200 Series	Not Supported
	VSP 7400 Series	Not Supported
	VSP 8200 Series	Not Supported
	VSP 8400 Series	Not Supported
	VSP 8600 Series	Not Supported
	XA1400 Series	Not Supported

CFM Fundamentals

The Shortest Path Bridging MAC (SPBM) network needs a mechanism to debug connectivity issues and to isolate faults. This is performed at Layer 2, not Layer 3. Connectivity Fault Management (CFM) operates at Layer 2 and provides an equivalent of ping and traceroute. To support troubleshooting of the SPBM cloud, the switch supports a subset of CFM functionality. Configure CFM on all SPBM VLANs.

CFM is based on the IEEE 802.1ag standard.

IEEE 802.1ag Connectivity Fault Management (CFM) provides OAM tools for the service layer, which allows you to monitor and troubleshoot an end-to-end Ethernet service instance. CFM is the standard for Layer 2 ping, Layer 2 traceroute, and the end-to-end connectivity check of the Ethernet network.

The 802.1ag feature divides or separates a network into administrative domains called Maintenance Domains (MD). Each MD is further subdivided into logical groupings called Maintenance Associations (MA). A single MD can contain several MAs.

Each MA is defined by a set of Maintenance Points (MP). An MP is a demarcation point on an interface that participates in CFM within an MD. Two types of MP exist:

- Maintenance End Point (MEP)
- Maintenance Intermediate Point (MIP)

CFM supports three kinds of standard CFM messages: Continuity Check Message (CCM), Loopback Message (LBM), and Linktrace Message (LTM). Messages are sent between Maintenance Points (MP) in the system.

On the switch, CFM is implemented using the LBM and LTM features only to debug SPBM. CCM messages are not required or supported.

You can assign maintenance levels for each CFM SPBM MEP and MIP to each SPBM B-VLAN individually or you can assign maintenance levels and global MEPs for all SPBM VLANs.

Autogenerated CFM and Explicitly Configured CFM

The switch simplifies CFM configuration with autogenerated CFM. With autogenerated CFM, you use the commands **`cfm spbm enable`** and **`cfm cmac enable`** and the switch creates default MD, MA, MEPs, and MIPs for SPBM B-VLANs and C-VLANs respectively.

If you choose to configure CFM explicitly, you must configure an MD, MA, MEPs, and MIPs.

- For SPBM B-VLANs, the switch provides two methods to configure CFM, namely, autogenerated and explicitly configured. You cannot use both.
- For C-VLANs, you can only use autogenerated CFM.

Autogenerated CFM

You can use autogenerated CFM at a global level to create a MEP and a MIP at a specified level for every SPBM B-VLAN and C-VLAN on the chassis. If you use autogenerated CFM commands, you do not have to configure explicit MDs, MAs, MEPs, or MIPs, and associate them with multiple VLANs.

If you do not want to use autogenerated CFM commands, you can choose to configure explicit MDs, MAs, MEPs, and MIPs for SPBM B-VLANs. However, you cannot use both an autogenerated CFM configuration and an explicit CFM configuration together.

Note:

Previous explicit CFM configurations of MDs, MAs, and MEPs on SPBM B-VLANs continue to be supported. However, if you want to enable the autogenerated commands you must first

remove the existing MEP and MIP on the SPBM B-VLANs. The switch only supports one type of MEP or MIP for each SPBM B-VLAN.

For information on autogenerated CFM configuration using the CLI see:

- Configuring autogenerated CFM on SPBM B-VLANs on page 83
- Configuring autogenerated CFM on C-VLANs on page 85

For information on autogenerated CFM configuration using the EDM see:

- Configuring autogenerated CFM on SPBM B-VLANs on page 111
- Configuring autogenerated CFM on C-VLANs on page 112

Explicitly configured CFM

If you choose to explicitly configure CFM, you must configure an MD, MA, MEPs, and MIPs. You can configure explicit CFM only on SPBM B-VLANs.

For explicit configuration information for CLI see Configuring explicit mode CFM on page 87.

For explicit configuration information for EDM see Configuring explicit CFM in EDM on page 114.

Using CFM

For SPBM B-VLANs, the autogenerated MEPs and MIPs respond to **`l2 ping`**, **`l2 traceroute`**, and **`l2 tracetree`** in the same manner as the MEPs and MIPs created explicitly. For C-VLANs, the autogenerated MEPs and MIPs respond to **`l2 ping`** and **`l2 traceroute`**, but not to **`l2 tracetree`** because no multicast trees exist on C-VLANs. The CFM show commands that display MD, MA, and MEP information work for both autogenerated and explicitly configured CFM MEPs.

You can use CFM to troubleshoot networks and hosts that support the CFM protocol. Once you configure CFM, CFM works in the network whether or not SPBM is in use.

You cannot use CFM to troubleshoot networks and hosts that do not support the CFM protocol, such as a customer domain that does not support CFM. Only devices that support the CFM protocol respond to **`l2 ping`** and **`l2 traceroute`** requests.

Maintenance Domain (MD)

A Maintenance Domain (MD) is the part of a network that is controlled by a single administrator. For example, a customer can engage the services of a service provider, who, in turn, can engage the services of several operators. In this scenario, there can be one MD associated with the customer, one MD associated with the service provider, and one MD associated with each of the operators.

You assign one of the following eight levels to the MD:

- 0–2 (operator levels)
- 3–4 (provider levels)
- 5–7 (customer levels)

The levels separate MDs from each other and provide different areas of functionality to different devices using the network. An MD is characterized by a level and an MD name (optional).

A single MD can contain several Maintenance Associations (MA).

Maintenance Association (MA)

An MA represents a logical grouping of monitored entities within its Domain. It can therefore represent a set of Maintenance association End Points (MEPs), each configured with the same Maintenance Association ID (MAID) and MD Level, established to verify the integrity of a single service instance.

The following figure shows MD level assignment in accordance with the 802.1ag standard. As shown in the figure, MIPs can be associated with MEPs. However, MIPs can also function independently of MEPs.

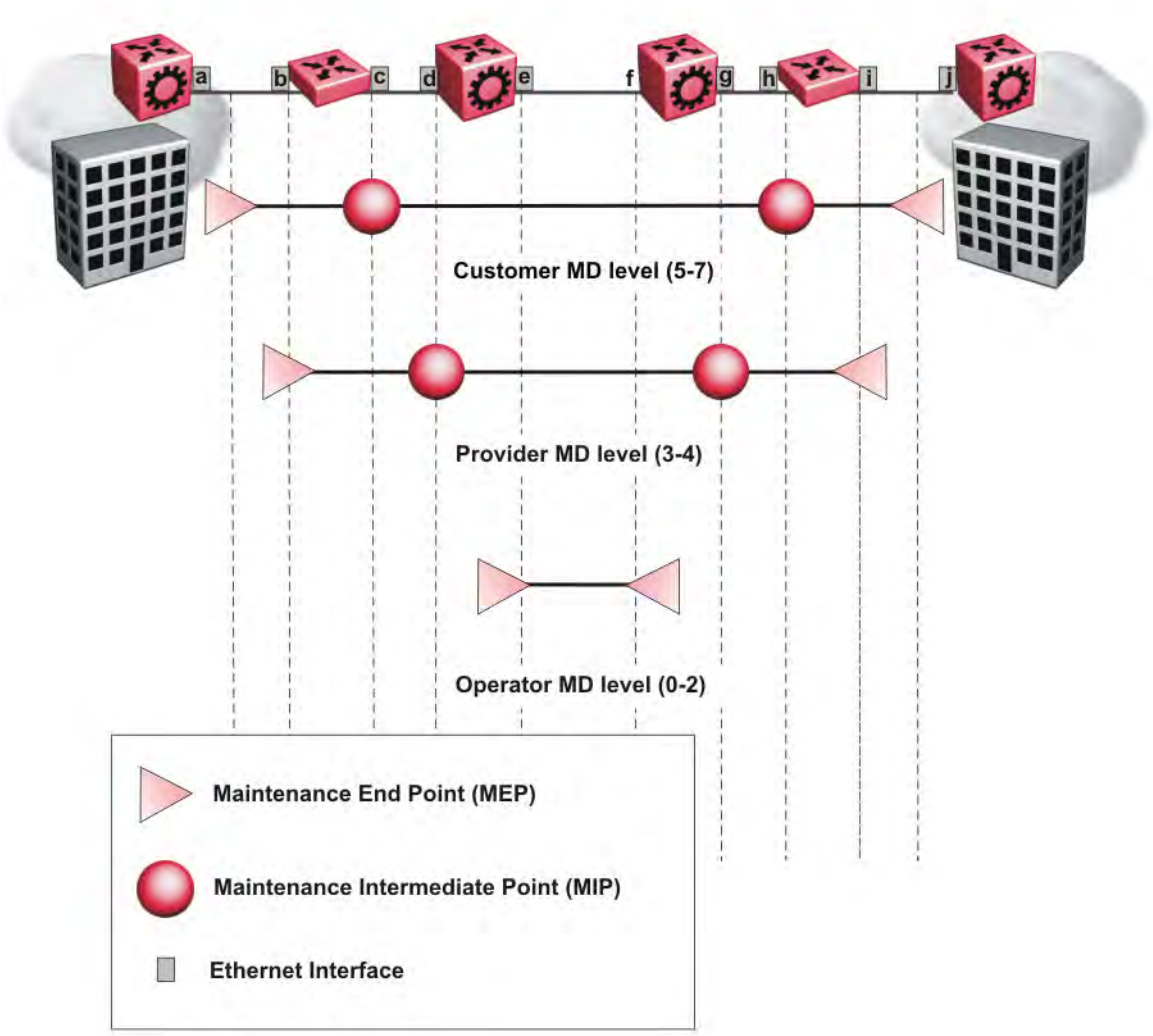

Figure 1: MD level assignment

Maintenance Association Endpoint (MEP)

A Maintenance Endpoint (MEP) represents a managed CFM entity, associated with a specific Domain Service Access Point (DoSAP) of a service instance, which can generate and receive CFM Protocol Data Units (PDU) and track any responses. A MEP is created by MEP ID under the context of an MA. MEP functionality can be divided into the following functions:

- Fault Detection
- Fault Verification
- Fault Isolation
- Fault Notification

Fault detection and notification are achieved through the use of Continuity Check Messages (CCM). CCM messages are not supported.

Fault Verification

Fault verification is achieved through the use of Loopback Messages (LBM). An LBM is a unicast message triggered by the operator issuing an operational command. LBM can be addressed to either a MEP or Maintenance Intermediate Point (MIP) but only a MEP can initiate an LBM. The destination MP can be addressed by its MAC address. The receiving MP responds with a Loopback Response (LBR). LBM can contain an arbitrary amount of data that can be used to diagnose faults as well as performance measurements. The receiving MP copies the data to the LBR.

LBM Message

The LBM packet is often compared to a ping. A MEP transmits the LBM packet. This packet can be addressed to another MEP or to the MAC address of the MP; in the case of SPBM, this is the SPBM system ID or its virtual SMLT MAC. Only the MP for which the packet is addressed responds with an LBR message.

- Provides “ICMP ping like” functionality natively at Layer-2.
- DA is the MAC address of the target.
- Includes a transaction identifier that allows the corresponding LBR to be identified when more than one LBM request is waiting for a response.
- Bridges forward the frame using the normal FDB rules.
- Only the target (MIP or MEP) responds.
- Initiator can choose the size and contents data portion of the LBM frame.
- Can be used to check the ability of the network to forward different sized frames.

L2 Ping

The **`l2 ping`** command is a proprietary command that allows a user to trigger an LBM message.

For B-VLANs, specify either the destination MAC address or node name.

This provides a simpler command syntax than the standard LBM commands, which require the user to specify the MD, MA, and MEP ID information. The **`l2 ping`** command provides a ping equivalent at Layer 2 for use with nodes on the SPBM B-VLAN in the customer domain. SPBM B-VLANs support the SMLT virtual option for the source mode.

Important:

CFM CMAC L2 ping and L2 traceroute traversing over Fabric Extend tunnels are not supported when originating from a VSP 4000.

Fault Isolation

Fault isolation is achieved through the use of Linktrace Messages (LTM). LTM is intercepted by all the MPs on the way to the destination MP. The switch supports two types of LTM.

The first type, the unicast LTM, can be addressed to either MEP or MIP MAC address. Each MP on the way decrements the TTL field in the LTM frame, sends Linktrace Reply (LTR), and forwards the original LTM to the destination. The LTM is forwarded until it reaches its destination or the TTL value is decremented to zero. LTR is a unicast message addressed to the originating MEP.

The second type, the proprietary LTM, is used to map the MAC addresses of the SPBM network; in this case the target MAC is not an MP, but rather a service instance identifier (I-SID).

Link Trace Message

Connectivity Fault Management offers link trace messaging for fast fault detection. Link trace messages allow operators, service providers and customers to verify the connectivity that they provide or use and to debug systems.

Link trace message — unicast

The link trace message (LTM) is often compared to traceroute. A MEP transmits the LTM packet. This packet specifies the target MAC address of an MP which is the SPBM system id or the virtual SMLT MAC. MPs on the path to the target address respond with an LTR.

- Trace the path to any given MAC address.
- DA is unicast
- LTM contains:
  - Time to live (TTL)
  - Transaction Identifier

  - Originator MAC address
  - Target MAC address
- CFM unaware entities forward the frame as is like any other data frame.
- MIP or MEP that is not on the path to the target discards the LTM and does not reply.
- MIP that is on the path to the target
  - Forwards the LTM after decrementing the TTL and replacing the SA with its own address.
  - Sends a reply (LTR) to the originator.
  - Identifies itself in the forwarded LTM and LTR by modifying TLV information.
- If the MIP or MEP is a target
  - Sends an LTR to the originator.
  - Identifies itself in the forwarded LTM and LTR by modifying TLV information.
- A MEP that is not the target but is on the path to the target
  - Generates a reply as described above.
  - It also sets one of the flags fields in the reply to indicate that it is the terminal MEP.

Link trace message — multicast

The multicast link trace message (LTM) can be used to trace the multicast tree from any node on any I- SID using the nickname MAC address and the I-SID multicast address.

Specifying a multicast target address for an LTM allows for the tracing of the multicast tree corresponding to that destination address (DA). With a multicast target every node that is in the active topology for that multicast address responds with a Linktrace reply and also forwards the LTM frame along the multicast path. Missing Linktrace replies (LTRs) from the nodes in the path indicate the point of first failure.

This functionality allows you to better troubleshoot I-SID multicast paths in a SPBM network.

L2 Traceroute

The **`l2 traceroute`** command is a proprietary command that allows you to trigger an LTM message. Use this command as follows:

- For B-VLANs, specify either the destination MAC address or node name.
- For C-VLANs, specify the destination MAC address.

This command provides a simpler command syntax than the standard LTM commands, which require the user to specify the MD, MA, and MEP ID information. The **`l2 traceroute`** command provides a trace equivalent at Layer 2 for use with nodes on the SPBM B-VLAN in the customer domain.

Important:

Only the VSP 4000 Series switch supports CFM configuration on C-VLANs.

CFM CMAC L2 ping and L2 traceroute traversing over Fabric Extend tunnels are not supported when originating from a VSP 4000.

You can use CFM to troubleshoot networks and hosts that support the CFM protocol. After you configure CFM, CFM works in the network whether or not SPBM is in use.

You cannot use CFM to troubleshoot networks and hosts that do not support the CFM protocol, such as a customer domain that does not support CFM. Only devices that support the CFM protocol respond to **`l2 ping`** and **`l2 traceroute`** requests.

l2 traceroute with IP address

The **`l2 traceroute`** command allows you to specify an IP address as the destination address. In this case, the IP address can be either a C-VLAN or a B-VLAN in the SPBM cloud.

The **`l2 traceroute`** command converts Layer 3 IP information to an appropriate Layer 2 VLAN and MAC combination. The system can also target IP addresses that are not SPBM derived routes.

If ECMP is enabled, **`l2 traceroute`** runs internally for each of the VLAN paths returned, and displays a summary of the results. If ECMP is disabled, the results display only one path.

Note:

If you use the **`l2 traceroute ip-address`** command on a DvR Leaf node, the output only shows DvR Controller IP addresses if the IP address or host route specified is unknown in the DvR domain.

Destination addresses for C-VLAN l2 traceroute and linktrace messages

For C-VLANs, CFM uses the following destination MAC addresses for the corresponding maintenance domain (MD) levels for **`l2 traceroute`** and **`linktrace`** messages.

The switch supports both **`l2 traceroute`** and **`linktrace`** for C-VLANs, but It is recommended that you use **`l2 traceroute`**.

Table 9: MD levels and corresponding destination addresses for CFM for C-VLANs

CFM MD Level	Destination MAC address
0	01:80:c2:00:00:38
1	01:80:c2:00:00:39
2	01:80:c2:00:00:3a
3	01:80:c2:00:00:3b
4	01:80:c2:00:00:3c
5	01:80:c2:00:00:3d
6	01:80:c2:00:00:3e
7	01:80:c2:00:00:3f

L2 Tracetree

The **`l2 tracetree`** command is a proprietary command that allows a user to trigger a multicast LTM message by specifying the B-VLAN and I-SID. This command allows the user to view a multicast tree on the SPBM B-VLAN from the source node to the destination nodes for a particular I-SID.

Note:

The Virtual Services Platform 4000 Series does not support the **`l2 tracetree`** command on C-VLANs because no multicast tree exists on C-VLANs.

L2 Tracetree-fan

The **`l2 tracetree-fan`** command allows a user to trigger an LTM on the internal Fabric Area Network I-SID. This command allows the user to trace the FAN tree.

Maintenance Domain Intermediate Point (MIP)

MIPs do not initialize any CFM messages. MIPs passively receive CFM messages, process the messages received and respond back to the originating MEP. By responding to received CFM messages, MIPs can support discovery of hop-by-hop path among MEPs, allow connection failures to be isolated to smaller segments of the network to help discover location of faults along the paths. MIPs can be created independent of MEPs. MIP functionality can be summarized as:

- Respond to Loopback (ping) messages at the same level as itself and addressed to it.
- Respond to Linktrace (traceroute) messages.
- Forward Linktrace messages after decrementing the TTL.

Layer 2 Tracemroute

The **`l2tracemroute`** command is a proprietary command that allows the user to trace the multicast tree for a certain multicast flow. The user specifies source, group, and service context (either VLAN or VRF) for the multicast flow to trace.

CFM sends a multicast LTM using an internal calculation to map the source, group, and context to the corresponding target address. The LTR comes from all leaves of the multicast tree for that flow, as well as transit nodes. The target MAC used in the LTM is a combination of the data I-SID and the nickname and the packet is sent on the appropriate SPBM B-VLAN. The user can see the generated multicast tree for that flow, which includes the data I-SID and nickname.

Nodal MPs

Nodal MPs provide both MEP and MIP functionality for SPBM deployments. Nodal MPs are associated with a B-VLAN and are VLAN encapsulated packets. The Nodal MEP provides traceability and troubleshooting at the system level for a given B-VLAN. Each node (chassis) has a given MAC address and communicates with other nodes. The SPBM instance MAC address is used as the MAC address of the Nodal MP. The Nodal B-VLAN MPs supports eight levels of CFM and you configure the Nodal B-VLAN MPs on a per B-VLAN basis. Virtual SMLT 10 MAC addresses are also able to respond for LTM and LBM.

Nodal B-VLAN MEPs

The Nodal B-VLAN MEPs created on the CP and function as if they are connected to the virtual interface of the given B-VLAN. Because of this they are supported for both port and MLT based B-VLANs. To support this behavior a MAC Entry is added to the FDB and a new CFM data-path table containing the B-VLAN and MP level are added to direct CFM frames to the CP as required.

Nodal B-VLAN MIPs

The Nodal MIP is associated with a B-VLAN. VLAN and level are sufficient to specify the Nodal MIP entity. The Nodal MIP MAC address is the SPBM system ID for the node on which it resides. If the fastpath sends a message to the CP, the MIP responds if it is not the target and the MEP responds if it is the target.

Nodal B-VLAN MIPs with SMLT

When Nodal MEPs or MIPs are on SPBM B-VLANs the LTM code uses a unicast MAC DA. The LTM DA is the same as the target MAC address, which is the SPBM MAC address or the SMLT MAC address of the target node.

The switch supports SMLT interaction with SPBM. This is accomplished by using two B-VIDs into the core from each pair of SMLT terminating nodes. Both nodes advertise the Nodal B-MAC into the core on both B-VIDS. In addition each node advertises the SMLT virtual B-MAC on one of the two B-VLANs.

The Nodal MEP and MIP are expanded to respond to both the Nodal MAC address as well as the Virtual SMLT MAC address if both MACs are being advertised on its B-VLAN. In addition a source mode is added to the LTM and LBM command to use either the Nodal MAC or the SMLT virtual MAC address as the source MAC in the packet.

Configuration Considerations

When you configure CFM, be aware of the following configuration considerations.

General CFM

- A single switch has a limit of one MEP and one MIP on a C-VLAN or B-VLAN.

- The maintenance level for MEPs and MIPs on a given B-VID (in a network) must be configured to the same level for them to respond to a given CFM command.
- You can configure global CFM at only one MD level for each switch for each VLAN type.
- All nodal MEPs and MIPs are restricted to SPBM B-VIDs.
- SMLT Virtual MAC for C-VLAN does not exist, so the switch does not support this option for the **`l2 ping`** and **`l2 traceroute`** commands.

Autogenerated CFM

- Autogenerated MEPs are not unique across the entire network unless you configure the global MEP ID on each switch to a different value. You must configure a unique MEP ID at a global level, for CFM.
- A single switch can have only one autogenerated MEP or MIP for each B-VLAN or C-VLAN.

Explicit CFM

- Previous explicit CFM configurations of MDs, MAs and MEPs on SPBM B-VLANs continue to be supported. However, if you want to enable autogenerated CFM you must first remove the existing MEP and MIP on the SPBM B-VLAN.
- You can assign maintenance levels for each CFM SPBM MEP and MIP to each SPBM B-VLAN individually or you can assign maintenance levels and global MEPs for all SPBM VLANs by following the appropriate procedure:
  - Assigning a MEP MIP level to an SPBM B-VLAN on page 91
  - Assigning MEP MIP levels to SPBM B-VLANs globally on page 93
  - Configuring CFM nodal MEP on page 117

C-VLAN versus SPBM B-VLAN considerations

Important:

Only VSP 4000 Series supports CFM configuration on C-VLANs.

CFM breaks the network into sections, called MEPs, so you can determine exactly where the problem exists.

The MEPs and MIPs configured for SPBM VLANs do not respond to CFM messages sent from C-MAC VLANs because the VLAN and packet encapsulation are different.

To forward customer traffic across the core network backbone, SPBM uses IEEE 802.1ah Provider Backbone Bridging (PBB) MAC-in-MAC encapsulation, which hides the customer MAC (C-MAC) addresses in a backbone MAC (B-MAC) address pair. MAC-in-MAC encapsulation defines a B-MAC destination address (BMAC-DA) and a B-MAC source address (BMAC-SA). In SPBM, each node populates its forwarding database (FDB) with the B-MAC information derived from the IS-IS shortest path tree calculations.

Typically the SPBM Backbone Core Bridges (BCBs) in the SPBM cloud only learn the B-MAC addresses. The Backbone Edge Bridges (BEBs) know the Customer MACs on the appropriate BEBs that terminate the virtual services networks (VSNs). As such, the nodes within the SPBM cloud have no knowledge of the C-MAC addresses in the VSNs.

Important:

To trace a route to a MAC address, the MAC address must be in the VLAN FDB table.

- For C-VLANs, you have to trigger an **`l2 ping`** to learn the C-MAC address.
- For B-VLANs, you do not have to trigger an **`l2 ping`** to learn the C-MAC address because IS-IS populates the MAC addresses in the FDB table.

In both cases, linktrace traces the path up to the closest device to that MAC address that supports CFM in the SPBM cloud.

C-VLAN source addresses

CFM uses either the VLAN MAC or the CFM C-MAC for the BMAC-SA for the C-VLANs. The CFM C-MAC is the value of the management base MAC, which ends in 0x64. The system creates the VLAN MAC after a user adds an IP address to a VLAN.

If a VLAN has a MAC address, the system uses the VLAN MAC as the BMAC-SA by default. If a VLAN does not have a MAC address, the system uses the CFM C-MAC for the BMAC-SA. You may also configure the system to use the CFM C-MAC, even if a VLAN MAC exists.

CFM Configuration Using CLI

This section provides procedures to configure and use Connectivity Fault Management (CFM) using Command Line Interface (CLI). The Shortest Path Bridging MAC (SPBM) network needs a mechanism to debug connectivity issues and to isolate faults. This is performed at Layer 2, not Layer 3. To support troubleshooting of the SPBM cloud, the switch supports a subset of CFM functionality.

Note:

When you enable CFM in an SBPM network, it is recommended that you enable CFM on the Backbone Edge Bridges (BEB) and on all Backbone Core Bridges (BCB). If you do not enable CFM on a particular node, you cannot obtain CFM debug information from that node.

You can configure CFM using one of two modes: simplified or explicit. Both modes are described in the following sections, but the simplified mode is recommended.

Note:

If you enable the **`cfm spbm enable`** command, you cannot assign a MEP/MIP level to an individual SPBM B-VLAN or configure CFM MD maintenance levels individually.

Regardless of whether you have chosen to configure individually or globally, there is one MEP per SPBM B-VLAN and one MIP level per SPBM B-VLAN.

Autogenerated CFM

CFM provides two methods for configuration; autogenerated and explicit. You cannot use both. You must choose one or the other. Use the procedures in this section to configure autogenerated MEPs that eliminate the need to configure a MD, MA, and MEP ID to create a MEP.

- For SPBM B-VLANs, you can use either autogenerated or explicitly configured CFM MEPs.
- For C-VLANs, you can only use autogenerated CFM MEPs.

The CFM show commands that display MD, MA, and MEP information work for both autogenerated and explicitly configured CFM MEPs.

Previous explicit CFM configurations of MDs, MAs and MEPs on SPBM B-VLANs continue to function. However, if you want to enable the autogenerated commands you must first remove the existing MEP and MIP on the SPBM B-VLAN.

The switch only supports one MEP and one MIP, either autogenerated or explicitly configured, on the SPBM B-VLAN. Similarly, the switch only supports one MEP and one MIP on the C-VLAN. This means that if you want to use these autogenerated MEPs, you cannot use your existing CFM configuration. You must first remove the existing MEP or MIP on the SPBM B-VLAN.

For information on configuring autogenerated CFM using the CLI, see:

- Configuring autogenerated CFM on SPBM B-VLANs on page 83
- Configuring autogenerated CFM on C-VLANs on page 85

Configuring Autogenerated CFM on SPBM B-VLANs

Use this procedure to configure the autogenerated CFM MEP and MIP level for every SPBM B-VLAN on the chassis. This eliminates the need to explicitly configure an MD, MA, and MEP ID, and to associate the MEP and MIP level to the SPBM B-VLAN.

When you enable this feature, the device creates a global MD (named spbm) for all the SPBM Nodal MEPs. This MD has a default maintenance level of 4, which you can change with the level attribute. All the MEPs that the device creates use the MEP ID configured under the global context, which has a default value of 1.

The nodal MEPs are automatically associated with the SPBM B-VLANs configured. The MIP level maps to the global level. When you enable the feature, the device automatically associates the MIP level with the SPBM B-VLANs configured. The feature is disabled by default.

Important:

CFM supports one MEP or MIP for each SPBM B-VLAN only. This means that if you want to use these autogenerated MEPs, you cannot use your existing CFM configuration. You must first remove the existing MEP or MIP on the SPBM B-VLAN. If you want to continue configuring MEPs manually, skip this procedure.

Procedure

1. Enter Global Configuration mode:

```
enable

configure terminal
```

2. Configure the maintenance level for every CFM SPBM MEP and MP level on all the SPBM B-VLANs:

   ```
   cfm spbm level <0-7>
   ```

   You can change this level from the default of 4 either before or after the feature is enabled.

   Only configure global CFM at one MD level for each chassis for each VLAN type.

3. Assign a global CFM MEP ID for all CFM SPBM MEPs:

   ```
   cfm spbm mepid <1-8191>
   ```

4. Enable the autogenerated CFM for SPBM B-VLANs globally:

   ```
   cfm spbm enable
   ```

5. **(Optional)** Configure the maintenance level for every CFM SPBM MEP and MP level on all the SPBM B-VLANs to the default:

   ```
   default cfm spbm level
   ```

6. **(Optional)** Assign a global CFM MEP ID for all CFM SPBM MEPs to the default:

   ```
   default cfm spbm mepid
   ```

7. **(Optional)** Disable the global CFM MEPs and MIPs:

   ```
   no cfm spbm enable
   ```

8. Display the global CFM MEP configuration:

   ```
   show cfm spbm
   ```

Example

Configure autogenerated CFM MEPs and MIPs:

```
Switch>enable
Switch#configure terminal
Switch(config)#cfm spbm level 6
Switch(config)#cfm spbm mepid 4
Switch(config)#cfm spbm enable
Switch(config)#show cfm spbm

LEVEL ADMIN    MEPID    MAC
=======================================
6        enable        4        00:15:e8:b8:a3:df
```

Variable Definitions

Use the data in the following table to use the **`cfm spbm`** command.

Variable	Value
level<*0-7*>	Specifies the global SPBM CFM maintenance level for the chassis within the range of 0 to 7. The default is 4. Only configure global CFM at one MD level for each chassis for each VLAN type.
mepid<*1-8191*>	Specifies the global MEP ID within the range of 1–8191. Select a unique ID for each switch to ensure that the MEPs are unique across the network. The default is 1.
enable	Enables autogenerated CFM on all SPBM B-VLANs.

Job Aid

The following table describes the fields for the **`show cfm spbm`** command.

Parameter	Description
LEVEL	Specifies the global SPBM CFM maintenance level for the chassis. The default is 4.
ADMIN	Specifies if CFM MEPs and MIPs are globally enabled.
MEP ID	Specifies the global MEP ID. The default is 1.
MAC	Specifies the MAC address.

Configure Autogenerated CFM on C-VLANs

Use this procedure to configure the autogenerated CFM MEP and MIP level for every C-VLAN on the chassis.

Note:

For C-VLANs, you can only use autogenerated CFM MEPs.

Important:

CFM supports one MEP or MIP for each C-VLAN. Only autogenerated CFM provides support for configuring MEP and MIPs on C-VLANs. You cannot explicitly configure C-VLANs.

About this task

When you enable this feature, you create a global MD (named cmac) for all the customer MAC (C-MAC) MEPs. This global MD has a default maintenance level of 4, which you can change with the level attribute. The autogenerated CFM commands also create an MA for each C-VLAN, a MEP for each C-VLAN, associate the MEP with the corresponding C-VLAN, and a MIP with the C-VLAN.

All the MEPs that the device creates use the MEP ID configured under the global context, which has a default value of 1. The device automatically associates the MEPs with the C-VLANs configured. The MIP level maps to the global level. The device automatically associates the MIP level with the C-VLANs configured when you enable the feature.

The feature is disabled by default.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure the maintenance level for every CFM C-MAC MEP and MP level on all the C-VLANs:

   ```
   cfm cmac level <0-7>
   ```

   Only configure global CFM at one MD level for each chassis for each VLAN type.

3. Assign a global CFM MEP ID for all CFM C-MAC MEPs:

   ```
   cfm cmac mepid <1-8191>
   ```

4. Enable the autogenerated CFM for C-VLANs:

   ```
   cfm cmac enable
   ```

5. **(Optional)** Configure the maintenance level for every CFM C-MAC MEPs and MP level on all the C-VLANs to the default:

   ```
   default cfm cmac level
   ```

6. **(Optional)** Assign a global CFM MEP ID for all CFM C-MAC MEPs to the default:

   ```
   default cfm cmac mepid
   ```

7. **(Optional)** Disable the global CFM MEPs and MIPs:

   ```
   no cfm cmac enable
   ```

8. Display the global CFM MEP configuration:

   ```
   show cfm cmac
   ```

Example

Configure autogenerated CFM MEPs and MIP level:

```
Switch>enable
Switch#configure terminal
Switch(config)#cfm cmac level 0
Switch(config)#cfm cmac mepid 4
Switch(config)#cfm cmac enable
Switch(config)#show cfm cmac

LEVEL ADMIN    MEPID    MAC
=======================================
0        enable        4        00:15:e8:b8:a3:de
```

Variable Definitions

Use the data in the following table for the **`cfm cmac`** command.

Variable	Value
level<*0-7*>	Specifies the global C-MAC CFM maintenance level for the chassis within the range of 0 to 7. The default is 4. Only configure global CFM at one MD level for each chassis for each VLAN type.
mepid<*1-8191*>	Specifies the global MEP ID within the range of 1–8191. Select a unique ID for each switch to ensure that the MEPs are unique across the network. The default is 1. **Note:** The MA takes its name from this value for autogenerated CFM. For example, if you specify 500 as the MEP ID, the MA will also be 500.
enable	Enables autogenerated CFM for all C-MAC VLANs.

Job Aid

The following table describes the fields for the **`show cfm cmac`** command.

Parameter	Description
LEVEL	Specifies the global C-VLAN CFM maintenance level for the chassis. The default is 4.
ADMIN	Specifies if CFM C-VLAN MEPs and MIPs are globally enabled.
MEP ID	Specifies the global MEP ID. The default is 1.
MAC	Specifies the MAC address.

Configuring Explicit Mode CFM

In the explicit mode of configuring CFM, you can manually configure an MD, MA, MEP and then associate the MEP to a B-VLAN and assign a MIP level to a B-VLAN.

Note:

If you use autogenerated CFM, these steps are unnecessary.

Configuring CFM MD

Use this procedure to configure the Connectivity Fault Management (CFM) Maintenance Domain (MD). An MD is the part of a network that is controlled by a single administrator. A single MD can contain several Maintenance Associations (MA).

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Create the CFM MD:

```
cfm maintenance-domain WORD<0-22> [index <1-2147483647>]
[maintenance-level <0-7>] [level <0-7>]
```

3. Display the CFM MD configuration:

```
show cfm maintenance-domain
```

4. Delete the CFM MD:

```
no cfm maintenance-domain WORD<0-22>
```

Example

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# cfm maintenance-domain md1 index 99 maintenance-level 3
Switch:1(config)# show cfm maintenance-domain
================================================================================
                          Maintenance Domain
================================================================================
Domain Name             Domain Index    Level Domain Type
--------------------------------------------------------------------------------
md1                     99              3     NODAL

 Total number of Maintenance Domain entries: 1.
Switch:1(config)# no cfm maintenance-domain md1
Switch:1(config)# show cfm maintenance-domain

================================================================================
                          Maintenance Domain
================================================================================
Domain Name             Domain Index    Level Domain Type
--------------------------------------------------------------------------------

 Total number of Maintenance Domain entries: 0.
```

Variable Definitions

Use the data in the following table to use the **`cfm maintenance-domain`** command.

Variable	Value
WORD<0–22>	Specifies the maintenance domain name.
index *<1–2147483647>*	Specifies a maintenance domain entry index.
maintenance-level *<0–7>*	Specifies the MD maintenance level when creating the MD. The default is 4.
level *<0–7>*	Modifies the MD maintenance level for an existing MD. The default is 4.

Configuring CFM MA

Use this procedure to configure the CFM Maintenance Association (MA). An MA represents a logical grouping of monitored entities within its domain. It can therefore represent a set of Maintenance Association End Points (MEPs), each configured with the same Maintenance Association ID (MAID) and MD Level, established to verify the integrity of a single service instance.

Procedure

1. Enter Global Configuration mode:

   ```
   enable

   configure terminal
   ```

2. Create the CFM MA:

   ```
   cfm maintenance-association WORD<0-22> WORD<0-22> [index <1-
   2147483647>]
   ```

3. Display the CFM MA configuration:

   ```
   show cfm maintenance-association
   ```

4. Use the following command, if you want to delete the CFM MA:

   ```
   no cfm maintenance-association WORD<0-22> WORD<0-22>
   ```

Example

```
Switch:1> enable

Switch:1# configure terminal

Switch:1(config)# cfm maintenance-association md1 ma1 index 98

Switch:1(config)# show cfm maintenance-association
================================================================================
                          Maintenance Association Status

================================================================================
Domain Name             Assn Name               Domain Idx  Assn Idx
--------------------------------------------------------------------------------
md1                     ma1                     1           98

 Total number of Maintenance Association entries: 1.

================================================================================
                          Maintenance Association config
================================================================================
Domain Name             Assn Name
--------------------------------------------------------------------------------
md1                     ma1

 Total number of MA entries: 1.
```

Variable Definitions

Use the data in the following table to use the **`cfm maintenance-association`** command.

Variable	Value
WORD<0–22> WORD<0–22>	Creates the CFM MA. The first parameter, specifies the MD name. The second parameter, specifies the MA short name.
index *<1–2147483647>*	Specifies a maintenance association entry index.

Configuring CFM MEP

Use this procedure to configure the CFM Maintenance Endpoint (MEP). A MEP represents a managed CFM entity, associated with a specific Domain Service Access Point (DoSAP) of a service instance, which can generate and receive CFM Protocol Data Units (PDU) and track any responses. A MEP is created by MEP ID under the context of an MA.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Create the CFM MEP:

   ```
   cfm maintenance-endpoint WORD<0-22> WORD<0-22> <1-8191> enable
   [state <enable>]
   ```

3. Enable an existing CFM MEP:

   ```
   cfm maintenance-endpoint WORD<0-22> WORD<0-22> <1-8191> enable
   ```

4. Disable an existing CFM MEP:

   ```
   no cfm maintenance-endpoint WORD<0-22> WORD<0-22> <1-8191> enable
   ```

5. Display the CFM MEP configuration:

   ```
   show cfm maintenance-endpoint
   ```

6. Delete an existing CFM MEP:

   ```
   no cfm maintenance-endpoint WORD<0-22> WORD<0-22> <1-8191>
   ```

Example

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# cfm maintenance-endpoint md1 ma1 1 state enable
Switch:1(config)# show cfm maintenance-endpoint
================================================================================
                          Maintenance Endpoint Config
================================================================================
DOMAIN                  ASSOCIATION             MEP  ADMIN
NAME                    NAME                    ID
--------------------------------------------------------------------------------
md1                     ma1                     1    enable
```

```
 Total number of MEP entries: 1.

================================================================================
                          Maintenance Endpoint Service
================================================================================
DOMAIN_NAME             ASSN_NAME               MEP_ID TYPE    SERVICE_DESCRIPTION
--------------------------------------------------------------------------------
md1                     ma1                     1      nodal   Vlan 1, Level 4

 Total number of MEP entries: 1.
```

Variable Definitions

Use the data in the following table to use the **`cfm maintenance-endpoint`** command.

Variable	Value
WORD<0–22>	The first parameter, specifies the MD name.
WORD<0–22>	The second parameter, specifies the MA short name.
<1–8191>	Specifies the MEP ID.
enable	Enables an existing MEP. Use this parameter with the no option to disable an existing MEP.
state {enable | disable}	Enables or disables the MEP when creating the MEP. The default is disabled.

Assigning a MEP/MIP Level to an SPBM B-VLAN

Use this procedure to assign a nodal MEP to an SPBM B-VLAN. The Nodal MEP provides traceability and troubleshooting at the system level for a given B-VLAN. The Nodal B-VLAN MEPs created on the CP and function as if they are connected to the virtual interface of the given B-VLAN. Because of this they are supported for both port and MLT based B-VLANs.

Nodal MPs provide both MEP and MIP functionality for SPBM deployments. Nodal MPs are associated with a B-VLAN and are VLAN encapsulated packets. Each node (chassis) has a given MAC address and communicates with other nodes. The SPBM instance MAC address is used as the MAC address of the Nodal MP.

Before you begin

- You must configure a CFM MD, MA, and MEP.

Procedure

1. Add nodal MEPs to the B-VLAN:

   ```
   vlan nodal-mep <1-4059> WORD<0-22> WORD<0-22> <1-8191>
   ```

2. Display the nodal MEP configuration:

   ```
   show vlan nodal-mep <1-4059>
   ```

3. Remove the nodal MEPs from the B-VLAN:

   ```
   no vlan nodal-mep <1-4059> WORD<0-22> WORD<0-22> <1-8191>
   ```

4. Add nodal MIP level to the B-VLAN:

```
vlan nodal-mip-level <1-4059> WORD<0-15>
```

5. Display the nodal MIP level configuration:

```
show vlan nodal-mip-level [<1-4059>]
```

6. Remove the nodal MIP level from the B-VLAN:

```
no vlan nodal-mip-level <1-4059> WORD<0-15>
```

Example

```
Switch:1> enable

Switch:1# configure terminal

Switch:1(config)# vlan nodal-mep 100 md1 ma1 2

Switch:1(config)# show vlan nodal-mep

================================================================================
                                Vlan Nodal Mep
================================================================================
VLAN_ID    DOMAIN_NAME.ASSOCIATION_NAME.MEP_ID
--------------------------------------------------------------------------------
100       spbm.100.6
200       spbm.200.6

Switch:1(config)# vlan nodal-mip 100 6

Switch:1(config)# show vlan nodal-mip

================================================================================
                            Vlan Nodal Mip Level
================================================================================
VLAN_ID    NODAL_MIP_LEVEL_LIST
--------------------------------------------------------------------------------
1
100            6
216
304
41000
1001
```

Variable Definitions

Use the data in the following table to use the **`vlan nodal-mep`** command.

Variable	Value
<1-4059>	Specifies the VLAN ID.
WORD<0–22>	The first parameter, specifies the Maintenance Domain name.
WORD<0–22>	The second parameter, specifies the Maintenance Association name.
<1–8191>	Specifies the nodal MEPs to add to the VLAN.

Use the data in the following table to use the **`vlan nodal-mip-level`** command.

Variable	Value
<1-4059>	Adds the nodal MIP level. Specifies the VLAN ID.
WORD<0–15>	Adds the nodal MIP level, which has up to eight levels, ranging from 0 to 7.

Assigning MEP/MIP Levels to SPBM B-VLANs Globally

Note:

If you enable the **`cfm spbm enable`** command, you cannot assign a MEP/MIP level to an individual SPBM B-VLAN or configure CFM MD maintenance levels individually.

About this task

Enables the global CFM MEP and MIPs for all SPBM B-VLANs.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enable simplified CFM configuration for SPBM VLANs:

   ```
   cfm spbm enable
   ```

3. Enter the CFM SPBM level:

   ```
   cfm spbm level <0-7>
   ```

4. Enter the CFM SPBM MEPID level:

   ```
   cfm spbm mepid <1-8191>
   ```

Example

```
Switch:1(config)# cfm spbm level 7
Switch:1(config)# cfm spbm mepid 12
Switch:1(config)# cfm spbm enable
```

Variable Definitions

Use the data in the following table to use the simplified CFM commands.

Variable	Value
spbm level *<0–7>*	Configures the maintenance level for every CFM SPBM MEP and MIP level on all SPBM B-VLANs. The default is 4.
mepid *<1–8191>*	Assigns a global MEP ID for all CFM SPBM MEPs. The default is 1.
no cfm spbm enable	Disables global configuration of CFM SPBM MEP and MIP levels on all SPBM B-VLANs.

Table continues…

Variable	Value
default cfm spbm level	Returns maintenance level to default for all CFM SPBM MEP and MIP level on all SPBM B-VLANs.
default cfm spbm mepid	Returns MEP ID for all CFM SPBM MEPs to default.
show cfm spbm	Displays the global CFM MEP configuration for SPBM B-VLANs.

Triggering a Loopback Test (LBM)

Use this procedure to trigger a loopback test.

The LBM packet is often compared to ping. An MEP transmits the loopback message to an intermediate or endpoint within a domain for the purpose of fault verification. This can be used to check the ability of the network to forward different sized frames.

Before you begin

- You must have a MEP that is associated with a B-VLAN.

Procedure

Trigger the loopback test:

```
loopback WORD<0-22> WORD<0-22> <1-8191> <0x00:0x00:0x00:0x00:0x00:0x00>
[burst-count <1-200>] [data-tlv-size <0-400>] [frame-size <64-1500>]
[interframe-interval <msecs>] [priority <0-7>] [source-mode {nodal|
noVlanMac|smltVirtual}][testfill-pattern <all-zero|all-zero-crc|pseudo-
random-bit-sequence|pseudo-random-bit-sequence-crc>] [time-out <1-10>]
```

Example

```
Switch:1# loopback md1 4001 13 00:14:0D:A2:B3:DF burst-count 10 priority 3
time-out 5
```

```
Result of LBM from mep: spbm.bvlan1000.8 to MAC address: 00:66:00:66:00:66 :
  Sequence number of the first LBM is 150404162
  The total number of LBMs sent out is 1
  The number of LBRs received is 1
  The number of LBRs lost is 0
  The percentage of LBMs lost is 0.00%
  The RTT Min is 15071 microsecs, Max is 15071 microsecs, Average is 15071.00 microsecs
  The Standard Deviation of RTT is 0.00 microsecs
```

Variable Definitions

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

Variable	Value
WORD<0–22>	The first parameter, specifies the MD name.
WORD<0–22>	The second parameter, specifies the MA name.
<1–8191>	Specifies the MEP ID.

Table continues…

Variable	Value
<0x00:0x00:0x00:0x00:0x00:0x00>	Specifies the remote MAC address to reach the MEP/MIP.
burst-count *<1–200>*	Specifies the burst-count.
data-tlv-size *<0–400>*	Specifies the data TLV size.
frame-size *<64–1500>*	Specifies the frame-size. The default is 0.
priority *<0–7>*	Specifies the priority. The default is 7.
source-mode*{nodal|noVlanMac|smltVirtual}]*|	Specifies the source mode: • nodal • noVlanMac — Use this value with C-VLAN only. When you select this option, even if a VLAN MAC address exists, the system uses the CFM C-MAC as the BMAC-SA. • smltVirtual—Use this value with B-VLANs only. The default is nodal.
testfill-pattern *{all-zero|all-zero-crc|pseudo-random-bit-sequence|pseudo-random-bit-sequence-crc}*	Specifies the testfill pattern: • all-zero — null signal without cyclic redundancy check • all-zero-crc — null signal with cyclic redundancy check with 32-bit polynomial • pseudo-random-bit-sequence — pseudo-random-bit-sequence without cyclic redundancy check • pseudo-random-bit-sequence-crc — pseudo-random-bit-sequence with cyclic redundancy check with 32-bit polynomial. A cyclic redundancy check is a code that detects errors. The default is 1: all-zero.
time-out *<1–10>*	Specifies the time-out interval in seconds. The default is 3.

Triggering Linktrace (LTM)

Use the following procedure to trigger a linktrace.

The Linktrace Message is often compared to traceroute. An MEP transmits the Linktrace Message packet to a maintenance endpoint with intermediate points responding to indicate the path of the traffic within a domain for the purpose of fault isolation. The packet specifies the target MAC address of an MP, which is the SPBM system ID or the virtual SMLT MAC. MPs on the path to the target address respond with an LTR.

Before you begin

- You must have a MEP that is associated with a VLAN.

Procedure

Trigger the linktrace:

```
linktrace WORD<0-22> WORD<0-22> <1-8191> <0x00:0x00:0x00:0x00:0x00:0x00>
[detail] [priority <0-7>] [source-mode <nodal|noVlanMac|smltVirtual>]
[ttl-value <1-255>]
```

Example

```
Switch:1# linktrace md1 4001 13 00:bb:00:00:14:00 priority 7
Please wait for LTM to complete or press any key to abort

Received LTRs:

SeqNum: 10575 MD: md1 MA:4001    MepId:     13     Priority: 7
-------------------------------------------------------------
TTL SRC MAC                    FWDYES    TERMMEP    RELAY        ACTION
--- ------------------ ------  -------- ----------------------
63    00:bb:00:00:10:00         true        false        Fdb
62     00:bb:00:00:14:00      false        true        Hit
```

Variable Definitions

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

Variable	Value
WORD<0–22>	The first parameter, specifies the MD name.
WORD<0–22>	The second parameter, specifies the MA name.
<1–8191>	Specifies the MEP ID.
<0x00:0x00:0x00:0x00:0x00:0x00>	Specifies the target MAC address to reach the MEP.
detail	Displays linktrace result details.
priority *<0–7>*	Specifies the priority. The default is 7.
source-mode<*nodal*|*noVlanMac*|*smltVirtual*>	Specifies the source mode: • 1: nodal • noVlanMac — Use this value with C-VLAN only. When you select this option, even if a VLAN MAC address exists, the system uses the CFM C-MAC as the BMAC-SA. • 2: smltVirtual—Use this value with B-VLANs only. The default is 1: nodal.
ttl-value *<1–255>*	Specifies the Time-to-Live value. The default is 64.

Triggering a Layer 2 Ping

Use this procedure to trigger a Layer 2 ping, inside an SPBM cloud or network, which acts like native **ping**. This feature enables CFM to debug Layer 2. It can also help you debug IP shortcuts and the record for the shortcuts' ARP.

Before you begin

- You must have a MEP that is associated with a VLAN.

Procedure

Trigger a Layer 2 ping:

```
l2 ping {vlan <1-4059> routernodename WORD<0-255> | vlan <1-4059> mac
<0x00:0x00:0x00:0x00:0x00:0x00>} [burst-count <1-200>] [data-tlv-size <0-
400>] [frame-size <64-1500>] [priority <0-7>] [source-mode <nodal|
noVlanMac|smltVirtual>] [testfill-pattern <all-zero|all-zero-crc|pseudo-
random-bit-sequence|pseudo-random-bit-sequence-crc>] [time-out <1-10>]

l2 ping {ip-address WORD<0-255>} [burst-count <1-200>] [data-tlv-size <0-
400>] [frame-size <64-1500>] [priority <0-7>] [source-mode <nodal|
noVlanMac|smltVirtual>] [testfill-pattern <all-zero|all-zero-crc|pseudo-
random-bit-sequence|pseudo-random-bit-sequence-crc>] [time-out <1-10>]
[vrf WORD<1-16>]
```

Example

```
Switch:1# l2 ping vlan 2 mac 00.14.0d.bf.a3.df

Please wait for l2ping to complete or press any key to abort
----00:14:0d:bf:a3:df  L2 PING Statistics----  0(68) bytes of data
1 packets transmitted, 0 packets received,   100.00% packet loss

Switch:1# l2 ping vlan 2 routernodename MONTI0

Please wait for l2ping to complete or press any key to abort

----00:14:0d:a2:b3:df    L2 PING Statistics----  0(68) bytes of data
1 packets transmitted, 1 packets received,  0.00% packet loss
  round-trip (us)        min/max/ave/stdv =  26895/26895/26895.00/ 0.00

Switch:1# l2 ping ip-address 192.0.2.10

Please wait for l2ping to complete or press any key to abort

L2 PING  Statistics : IP 192.0.2.10, paths found 1, paths attempted 1
================================================================================
TX    RX    PERCENT  ROUND TRIP TIME
VLAN NEXT HOP                              PKTS   PKTS   LOSS     MIN/MAX/AVE
(us)
================================================================================
2  SHAMIM        (00:1a:8f:08:53:df)  1     0     100.00%  0/0/0.00
```

Variable Definitions

Use the data in the following table to configure the L2 ping parameters.

Variable	Value
{vlan *<1-4059>* routernodename *WORD<0–255>* } (vlan *<1-4059>* mac *<0x00:0x00:0x00:0x00:0x00:0x00>* } {ip-address *WORD<0–255>* }	Specifies the destination for the L2 ping: • *<1-4059>* — Specifies the VLAN ID. • WORD<0–255> — Specifies the Router node name. • <XX:XX:XX:XX:XX:XX> — Specifies the MAC address. • <A.B.C.D> — Specifies the IP address.

Table continues…

Variable	Value
burst-count *<1–200>*	Specifies the burst count.
data-tlv-size *<0–400>*	Specifies the data TLV size. The default is 0.
frame-size *<64–1500>*]	Specifies the frame size. The default is 0.
testfill-pattern *<all-zero|all-zero-crc|pseudo-random-bit-sequence|pseudo-random-bit-sequence-crc>*	Specifies the testfill pattern: • all-zero — null signal without cyclic redundancy check • all-zero-crc — null signal with cyclic redundancy check with 32-bit polynomial • pseudo-random-bit-sequence — pseudo-random-bit-sequence without cyclic redundancy check • pseudo-random-bit-sequence-crc — pseudo-random-bit-sequence with cyclic redundancy check with 32-bit polynomial. A cyclic redundancy check is a code that detects errors. The default is all-zero.
priority *<0–7>*	Specifies the priority. The default is 7.
time-out *<1–10>*	Specifies the interval in seconds. The default is 3.
source-mode<*nodal*|*noVlanMac*|*smltVirtual*>	Specifies the source mode: • 1: nodal • noVlanMac — Use this value with C-VLAN only. When you select this option, even if a VLAN MAC address exists, the system uses the CFM C-MAC as the BMAC-SA. • 2: smltVirtual—Use this value with B-VLANs only. The default is 1: nodal.
vrf *WORD<1–16>*	Specifies the VRF name.

Trigger a Layer 2 Traceroute

Note:

This procedure only applies to VSP 4000 Series, VSP 4900 Series, VSP 7200 Series, VSP 7400 Series, VSP 8200 Series, VSP 8400 Series, and VSP 8600 Series.

Use this procedure to trigger a Layer 2 traceroute, which acts like native **`traceroute`**. This feature enables CFM to debug Layer 2 in an SPBM cloud or network. It can determine the path used by IS—IS to get from one MEP to another, by showing all the hops between. Therefore, it can show where connectivity is lost. It can also work for IP shortcuts.

Important:

To trace a route to a MAC address, the MAC address must be in the VLAN FDB table.

- For B-VLANs, you do not have to trigger an **l2ping** to learn the MAC address because IS-IS populates the MAC addresses in the FDB table.

linktrace traces the path up to the closest device to that MAC address that supports CFM.

Before you begin

- You must have a MEP that is associated with a VLAN.

Procedure

Trigger a Layer 2 traceroute:

```
l2 traceroute {<vlan <1-4059> routernodename WORD<0-255> | <vlan <1-4059>
mac <0x00:0x00:0x00:0x00:0x00:0x00>} [priority <0-7>] [source-mode
<nodal|noVlanMac|smltVirtual>] [ttl <1-255>]

l2 traceroute {ip-address WORD<0-255>} [priority <0-7>][source-mode
<nodal|noVlanMac|smltVirtual>][ttl <1-255>] [vrf WORD<1-16]
```

Examples

```
Switch:1# l2 traceroute vlan 2 routernodename Switch-MONTIO

Please wait for l2traceroute to complete or press any key to abort

l2traceroute to Switch-MONTI0  (00:14:0d:a2:b3:df),  vlan 2
0    Switch-PETER4         (00:15:9b:11:33:df)
1    Switch-MONTI0          (00:14:0d:a2:b3:df)

Switch:1# l2 traceroute ip-address 192.0.2.10

Please wait for l2trace to complete or press any key to abort

L2 Trace  Statistics : IP 192.0.2.10, paths found 1
======================================================================
Switch-SHAMIM  (00:1a:8f:08:53:df),  vlan 2
0    Switch-PETER4          (00:15:9b:11:33:df)
1    Switch-MONTI0          (00:14:0d:a2:b3:df)
```

Variable Definitions

Use the data in the following table to use the **l2 traceroute** command.

Variable	Value
{vlan *<1-4059>* routernodename *WORD<0–255>* } (vlan *<1-4059>* mac *<0x00:0x00:0x00:0x00:0x00:0x00>* } {ip-address *WORD<0–255>* }	Specifies the destination for the L2 traceroute: • *<1-4059>* — Specifies the VLAN ID • WORD<0–255> — Specifies the Router Node Name • <XX:XX:XX:XX:XX:XX> — Specifies the MAC address • WORD<0–255> — Specifies the IP address

Table continues…

Variable	Value
ttl-value *<1–255>*	Specifies the TTL value. The default is 64.
priority *<0–7>*	Specifies the priority. The default is 7.
source-mode<*nodal*|*noVlanMac*| *smltVirtual*>	Specifies the source mode: • 1: nodal • noVlanMac — Use this value with C-VLAN only. When you select this option, even if a VLAN MAC address exists, the system uses the CFM C-MAC as the BMAC-SA. • 2: smltVirtual—Use this value with B-VLANs only. The default is 1: nodal.
vrf *WORD<1–16*	Specifies the VRF name.

Trigger a Layer 2 Tracetree

Note:

This procedure only applies to VSP 4000 Series, VSP 4900 Series, VSP 7200 Series, VSP 7400 Series, VSP 8200 Series, VSP 8400 Series, and VSP 8600 Series.

Use this procedure to trigger a Layer 2 tracetree. Layer 2 tracetree allows a user to trigger a multicast LTM message by specifying the B-VLAN and I-SID. The command allows the user to view a multicast tree on the SPBM B-VLAN from the source node to the destination nodes for a particular I-SID.

Note:

This command is supported on SPBM B-VLANs only, not C-VLANs.

Before you begin

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP.
- Enable the MEP.
- Assign a nodal MEP to the B-VLAN.

Procedure

Trigger a Layer 2 tracetree:

```
l2 tracetree {<1-4059> <1-16777215> [routernodename WORD<0-255> |
<1-4059> <1-16777215>] [mac <0x00:0x00:0x00:0x00:0x00:0x00>}] [priority
<0-7>] [source-mode <nodal|noVlanMac|smltVirtual>] [ttl—value <1-255>]
```

Example

```
Switch:1# l2 tracetree 500 1
```

```
Switch:1# l2 tracetree 500 1

Please wait for l2tracetree to complete or press any key to abort
```

```
l2tracetree to 53:55:10:00:00:01, vlan 500 i-sid 1 nickname 5.55.10
hops 64
1   Switch-PETER4        00:15:9b:11:33:df -> Switch-MONTI0        00:14:0d:a2:b3:df
2   Switch-MONTI0        00:14:0d:a2:b3:df -> Switch-LEE2          00:15:e8:b8:a3:df
```

Variable Definitions

Use the data in the following table to use the **`l2 tracetree`** command.

Variable	Value
{ *<1-4059><1–16777215> routernodename WORD<0–255>* | *<1-4059><1–16777215> mac <0x00:0x00:0x00:0x00:0x00:0x00>*}	• *<1-4059>* — Specifies the VLAN ID. • <1–16777215> — Specifies the I-SID. • WORD<0–255> — Specifies the Router Node Name. • <0x00:0x00:0x00:0x00:0x00:0x00> — Specifies the MAC address.
ttl-value *<1–255>*	Specifies the TTL value. The default is 64.
priority *<0–7>*	Specifies the priority value. The default is 7.
source-mode<*nodal*|*noVlanMac*| *smltVirtual*>	Specifies the source mode: • 1: nodal • 2: smltVirtual The default is nodal.

Triggering a Layer 2 Tracetree-fan

About this task

Use this procedure to trigger a Layer 2 tracetree-fan from the nickname server to make sure that all the nodes towards the nickname client support the FAN protocol. Layer 2 tracetree-fan allows a user to trigger an LTM on the internal Fabric Area Network (FAN) I-SID. This command allows the user to trace the FAN tree.

Procedure

Trigger a Layer 2 tracetree-fan:

```
l2 tracetree—fan [mac <0x00:0x00:0x00:0x00:0x00:0x00>] [priority <0-7>]
[routernodename WORD<0-255>] [ttl—value <1-255>]
```

Example

```
Switch:1# l2 tracetree-fan
```

```
Switch:1# l2 tracetree-fan

Please wait for l2tracetree to complete or press any key to abort


l2tracetree to b1:ad:aa:41:b0:84, vlan 4051 i-sid 16777001 nickname 0.00.00 hops 64
```

```
1    Switch-PETER4        b0:ad:aa:41:b0:84 -> Switch-MONTI0        b0:ad:aa:41:48:84
1    Switch-PETER4        b0:ad:aa:41:b0:84 -> Switch-MONTI1        b0:ad:aa:42:88:84
2    Switch-MONTI0        b0:ad:aa:41:48:84 -> Switch-LEE2          b0:ad:aa:43:3c:84
```

Trigger a Layer 2 Tracemroute

Note:

This procedure only applies to VSP 4000 Series, VSP 4900 Series, VSP 7200 Series, VSP 7400 Series, VSP 8200 Series, VSP 8400 Series, and VSP 8600 Series.

Use this procedure to debug the IP Multicast over Fabric Connect stream path using **l2 tracemroute** on the VLAN (Layer 2) or the VRF (Layer 3). This procedure queries the SPBM multicast module to determine the B-VLAN, I-SID and nickname for the S and G streams. The nickname and I-SID are used to create a multicast MAC address.

Note:

The VLAN option is only valid for a VLAN that has an I-SID configured and IGMP snooping enabled.

Before you begin

- On the source and destination nodes, you must configure an autogenerated or an explicit CFM MD, MA, and MEP.
- Enable the MEP.
- Assign a nodal MEP to the B-VLAN.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Trigger a Layer 2 tracemroute on the VLAN:

   ```
   l2 tracemroute source <A.B.C.D> group <A.B.C.D> vlan
   <1-4059>[priority <0-7>] [ttl—value <1-255>]
   ```

   Note:

   For the preceding command, if you do not specify a VLAN, **l2 tracemroute** uses the global default VRF.

   Wait for the l2 tracemroute to complete or press any key to abort.
3. Trigger a Layer 2 tracemroute on the VRF:

   ```
   l2 tracemroute source <A.B.C.D> group <A.B.C.D> vrf WORD<1-16>
   [priority <0-7>] [ttl—value <1-255>]
   ```

Note:

For the preceding command, if you do not specify a VRF, **l2 tracemroute** uses the global default VRF.

Wait for the l2 tracemroute to complete or press any key to abort.

Examples

The following is a sample output for a Layer 2 tracemroute on a VLAN:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#l2 tracemroute source 192.0.2.81 233.252.0.1 vlan 201

Please wait for l2 tracemroute to complete or press any key to abort.

Source 192.0.2.81

Group: 233.252.0.1

VLAN:201

BMAC: 03:00:03:f4:24:01

B-VLAN: 10

I-SID: 16000001

================================================================
1 PETER4 00:03:00:00:00:00 -> LEE1 00:14:0d:bf:a3:df
2 LEE1 00:14:0d:bf:a3:df -> LEE2 00:15:e8:b8:a3:df
```

The following is a sample output for a Layer 2 tracemroute on a VRF:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#l2 tracemroute source 192.0.2.10 group 233.252.0.1 vrf red

Please wait for l2 tracemroute to complete or press any key to abort.

Source 192.0.2.10

Group: 233.252.0.1

VRF: redID 1

BMAC: 03:00:04:f4:24:01

B-VLAN: 20

I-SID: 16000001

================================================================
1 PETER4 00:03:00:00:00:00 -> LEE1 00:14:0d:bf:a3:df
2 LEE1 00:14:0d:bf:a3:df -> LEE2 00:15:e8:b8:a3:df
```

Variable Definitions

Use the data in the following table to use the **l2 tracemroute** command.

Variable	Value
source *<A.B.C.D>*	Specifies the source IP address.
group *<A.B.C.D>*	Specifies the IP address of the multicast group.
vlan *<1–4084>*	Specifies the VLAN value.
vrf *WORD<1–16>*	Specifies the VRF name. If you do not specify a VRF name, then the results are shown for the flow in the Global Router (default) context.
priority *<0–7>*	Specifies the priority value.
ttl *<1–255>*	Specifies the time-to-live (TTL) for the trace packet, which is how many hops the trace packet takes before it is dropped.

Job Aid

The following table describes the fields in the output for **`l2 tracemroute`** command for a VLAN.

Parameter	Description
Source	Specifies the source IP address of the flow where the multicast trace tree originates.
Group	Specifies the IP address of the multicast group.
VLAN	Specifies the VLAN.
BMAC	Specifies the backbone MAC address.
B-VLAN	Specifies the backbone VLAN.
I-SID	Specifies the service identifier.

The following table describes the fields in the output for **`l2 tracemroute`** command for a VRF.

Parameter	Description
Source	Specifies the source IP address of the flow where the multicast trace tree originates.
Group	Specifies the IP address of the multicast group.
VRF	Specifies the VRF.
BMAC	Specifies the backbone MAC address.
B-VLAN	Specifies the backbone VLAN.
I-SID	Specifies the service identifier.

Using trace CFM to Diagnose Problems

Use the following procedure to display trace information for CFM.

About this task

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

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

Use the **`trace level 120 <0-4>`** command to trace CFM module information, including CLI, instrumentation, show config, and platform dependent code. The CFM module ID is 120.

Use the **`trace cfm level <0-4>`** command to trace platform independent code and CFM protocol code.

Caution:

Risk of traffic loss

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

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Clear the trace:

   ```
   clear trace
   ```

3. Begin the trace operation:

   ```
   trace cfm level <0-4>
   ```

   Wait approximately 30 seconds, and then stop trace.

4. Stop tracing:

   ```
   trace shutdown
   ```

5. View the trace results:

   ```
   show trace cfm
   ```

6. Begin the trace operation for the CFM module:

   ```
   trace level 120 <0-4>
   ```

   Wait approximately 30 seconds, and then stop trace.

7. View trace results:

   ```
   trace screen enable
   ```

   Important:

   If you use trace level 3 (verbose) or trace level 4 (very verbose), it is recommended that you do not use the screen to view commands due to the volume of information the system generates and the effect on the system.

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

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

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

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

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

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

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)# clear trace
Switch:1(config)# trace cfm level 3
Switch:1(config)# trace shutdown
Switch:1(config)# show trace cfm
================================================================================
                                CFM Tracing Info
================================================================================
 Status    : Enabled
 Level     : VERBOSE
Switch:1(config)#trace level 120 3
Switch:1(config)# save trace
Switch:1(config)# trace grep error
Switch:1(config)#trace grep 00-1A-4B-8A-FB-6B
```

Variable Definitions

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

Variable	Value
cfm level *[<0-4>]*	Starts the trace by specifying the level. • *<0–4>* specifies the trace level from 0–4, where 0 is disabled; 1 is very terse; 2 is terse; 3 is verbose, 4 is very verbose.
filter	Configures a filter trace for a file or module.
flags	Configures trace flags for IS-IS or OSPF.
grep *[WORD<0-128>]*	Searches trace results for a specific string value, for example, the word error. Performs a comparison of trace messages.
level [<*Module_ID*>] *[<0-4>]*	Starts the trace by specifying the module ID and level. • *<Module_ID>* specifies the module for the trace. Different hardware platforms support different ID ranges because of feature support differences. To see which module IDs are available on the switch, use the **`show trace modid-list`** command or CLI command completion Help. • <0-4> specifies the trace level from 0–4, where 0 is disabled; 1 is very terse; 2 is terse; 3 is verbose, 4 is very verbose.

Table continues…

Variable	Value
route-map	Enables or disables the trace route-map. The values are on and off.
screen {disable|enable}	Enables the display of trace output to the screen.
shutdown	Stops the trace operation.
spbm isis level *[<0-4>]*	Starts the trace by specifying the level. • <0-4> specifies the trace level from 0–4, where 0 is disabled; 1 is very terse; 2 is terse; 3 is verbose, 4 is very verbose. The default is 1, very terse.

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

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

Using trace SPBM to Diagnose Problems

Use the following procedure to display trace information for SPBM IS-IS. In the case of IS-IS, this procedure also provides information related to the flags set.

About this task

Use the **`trace level 119 <0–4>`** command to trace IS-IS module information, including CLI, instrumentation, show config and platform dependent code. The IS-IS module ID is 119.

Use the **`trace level 125 <0–4>`** command to trace SPBM module information, including CLI, instrumentation, show config and platform dependent code. The SPBM module ID is 125.

Use the **`trace spbm isis level`** command to trace platform independent code, IS-IS protocol, IS-IS hello, IS-IS adjacency, LSP processing, and IS-IS computation.

Caution:

Risk of traffic loss

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

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Clear the trace:

   ```
   clear trace
   ```

3. Begin the trace operation:

   ```
   trace spbm isis level <0-4>
   ```

   Wait approximately 30 seconds, and then stop trace.

4. Stop tracing:

   ```
   trace shutdown
   ```

5. Display the trace information for SPBM IS-IS:

   ```
   show trace spbm isis
   ```

6. Begin the trace operation for the SPBM module:

   ```
   trace level 125 <0-4>
   ```

   Wait approximately 30 seconds, and then stop trace.

7. Begin the trace operation for the IS-IS module:

   ```
   trace level 119 <0-4>
   ```

   Wait approximately 30 seconds, and then stop trace.

8. View trace results:

   ```
   trace screen enable
   ```

   Important:

   If you use trace level 3 (verbose) or trace level 4 (very verbose), it is recommended that you do not use the screen to view commands due to the volume of information the system generates and the effect on the system.

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

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

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

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

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

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

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)# clear trace
Switch:1(config)# trace spbm isis level 3
Switch:1(config)# trace shutdown
Switch:1(config)# show trace spbm isis
================================================================================
                              SPBM ISIS Tracing Info
================================================================================
 Status    : Enabled
 Level     : VERY_TERSE
 Flag Info :
Switch:1(config)#trace level 125 3
Switch:1(config)#trace level 119 3
Switch:1(config)# save trace
Switch:1(config)# trace grep error
Switch:1(config)#trace grep 00-1A-4B-8A-FB-6B
```

Variable Definitions

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

Variable	Value
cfm level *[<0-4>]*	Starts the trace by specifying the level. • *<0–4>* specifies the trace level from 0–4, where 0 is disabled; 1 is very terse; 2 is terse; 3 is verbose, 4 is very verbose.
filter	Configure a filter trace for a file or module.
flags	Configure trace flags for IS-IS or OSPF.
grep *[WORD<0-128>]*	Searches trace results for a specific string value, for example, the word error. Performs a comparison of trace messages.
level [<*Module_ID*>] *[<0-4>]*	Starts the trace by specifying the module ID and level. • *<Module_ID>* specifies the module for the trace. Different hardware platforms support different ID ranges because of feature support differences. To see which module IDs are available on the switch, use the **`show trace modid-list`** command or CLI command completion Help. • <0-4> specifies the trace level from 0–4, where 0 is disabled; 1 is very terse; 2 is terse; 3 is verbose, 4 is very verbose.
route-map	Enables or disables the trace route-map. The values are on and off.
screen {disable|enable}	Enables the display of trace output to the screen.
shutdown	Stops the trace operation.

Table continues…

Variable	Value
spbm isis level *[<0-4>]*	Starts the trace by specifying the level. • <0-4> specifies the trace level from 0–4, where 0 is disabled; 1 is very terse; 2 is terse; 3 is verbose, 4 is very verbose. The default is 1, very terse.

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

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

CFM Configuration Using EDM

This section provides procedures to configure Connectivity Fault Management (CFM) using Enterprise Device Manager (EDM).

Note:

When you enable CFM in an SPBM network, it is recommended that you enable CFM on the Backbone Edge Bridges (BEB) and on all Backbone Core Bridges (BCB). If you do not enable CFM on a particular node, you cannot obtain CFM debug information from that node.

Autogenerated CFM

CFM provides two methods for creating MEPs: autogenerated and explicit. You cannot use both. You must choose one or the other. Use the procedures in this section to configure autogenerated MEPs that eliminate the need to configure an MD, MA, and MEP ID to create a MEP.

- For SPBM B-VLANs, you can use either autogenerated or explicitly configured CFM MEPs.

- For C-VLANs, you can only use autogenerated CFM MEPs.

Previous explicit CFM configurations of MDs, MAs and MEPs on SPBM B-VLANs continue to function. However, if you want to enable the autogenerated commands, you must first remove the existing MEP and MIP on the SPBM B-VLAN. The switch only supports one MEP or MIP on the SPBM B-VLAN, either explicitly configured or autogenerated.

For autogenerated CFM configuration information for EDM, see the following tasks:

- Configuring autogenerated CFM on SPBM B-VLANs on page 111
- Configuring autogenerated CFM on C-VLANs on page 112

Configuring Autogenerated CFM on SPBM B-VLANs

Use this procedure to configure the autogenerated CFM MEP and MIP level for every SPBM B-VLAN on the chassis. This configuration eliminates the need to explicitly configure an MD, MA, and MEP ID and to associate the MEP and MIP level to the SPBM B-VLAN.

To configure autogenerated CFM on C-VLANs, see Configuring autogenerated CFM on C-VLANs on page 112.

About this task

When you enable this feature, the device creates a global MD (named spbm) for all the SPBM Nodal MEPs. This MD has a default maintenance level of 4, which you can change with the level attribute. All the MEPs that the device creates use the MEP ID configured under the global context, which has a default value of 1. The nodal MEPs are automatically associated with the SPBM B-VLANs configured. The MIP level maps to the global level. When you enable the feature, the device automatically associates the MIP level with the SPBM B-VLANs configured. The feature is disabled by default.

Important:

CFM supports one MEP or MIP for each SPBM B-VLAN only. This means that if you want to use these autogenerated MEPs, you cannot use your existing CFM configuration. You must first remove the existing MEP or MIP on the SPBM B-VLAN. If you want to continue configuring MEPs manually, skip this procedure.

Procedure

1. In the navigation pane, expand the **Configuration** > **Edit** > **Diagnostics** folders.
2. Click **CFM**.
3. Click the **Global** tab.
4. Select **enable** next to **SpbmAdminState**.
5. Click **Apply**.
6. To verify the values assigned to MA, MD, and MEP, perform the following steps:
   a. Click the **MD** tab.
   b. Select **SPBM**, and then check the MA and MEP values.

CFM Global Field Descriptions

Use the data in the following table to configure the global MEP and MIP parameters.

Name	Description
SpbmAdminState	Enables or disables autogenerated CFM for B-VLANs. The default is disable.
SpbmLevel	Specifies the global SPBM CFM maintenance level for the chassis within the range of 0 to 7. The default is 4. Only configure global CFM at one MD level for each chassis for each VLAN type.
SpbmMepId	Specifies the global MEP ID within the range of 1 to 8191. Select a unique ID for each switch to ensure that the MEPs are unique across the network. The default is 1.
CmacAdminState	Enables or disables autogenerated CFM for C-VLANs. The default is disable. This field does not appear for all hardware platforms.
CmacLevel	Specifies the global C-MAC CFM maintenance level for the chassis within the range of 0 to 7. The default is 4. Only configure global CFM at one MD level for each chassis for each VLAN type. This field does not appear for all hardware platforms.
CmacMepId	Specifies the global MEP ID within the range of 1 to 8191. Select a unique ID for each switch to ensure that the MEPs are unique across the network. The default is 1. This field does not appear for all hardware platforms.
Bmac	Displays the B-MAC address of the node. This field does not appear for all hardware platforms.
Cmac	Displays the C-MAC address of the node. This field does not appear for all hardware platforms.

Configure Autogenerated CFM on C-VLANs

Use this procedure to configure the autogenerated CFM MEP and MIP level for every C-VLAN on the chassis.

To configure autogenerated CFM on SPBM B-VLANs, see Configuring autogenerated CFM on SPBM B-VLANs on page 111.

Note:

For C-VLANs, you can only use autogenerated CFM MEPs.

Important:

CFM supports one MEP or MIP on each C-VLAN. Only autogenerated CFM provides support for configuring MEP and MIPs on C-VLANs. You cannot explicitly configure C-VLANs.

About this task

When you enable this feature, you create a global MD (named cmac) for all the customer MAC (C-MAC) MEPs. This MD has a default maintenance level of 4, which you can change with the level attribute. The autogenerated CFM commands also create an MA for each C-VLAN, a MEP for each C-VLAN, and associate the MEP with the corresponding C-VLAN and a MIP with the C-VLAN.

All the MEPs that the device creates use the MEP ID configured under the global context, which has a default value of 1. The device automatically associates the MEPs with the C-VLANs configured. The MIP level maps to the global level. The device automatically associates the MIP level with the C-VLANs configured when you enable the feature.

The feature is disabled by default.

Procedure

1. In the navigation pane, expand the **Configuration** > **Edit** > **Diagnostics** folders.
2. Select **CFM**.
3. Select the **Global** tab.
4. Select **enable** next to **CmacAdminState**.
5. In the fields provided, specify a maintenance level and a MEP ID.
6. Select **Apply**.

CFM Global Field Descriptions

Use the data in the following table to configure the global MEP and MIP parameters.

Name	Description
SpbmAdminState	Enables or disables autogenerated CFM for B-VLANs. The default is disable.
SpbmLevel	Specifies the global SPBM CFM maintenance level for the chassis within the range of 0 to 7. The default is 4. Only configure global CFM at one MD level for each chassis for each VLAN type.
SpbmMepId	Specifies the global MEP ID within the range of 1 to 8191. Select a unique ID for each switch to ensure that the MEPs are unique across the network. The default is 1.
CmacAdminState	Enables or disables autogenerated CFM for C-VLANs. The default is disable. This field does not appear for all hardware platforms.

Table continues…

Name	Description
CmacLevel	Specifies the global C-MAC CFM maintenance level for the chassis within the range of 0 to 7. The default is 4. Only configure global CFM at one MD level for each chassis for each VLAN type. This field does not appear for all hardware platforms.
CmacMepId	Specifies the global MEP ID within the range of 1 to 8191. Select a unique ID for each switch to ensure that the MEPs are unique across the network. The default is 1. This field does not appear for all hardware platforms.
Bmac	Displays the B-MAC address of the node. This field does not appear for all hardware platforms.
Cmac	Displays the C-MAC address of the node. This field does not appear for all hardware platforms.

Configuring Explicit CFM

For SPBM B-VLANs, CFM provides two methods for creating MEPs: autogenerated and explicit. You cannot use both. Use the procedures in this section to configure MEPs explicitly.

If you want to create autogenerated CFM MEPs that eliminate the need to configure an MD, MA, and MEP ID, see the procedures in Autogenerated CFM on page 110. For C-VLANs, you can only use the autogenerated method.

Note:

The CFM show commands that display MD, MA, and MEP information work for both autogenerated and explicitly-configured CFM MEPs.

Configuring CFM MD

Use this procedure to configure a Connectivity Fault Management (CFM) Maintenance Domain (MD). An MD is the part of a network that is controlled by a single administrator. A single MD can contain several Maintenance Associations (MA).

Procedure

1. In the navigation pane, expand the **Configuration** > **Edit** > **Diagnostics** folders.
2. Click **CFM**.
3. Click the **MD** tab.
4. Click **Insert**.
5. In the fields provided, specify an index value, name, and level for the MD.

6. Click **Insert**.

MD Field Descriptions

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

Name	Description
Index	Specifies a maintenance domain entry index.
Name	Specifies the MD name.
NumOfMa	Indicates the number of MAs that belong to this maintenance domain.
Level	Specifies the MD maintenance level. The default is 4.
NumOfMip	Indicates the number of MIPs that belong to this maintenance domain
Type	Indicates the type of domain.

Configuring CFM MA

Use this procedure to configure a CFM Maintenance Association (MA). An MA represents a logical grouping of monitored entities within its Domain. It can therefore represent a set of Maintenance Endpoints (MEPs), each configured with the same Maintenance Association ID (MAID) and MD Level, established to verify the integrity of a single service instance.

Before you begin

- You must configure a CFM MD.

Procedure

1. In the navigation pane, expand the **Configuration** > **Edit** > **Diagnostics** folders.
2. Click **CFM**.
3. Click the **MD** tab.
4. Highlight an existing MD, and then click **MaintenanceAssociation**.
5. In the **MA** tab, click **Insert**.
6. In the fields provided, specify an index value and name for the MA.
7. Click **Insert**.

MA Field Descriptions

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

Name	Description
DomainIndex	Specifies the maintenance domain entry index.
AssociationIndex	Specifies a maintenance association entry index.
DomainName	Specifies the MD name.

Table continues…

Name	Description
AssociationName	Specifies the MA name.
NumOfMep	Indicates the number of MEPs that belong to this maintenance association.

Configuring CFM MEP

Use this procedure to configure the CFM Maintenance Endpoint (MEP). A MEP represents a managed CFM entity, associated with a specific Domain Service Access Point (DoSAP) of a service instance, which can generate and receive CFM Protocol Data Units (PDU) and track any responses. A MEP is created by MEP ID under the context of an MA.

Procedure

1. In the navigation pane, expand the **Configuration** > **Edit** > **Diagnostics** folders.
2. Click **CFM**.
3. Click the **MD** tab.
4. Highlight an existing MD, and then click **MaintenanceAssociation**.
5. In the **MA** tab, highlight an existing MA, and then click **MaintenanceEndpoint**.
6. Click **Insert**.
7. In the fields provided, specify the ID and the administrative state of the MEP.
8. Click **Insert**.

MEP Field Descriptions

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

Name	Description
DomainIndex	Specifies the MD index.
AssociationIndex	Specifies the MA index.
Id	Specifies the MEP ID.
DomainName	Specifies the MD name.
AssociationName	Specifies the MA name.
AdminState	Specifies the administrative state of the MEP. The default is disable.
MepType	Specifies the MEP type: • trunk • sg • endpt • vlan • port

Table continues…

Name	Description
	• endptClient • nodal • remotetrunk • remotesg • remoteendpt • remoteVlan • remotePort • remoteEndptClient
ServiceDescription	Specifies the service to which this MEP is assigned.

Configuring CFM Nodal MEP

Use this procedure to configure the CFM nodal Maintenance Endpoint (MEP). The Nodal MEP provides traceability and troubleshooting at the system level for a given B-VLAN. The Nodal B-VLAN MEPs created on the CP and function as if they are connected to the virtual interface of the given B-VLAN. Because of this they are supported for both port and MLT based B-VLANs.

Nodal MPs provide both MEP and Maintenance Intermediate Point (MIP) functionality for SPBM deployments. Nodal MPs are associated with a B-VLAN and are VLAN encapsulated packets. Each node (chassis) has a given MAC address and communicates with other nodes. The SPBM instance MAC address is used as the MAC address of the Nodal MP.

Before you begin

- You must configure a CFM MD, MA, and MEP.

Procedure

1. In the navigation pane, expand the **Configuration** > **VLAN** folders.
2. Click **VLANs**.
3. Click the **Advanced** tab.
4. Select a VLAN with a type of spbm-bvlan.
5. Click **Nodal**.
6. In the **NodalMepList** field, specify the nodal MEPs to add to the VLAN.
7. Click **Apply**.

Nodal MEP/MIP Field Descriptions

Use the data in the following table to use the **Nodal MEP/MIP** tab.

Name	Description
NodalMepList	Specifies the nodal MEPs to add to the VLAN, in the format <mdName.maName.mepId>, for example md10.ma20.30.
NumOfNodalMep	Indicates the number of nodal MEPs assigned to this VLAN.
NodalMipLevelList	Specifies a MIP level list.
NumOfNodalMipLevel	Indicates the number of nodal MIP levels assigned to this VLAN that allows MIP functionality to be enabled on a per level per VLAN basis.

Configuring Layer 2 Ping

Use this procedure to configure a Layer 2 ping inside an SPBM cloud or network. This feature enables CFM to debug Layer 2. It can also help you debug IP shortcuts and the record for the shortcuts' ARP.

Before you begin

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP, and assign a nodal MEP to the B-VLAN.

Procedure

1. In the navigation tree, expand the **Configuration** > **Edit** > **Diagnostics** folders.
2. Click **L2Ping/L2Trace Route**.
3. From the **L2Ping** tab, configure the Layer 2 ping properties.
4. To initiate a Layer 2 ping, highlight an entry and click the **Start** button.
5. To update a Layer 2 ping, click the **Refresh** button.
6. To stop the Layer 2 ping, click the **Stop** button.

L2Ping Field Descriptions

Use the data in the following table to use the **L2Ping** tab.

Name	Description
VlanId	Identifies the backbone VLAN.
DestMacAddress	Specifies the target MAC address.
HostName	Specifies the target host name.
DestIsHostName	Indicates whether the host name is (true) or is not (false) used for L2Ping transmission.
Messages	Specifies the number of L2Ping messages to be transmitted. The default is 1.

Table continues…

Name	Description
Status	Specifies the status of the transmit loopback service: • ready: the service is available. • transmit: the service is transmitting, or about to transmit, the L2Ping messages. • abort: the service aborted or is about to abort the L2Ping messages. This field is also used to avoid concurrency or race condition problems that can occur if two or more management entities try to use the service at the same time. The default is ready.
ResultOk	Indicates the result of the operation: • true: the L2Ping Messages will be (or have been) sent. • false: the L2Ping Messages will not be sent. The default is true.
Priority	Specifies a 3–bit value to be used in the VLAN header, if present in the transmitted frame. The default is 7.
TimeoutInt	Specifies the interval to wait for an L2Ping time-out. The default value is 3 seconds.
TestPattern	Specifies the test pattern to use in the L2Ping PDU: • allZero: null signal without cyclic redundancy check • allZeroCrc: null signal with cyclic redundancy check with 32-bit polynomial • pseudoRandomBitSequence: pseudo-random-bit-sequence without cyclic redundancy check • pseudoRandomBitSequenceCrc: pseudo-random-bit-sequence with cyclic redundancy check with 32-bit polynomial. A cyclic redundancy check is a code that detects errors. The default value is allZero.
DataSize	Specifies an arbitrary amount of data to be included in the data TLV, if the data size is selected to be sent. The default is 0.
FrameSize	Specifies the frame size. If the frame size is specified then the data size is internally calculated and the calculated data size is included in the data TLV. The default is 0.

Table continues…

Name	Description
SourceMode	Specifies the source mode of the transmit loopback service: • nodal • noVlanMac — Use this value with C-VLAN only. When you select this option, even if a VLAN MAC address exists, the system uses the CFM C-MAC as the BMAC-SA. • smltVirtual — Use the smltVirtual option with B-VLANs only. The default is nodal.
SeqNumber	The transaction identifier/sequence number of the first loopback message (to be) sent. The default is 0.
Result	Displays the Layer 2 Ping result.

Initiate a Layer 2 Traceroute

Note:

This procedure only applies to VSP 4000 Series, VSP 4900 Series, VSP 7200 Series, VSP 7400 Series, VSP 8200 Series, VSP 8400 Series, and VSP 8600 Series.

Use this procedure to trigger a Layer 2 traceroute. This feature enables CFM to debug Layer 2 in an SPBM cloud or network. It can determine the path used by IS—IS to get from one MEP to another, by showing all the hops between. Therefore, it can show where connectivity is lost. It can also work for IP shortcuts.

If you configure **IsTraceTree** to false then EDM performs Traceroute on the unicast path. If you configure **IsTraceTree** to true then EDM performs TraceTree on the multicast tree.

For more information on configuring tracetree, see Configuring Layer 2 tracetree on page 136.

Important:

To trace a route to a MAC address, the MAC address must be in the VLAN FDB table.

For B-VLANs, you do not have to trigger an **L2Ping** to learn the MAC address because IS-IS populates the MAC addresses in the FDB table.

Linktrace traces the path up to the closest device to that MAC address that supports CFM.

Before you begin

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP, and assign a nodal MEP to the B-VLAN.

Procedure

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

2. Select **L2Ping/L2Trace Route**.
3. Select the **L2 Traceroute/TraceTree** tab.
4. To start the traceroute, highlight an entry, and then select **Start**.
5. To update the traceroute, select **Refresh**.
6. To stop the traceroute, select **Stop**.

L2 Traceroute Field Descriptions

Use the data in the following table to use the **L2 Traceroute/TraceTree** tab.

Name	Description
VlanId	Specifies a value that uniquely identifies the Backbone VLAN (B-VLAN).
Priority	Specifies a 3–bit value to be used in the VLAN header, if present in the transmitted frame. The default is 7.
DestMacAddress	Specifies the target MAC address.
HostName	Specifies the target host name.
DestIsHostName	Specifies whether the host name is (true) or is not (false) used for the L2Trace transmission.
Isid	Specifies the Service Instance Identifier (I-SID).
IsTraceTree	Specifies whether the multicast tree or unicast path is traced: • If you configure **IsTraceTree** to false then EDM performs Traceroute on the unicast path. • If you configure **IsTraceTree** to true then EDM performs TraceTree on the multicast tree.
Status	Indicates the status of the transmit loopback service: • ready: the service is available. • transmit: the service is transmitting, or about to transmit, the L2Trace messages. • abort: the service aborted or is about to abort the L2Trace messages. This field is also used to avoid concurrency or race condition problems that can occur if two or more management entities try to use the service at the same time. The default is ready.

Table continues…

Name	Description
ResultOk	Indicates the result of the operation: • true: the L2Trace messages will be (or have been) sent. • false: the L2Trace messages will not be sent. The default is true.
Ttl	Specifies the number of hops remaining to this L2Trace. This value is decremented by 1 by each Bridge that handles the L2Trace. The decremented value is returned in the L2Trace. If 0 on output, the L2Trace is not transmitted to the next hop. The value of the time-to-live (TTL) field in the L2Trace is defined by the originating MEP. The default value is 64.
SourceMode	Specifies the source mode: • 1: nodal • noVlanMac — Use this value with C-VLAN only. When you select this option, even if a VLAN MAC address exists, the system uses the CFM C-MAC as the BMAC-SA. • 2: smltVirtual—Use this value with B-VLANs only. The default is 1: nodal.
SeqNumber	Specifies the transaction identifier/sequence number of the first loopback message (to be) sent. The default is 0.
Flag	L2Trace result flag that indicates L2Trace status or error code: • none (1): No error • internalError (2): L2Trace internal error • invalidMac (3): Invalid MAC address • mepDisabled (4): MEP must be enabled in order to perform L2Trace • noL2TraceResponse (5): No L2Trace response received • l2TraceToOwnMepMac (6): L2Trace to own MEP MAC is not sent • l2TraceComplete (7): L2Trace completed

Table continues…

Name	Description
	• l2TraceLookupFailure (8): Lookup failure for L2Trace • l2TraceLeafNode (9): On a leaf node in the I-SID tree • l2TraceNotInTree (10): Not in the I-SID tree • l2TraceSmltNotPrimary (11): Requested SMLT source from non-primary node

View Layer 2 Traceroute Results

Note:

This procedure only applies to VSP 4000 Series, VSP 4900 Series, VSP 7200 Series, VSP 7400 Series, VSP 8200 Series, VSP 8400 Series, and VSP 8600 Series.

Use this procedure to view Layer 2 traceroute results. This feature enables CFM to debug Layer 2. It can also help you debug ARP problems by providing the ability to troubleshoot next hop ARP records.

About this task

You can display Layer 2 tracetree results to view a multicast tree on the SPBM B-VLAN from the source node to the destination nodes for a particular I-SID. For more information, see Viewing Layer 2 tracetree results on page 139.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **Diagnostics**.
2. Select **L2Ping/L2Trace Route**.
3. Select the **L2Traceroute/TraceTree** tab.
4. Select **Refresh** to update the results.
5. To view the traceroute results, highlight an entry, and then select **Result**.

L2 Traceroute/Tracetree Result Field Descriptions

Use the data in the following table to use the **L2 Traceroute/Tracetree Result** tab.

Name	Description
VlanId	A value that uniquely identifies the Backbone VLAN (B-VLAN).
SeqNumber	The transaction identifier/sequence number returned by a previous transmit linktrace message command, indicating which L2Trace's response of the L2Trace is going to be returned. The default is 0.

Table continues…

Name	Description
Hop	The number of hops away from L2Trace initiator.
ReceiveOrder	An index to distinguish among multiple L2Trace responses with the same Transaction Identifier field value. This value is assigned sequentially from 1, in the order that the Linktrace Initiator received the responses.
Ttl	Time-to-Live (TTL) field value for a returned L2Trace response.
SrcMac	MAC address of the MP that responds to the L2Trace request for this L2TraceReply.
HostName	The host name of the replying node.
LastSrcMac	The MAC address of the node that forwarded the L2Trace to the responding node.
LastHostName	The host name of the node that forwarded the L2Trace to the responding node.

Configuring Layer 2 IP Ping

Use this procedure to configure Layer 2 IP ping

Before you begin

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP, and assign a nodal MEP to the B-VLAN.
- If you want to run a Layer 2 IP Ping for a specific VRF, you must use EDM in the specific VRF context first. For more information, see the procedure for selecting and launching a VRF context view in Configuring IPv4 Routing for VOSS.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **Diagnostics**.
2. Click **L2Ping/L2Trace Route**.
3. Click the **L2 IP Ping** tab.
4. To add a new entry, click **Insert**, specify the destination IP address and optional parameters, and then click **Insert**.
5. To start the Layer 2 IP ping, highlight an entry, and then click **Start**.
6. To update the Layer 2 IP ping, click the **Refresh** button.
7. To stop the Layer 2 IP ping, click **Stop**.

L2 IP Ping Field Descriptions

Use the data in the following table to use the **L2 IP Ping** tab.

Name	Description
IpAddrType	Specifies the address type of destination IP Address (only IPv4 is supported).
IpAddr	Specifies the destination IP Address.
VrfId	Specifies the VRF ID.
VrfName	Specifies the name of the virtual router.
Messages	Specifies the number of L2IpPing messages to be transmitted per MAC/VLAN pair. Range is 1–200. The default is 1.
Status	Specifies the status of the transmit loopback service: • ready: the service is available. • transmit: the service is transmitting, or about to transmit, the L2IpPing messages. • abort: the service is aborted or about to abort the L2IpPing messages. This field is also used to avoid concurrency or race condition problems that could occur if two or more management entities try to use the service at the same time. The default is ready.
ResultOk	Indicates the result of the operation: • true: L2IpPing Messages will be or have been sent. • false: L2IpPing Messages will not be sent. The default is true.
TimeoutInt	Specifies the interval to wait for an L2IpPing time-out with a range of 1–10 seconds with a default value of 3 seconds.
TestPattern	Specifies the test pattern to use in the L2IPPing PDU: • **allZero** — null signal without cyclic redundancy check • **allZeroCrc** — null signal with cyclic redundancy check with 32-bit polynomial • **pseudoRandomBitSequence** — pseudo-random-bit-sequence without cyclic redundancy check • **pseudoRandomBitSequenceCrc** — pseudo-random-bit-sequence with cyclic redundancy check with 32-bit polynomial. A cyclic redundancy check is a code that detects errors.

Table continues…

Name	Description
	The default value is allZero.
DataSize	Specifies an arbitrary amount of data to be included in the data TLV, if the data size is selected to be sent. The range is 0–400. The default is 0.
PathsFound	Specifies the number of paths found to execute the command. The default is 0.

Viewing Layer 2 IP Ping Results

Use this procedure to view Layer 2 IP ping results.

Note:

After you trigger Layer 2 IP Ping, you must click the **Refresh** button to update the results.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **Diagnostics**.
2. Click **L2Ping/L2Trace Route**.
3. Click the **L2 IP Ping** tab.
4. To view the Layer 2 IP ping results, highlight an entry, and then click **Result**.

L2 IP Ping Result Field Descriptions

Use the data in the following table to use the **L2 IP Ping Result** tab.

Name	Description
IpAddrType	The address type of the destination IP Address.
IpAddr	Destination IP Address.
SendOrder	Specifies the order that sessions were sent. It is an index to distinguish among multiple L2Ping sessions. This value is assigned sequentially from 1. It correlates to the number of paths found.
VrfId	Specifies the VRF ID.
VlanId	Specifies the VLAN ID found from the Layer 3 lookup and used for transmission.
DestMacAddress	An indication of the target MAC Address transmitted.
PortNum	Either the value '0', or the port number of the port used for the l2 IP ping.
DestHostName	The host name of the responding node.
Size	The number of bytes of data sent.

Table continues…

Name	Description
PktsTx	Number of Packets transmitted for this VLAN/MAC.
PktsRx	Number of Packets received for this VLAN/MAC.
PercentLossWhole	Percentage of packet loss for this VLAN/MAC.
PercentLossFract	Percentage of packet loss for this VLAN/MAC.
MinRoundTrip	Minimum time for round-trip for this VLAN/MAC in us.
MaxRoundTrip	Maximum time for round-trip for this VLAN/MAC in us.
RttAvgWhole	Average time for round-trip for this VLAN/MAC in us.
RttAvgFract	Fractional portion of average time for round-trip.
Flag	Result flag indicating status or error code: • 1 - No error • 2 - Internal error • 3 - Invalid IP • 4 - L2Trace completed • 5 - Lookup failure for IP (no VLAN/MAC entries)

Configuring Layer 2 IP Traceroute

Use this procedure to configure Layer 2 IP traceroute.

Before you begin

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP, and assign a nodal MEP to the B-VLAN
- If you want to run a Layer 2 IP Traceroute for a specific VRF, you must use EDM in the specific VRF context first. For more information, see the procedure for selecting and launching a VRF context view in Configuring IPv4 Routing for VOSS.

Procedure

1. In the navigation pane, expand the **Configuration** > **Edit** > **Diagnostics** folders.
2. Click **L2Ping/L2Trace Route**
3. Click the **L2 IP Traceroute** tab.
4. To add a new entry, click **Insert**, specify the destination IP address and, optionally, the TTL value, and then click **Insert**.
5. To start the Layer 2 IP traceroute, highlight an entry, and then click the **Start** button.
6. To update the L2 IP traceroute, click the **Refresh** button.
7. To stop the Layer 2 IP traceroute, click the **Stop** button.

L2 IP Traceroute Field Descriptions

Use the data in the following table to use the **L2 IP Traceroute** tab.

Name	Description
IpAddrType	Specifies the address type of destination IP address (only IPv4 is supported).
IPAddr	Specifies the destination IP Address.
VrfId	Specifies the VRF ID.
VrfName	Specifies the name of the virtual router.
Ttl	Specifies the number of hops remaining to this L2Trace. This value is decremented by 1 by each Bridge that handles the L2Trace. The decremented value is returned in the L2Trace. If 0 on output, the L2Trace is not transmitted to the next hop. The default value is 64
Status	Indicates the status of the transmit loopback service: • ready: the service is available. • transmit: the service is transmitting, or about to transmit, the L2Trace messages. • abort: the service is aborted or about to abort the L2Trace messages. This field is also used to avoid concurrency or race condition problems that could occur if two or more management entities try to use the service at the same time. The default is ready.
ResultOk	Indicates the result of the operation: • true: the Trace Messages will be or have been sent. • false. the Trace Messages will not be sent The default is true.
PathsFound	Specifies the number of paths found to execute the L2trace. The default is 0.

Viewing Layer 2 IP Traceroute Results

Use this procedure to view Layer 2 IP traceroute results.

Note:

After you trigger Layer 2 IP traceroute, you must click the **Refresh** button to update the results.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **Diagnostics**.
2. Click **L2Ping/L2Trace Route**.
3. Click the **L2 IP Traceroute** tab.
4. To view the Layer 2 IP traceroute results, highlight an entry, and then click **Result**.

L2 IP Traceroute Result Field Descriptions

Use the data in the following table to use the **L2 IP Traceoute Result** tab.

Name	Description
IpAddrType	Specifies the address type of destination IP address.
IpAddr	Specifies the destination IP address.
SendOrder	Denotes the order that sessions are sent. It is an index to distinguish among multiple L2Trace sessions. It correlates to the number of paths found. This value is assigned sequentially from 1.
Hop	Specifies the number of L2 hops away from L2Trace initiator.
ReceiveOrder	Specifies the order that sessions are sent. It is an index to distinguish among multiple L2Trace responses with the same Send Transaction Identifier field value. This value is assigned sequentially from 1, in the order that the Linktrace Initiator received the responses.
Ttl	Specifies the time-to-live (TTL) field value for a returned L2Trace response.
VrfId	Specifies the VRF ID.
VlanId	Specifies the VLAN found from Layer 3 lookup and used for transmission.
DestMacAddress	Indicates the target MAC address transmitted.
PortNum	Specifies either the value '0', or the port number of the port used for the l2trace.
SeqNumber	Specifies the transaction identifier/sequence number used in linktrace message packet. The default is 0.
SrcMac	Specifies the MAC address of the MP that responded to L2Trace request for this L2traceReply.
HostName	Specifies the host name of the replying node.
LastSrcMac	Specifies the MAC address of the node that forwarded the L2Trace to the responding node.
LastHostName	Specifies the host name of the node that forwarded the L2Trace to the responding node.

Table continues…

Name	Description
Flag	L2Trace result flag indicating status or error code: • none (1): No error • internalError (2): L2Trace internal error • invalidMac (3): Invalid MAC address • mepDisabled (4): MEP must be enabled in order to perform L2Trace • noL2TraceResponse (5): No L2Trace response received • l2TraceToOwnMepMac (6): L2Trace to own MEP MAC is not sent • l2TraceComplete (7): L2Trace completed • l2TraceLookupFailure (8): Lookup failure for L2Trace

Triggering a Loopback Test

Use this procedure to trigger a loopback test.

The LBM packet is often compared to ping. An MEP transmits the loopback message to an intermediate or endpoint within a domain for the purpose of fault verification. This can be used to check the ability of the network to forward different sized frames.

Before you begin

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP.
- Enable the MEP.
- Assign a nodal MEP to the B-VLAN.

Procedure

1. In the navigation pane, expand the **Configuration** > **Edit** > **Diagnostics** folders.
2. Click **CFM**.
3. Click the **LBM** tab.
4. Configure the loopback test properties as required.
5. Click **Apply**.
6. To trigger the loopback test, double-click in the **Status** field, select **transmit**.
7. Click **Apply**.
8. To update the loopback test, click the **Refresh** button.

LBM Field Descriptions

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

Name	Description
DomainIndex	Specifies the MD index value.
AssociationIndex	Specifies the MA index value.
Index	Specifies the Maintenance Endpoint index value.
DomainName	Specifies the MD name.
AssociationName	Specifies the MA name.
DestMacAddress	Specifies the remote MAC address to reach the MEP/MIP.
Messages	Specifies the number of loopback messages to be transmitted. The default is 1.
VlanPriority	Specifies the priority. The default is 7.
SeqNumber	Specifies the transaction identifier/sequence number of the first loopback message (to be) sent. The default is 0.
ResultOk	Indicates the result of the operation: • true: The Loopback Messages will be (or have been) sent. • false: The Loopback Messages will not be sent. The default is true.
Status	Indicates the status of the transmit loopback service: • ready: The service is available. • transmit: The service is transmitting, or about to transmit, the Loopback messages. • abort: The service is aborted or about to abort the Loopback messages. The default is ready.
Result	Displays the LBM result.
TimeoutInt	Specifies the timeout interval in seconds. The default value is 3 seconds.
InterFrameInt	Specifies the interval between LBM frames with a range of (0..1000) msecs and a default value of 500 msecs. The value of 0 msecs indicates to send the frames as fast as possible. The default is 500.
TestPattern	Specifies the testfill pattern: • **allZero** — null signal without cyclic redundancy check

Table continues…

Name	Description
	• **allZeroCrc** — null signal with cyclic redundancy check with 32-bit polynomial • **pseudoRandomBitSequence** — pseudo-random-bit-sequence without cyclic redundancy check • **pseudoRandomBitSequenceCrc** — pseudo-random-bit-sequence with cyclic redundancy check with 32-bit polynomial. A cyclic redundancy check is a code that detects errors. The default value is allZero.
DataSize	Specifies the data type-length-value (TLV) size. The default is 0.
FrameSize	Specifies the frame-size. The default is 0.
Sourcemode	Specifies the source mode of the transmit loopback service: • nodal • noVlanMac — Use this value with C-VLAN only. When you select this option, even if a VLAN MAC address exists, the system uses the CFM C-MAC as the BMAC-SA. • smltVirtual — Use the smltVirtual option with B-VLANs only. The default is nodal.

Triggering Linktrace

Use the following procedure to trigger a linktrace. The link trace message is often compared to traceroute. An MEP transmits the Linktrace Message packet to a maintenance endpoint with intermediate points responding to indicate the path of the traffic within a domain for the purpose of fault isolation. The packet specifies the target MAC address of an MP, which is the SPBM system ID or the virtual SMLT MAC. MPs on the path to the target address respond with an LTR.

Before you begin

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP.
- Enable the MEP.
- Assign a nodal MEP to the B-VLAN.

Procedure

1. In the navigation pane, expand the **Configuration** > **Edit** > **Diagnostics** folders.
2. Click **CFM**.
3. Click the **LTM** tab.

4. Configure the linktrace test properties as required.
5. Click **Apply**.
6. To trigger the linktrace test, double-click in the Status field, select **transmit**, and then click **Apply**.

   OR

   Highlight an entry, and then click **Start**.
7. To update the linktrace, click the **Refresh** button.
8. To stop the linktrace, click **Stop**.
9. To view the results of the linktrace, click **Result**.

LTM Field Descriptions

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

Name	Description
DomainIndex	Specifies the MD index value.
AssociationIndex	Specifies the MA index value.
Index	Specifies the MEP index value.
DomainName	Specifies the MD name.
AssociationName	Specifies the MA name.
VlanPriority	Specifies the VLAN priority, a 3–bit value to be used in the VLAN tag, if present in the transmitted frame. The default is 7.
DestMacAddress	Specifies the remote MAC address to reach the MEP.
Ttl	Indicates the number of hops remaining to this LTM. This value is decremented by 1 by each bridge that handles the LTM. The decremented value is returned in the LTR. If the value is 0 on output, the LTM is not transmitted to the next hop. The value of the TTL field in the LTM is specified at the originating MEP. The default value is 64.
SeqNumber	Specifies the transaction identifier/sequence number of the first loopback message (to be) sent. The default is 0.
ResultOk	Indicates the result of the operation: • true: The Loopback Messages will be (or have been) sent. • false: The Loopback Messages will not be sent. The default is true.

Table continues…

Name	Description
Status	Indicates the status of the transmit loopback service: • ready: The service is available. • transmit: The service is transmitting, or about to transmit, the LTM messages. • abort: The service is aborted, or about to abort, the LTM message. The default is ready.
Flag	Displays the LTM result flag indicating LTM status or error code. Each value represents a status or error case: • 1 - No error • 2 - LTM internal error • 3 - Unknown Remote Maintenance Endpoint • 4 - Invalid Remote Maintenance Endpoint MAC Address • 5 - Unset Remote Maintenance Endpoint MAC address • 6 - MEP must be enabled in order to perform LTM • 7 - No LTR response received • 8 - Linktrace to own MEP MAC is not sent • 9 - Endpoint must be enabled in order to perform LTM • 10 - Pbt-trunk must be enabled in order to perform LTM • 11 - LTM completed • 12 - LTM leaf node
SourceMode	Specifies the source mode of the transmit loopback service: • nodal • noVlanMac — Use this value with C-VLAN only. When you select this option, even if a VLAN MAC address exists, the system uses the CFM C-MAC as the BMAC-SA. • smltVirtual — Use the smltVirtual option with B-VLANs only. The default is nodal.

Viewing Linktrace Results

Use this procedure to view linktrace results.

Note:

After you trigger linktrace, you must click the **Refresh** button to update the results.

Procedure

1. In the navigation pane, expand the **Configuration** > **Edit** > **Diagnostics** folders.
2. Click **CFM**.
3. Click the **LTM** tab.
4. Highlight an entry, and then click **Result**.

Link Trace Replies Field Descriptions

Use the data in the following table to use the **Link Trace Result** tab.

Name	Description
DomainIndex	Indicates the Maintenance Domain Index.
AssociationIndex	Indicates the Maintenance Association Index.
MepId	Indicates the Maintenance EndPoint ID.
SeqNumber	Indicates the transaction identifier/sequence number returned by a previous transmit linktrace message command, indicating which LTM response is going to be returned. The default is 0.
Hop	Indicates the number of hops away from the LTM initiator.
ReceiveOrder	Indicates the index value used to distinguish among multiple LTRs with the same LTR Transaction Identifier field value. This value is assigned sequentially from 1, in the order that the Linktrace Initiator received the LTRs.
Ttl	Indicates the TTL field value for a returned LTR.
DomainName	Indicates the Maintenance Domain Name.
AssociationName	Indicates the Maintenance Association Name.
Forwarded	Indicates if a LTM was forwarded by the responding MP, as returned in the FwdYes flag of the flags field.
TerminalMep	Displays a boolean value stating whether the forwarded LTM reached a MEP enclosing its MA, as returned in the Terminal MEP flag of the Flags field.

Table continues…

Name	Description
LastEgressIdentifier	Displays an octet field holding the Last Egress Identifier returned in the LTR Egress Identifier TLV of the LTR. The Last Egress Identifier identifies the MEP Linktrace Indicator that originated, or the Linktrace Responder that forwarded, the LTM to which this LTR is the response. This is the same value as the Egress Identifier TLV of that LTM.
NextEgressIdentifier	Displays an octet field holding the Next Egress Identifier returned in the LTR Egress Identifier TLV of the LTR. The Next Egress Identifier Identifies the Linktrace Responder that transmitted this LTR, and can forward the LTM to the next hop. This is the same value as the Egress Identifier TLV of the forwarded LTM, if any. If the FwdYes bit of the Flags field is false, the contents of this field are undefined, and the field is ignored by the receiver.
RelayAction	Indicates the value returned in the RelayAction field.
SrcMac	Displays the MAC address of the MP that responded to the LTM request for this LTR.
IngressAction	Displays the value returned in the IngressAction Field of the LTM. The value ingNoTlv indicates that no Reply Ingress TLV was returned in the LTM.
IngressMac	Displays the MAC address returned in the ingress MAC address field. If the rcCfmLtrReplyIngress object contains the value ingNoTlv(5), then the contents of this field are meaningless.
EgressAction	Displays the value returned in the Egress Action Field of the LTM. The value egrNoTlv(5) indicates that no Reply Egress TLV was returned in the LTM.
EgressMac	Displays the MAC address returned in the egress MAC address field. If the rcCfmLtrReplyEgress object contains the value egrNoTlv(5), then the contents of this field are meaningless.

Configure a Layer 2 Tracetree

Note:

This procedure only applies to VSP 4000 Series, VSP 4900 Series, VSP 7200 Series, VSP 7400 Series, VSP 8200 Series, VSP 8400 Series, and VSP 8600 Series.

Use this procedure to configure a Layer 2 Tracetree. This feature enables CFM to debug Layer 2. Layer 2 Tracetree allows a user to trigger a multicast LTM message by specifying the B-VLAN and I-SID. The command allows the user to view a multicast tree on the SPBM B-VLAN from the source node to the destination nodes for a particular I-SID.

Note:

Troubleshooting using ping and traceroute (including Layer 2 ping and Layer 2 traceroute) is not supported on EDM. For more information, see Release Notes for VOSS. As an alternative, use CLI.

If you configure **IsTraceTree** to false, then EDM performs Traceroute on the unicast path. If you configure **IsTraceTree** to true, then EDM performs TraceTree on the multicast tree.

Note:

This command is supported on SPBM B-VLANs only, not C-VLANs.

Before you begin

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP.
- Enable the MEP.
- Assign a nodal MEP to the B-VLAN.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **Diagnostics**.
2. Select **L2Ping/L2Trace Route**.
3. From the **L2 Traceroute/TraceTree** tab, configure the Layer 2 tracetree properties.
4. In the **IsTraceTree** field double-click and select **true** for EDM to perform Tracetree on the multicast tree.
5. Select **Apply**.
6. Select **Refresh** to update the results.

L2Tracetree Field Descriptions

Use the data in the following table to use the **L2Tracetree** tab.

Name	Description
VlanId	Identifies the Backbone VLAN.
Priority	Specifies a 3–bit value to be used in the VLAN header, if present in the transmitted frame. The default is 7.
DestMacAddress	Specifies the target MAC address.
HostName	Specifies the target host name.
DestIsHostName	Indicates whether the host name is (true) or is not (false) used for L2Tracetree transmission.
Isid	Specifies the service instance identifier (I-SID).
IsTraceTree	Specifies whether the multicast tree or unicast path is traced. If you configure **IsTraceTree** to false then EDM performs Traceroute on the unicast path. If you

Table continues…

Name	Description
	configure **IsTraceTree** to true then EDM performs TraceTree on the multicast tree.
Status	Specifies the status of the transmit loopback service: • ready: the service is available. • transmit: the service is transmitting, or about to transmit, the L2Tracetree messages. • abort: the service aborted or is about to abort the L2Tracetree messages. This field is also used to avoid concurrency or race condition problems that can occur if two or more management entities try to use the service at the same time. The default is ready.
ResultOk	Indicates the result of the operation: • true: the L2Tracetree Messages will be (or have been) sent. • false: the L2Tracetree Messages will not be sent The default is true.
Ttl	Specifies the Time-to-Live value. Indicates the number of hops remaining to this L2Tracetree. The tracetree is decremented by one by each bridge that handles the Layer 2 tracetree and the decremented value is returned to the tracetree. If the output is 0, then the L2Tracetree is not transmitted to the next hop. The value of the TTL field in the L2Tracetree is transmitted by the originating MEP is controlled by a managed object. The default is 64.
SourceMode	Specifies the source mode of the transmit loopback service: • nodal • noVlanMac — Use this value with C-VLAN only. When you select this option, even if a VLAN MAC address exists, the system uses the CFM C-MAC as the BMAC-SA. • smltVirtual — Use the smltVirtual option with B-VLANs only. The default is nodal.
SeqNumber	The transaction identifier/sequence number of the first loopback message (to be) sent. The default is 0.

Table continues…

Name	Description
Flag	Specifies the L2Tracetree result flag, which indicates the L2Tracetree status or error code. Each sum represents a status or error: • 1 — No error • 2 — L2Tracetree internal error • 3 — Invalid MAC address • 4 — MEP must be enabled in order to perform L2Tracetree • 5 — No L2Tracetree response received • 6 — L2Tracetree to own MEP MAC is not sent • 7 — L2Tracetree completed • 8 — Lookup failure for L2Tracetree • 9 — On a leaf node in the I-SID tree • 10 — Not in the I-SID tree • 11 — Requested SMLT source from nonprimary node

View Layer 2 Tracetree Results

Note:

This procedure only applies to VSP 4000 Series, VSP 4900 Series, VSP 7200 Series, VSP 7400 Series, VSP 8200 Series, VSP 8400 Series, and VSP 8600 Series.

Use this procedure to view Layer 2 Tracetree results. The Layer 2 Tracetree command is a proprietary command that allows a user to trigger a multicast LTM message by specifying the B-VLAN and I-SID. This command allows the user to view a multicast tree on the SPBM B-VLAN from the source node to the destination nodes for a particular I-SID.

Note:

Troubleshooting using ping and traceroute (including Layer 2 ping and Layer 2 traceroute) is not supported on EDM. For more information, see Release Notes for VOSS. As an alternative, use CLI.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **Diagnostics**.
2. Select **L2Ping/L2Trace Route**.
3. Select the **L2 Traceroute/TraceTree** tab.

4. In the **IsTraceTree** field double-click and select **true** for EDM to perform Tracetree on the multicast tree.
5. Select **Apply**.
6. Select **Refresh** to update the results.
7. To view the tracetree results, highlight an entry, and then select **Result**.

L2 Traceroute/Tracetree Result Field Descriptions

Use the data in the following table to use the **L2 Traceroute/Tracetree Result** tab.

Name	Description
VlanId	A value that uniquely identifies the Backbone VLAN (B-VLAN).
SeqNumber	The transaction identifier/sequence number returned by a previous transmit linktrace message command, that indicates which response of the L2Tracetree is going to be returned. The default is 0.
Hop	The number of hops away from L2Tracetree initiator.
ReceiveOrder	An index to distinguish among multiple L2Tracetree responses with the same Transaction Identifier field value. This value is assigned sequentially from 1, in the order that the Linktrace Initiator received the responses.
Ttl	Time-to-Live (TTL) field value for a returned L2Tracetree response.
SrcMac	MAC address of the MP that responds to the L2Tracetree request for this L2tractreeReply.
HostName	The host name of the replying node.
LastSrcMac	The MAC address of the node that forwarded the L2Tracetree to the responding node.
LastHostName	The host name of the node that forwarded the L2Tracetree to the responding node.

Configure Layer 2 Trace Multicast Route on a VLAN

Note:

This procedure only applies to VSP 4000 Series, VSP 4900 Series, VSP 7200 Series, VSP 7400 Series, VSP 8200 Series, VSP 8400 Series, and VSP 8600 Series.

Use this procedure to configure the Layer 2 tracemroute on the VLAN (Layer 2). This procedure queries the SPBM multicast module to determine the B-VLAN, I-SID, and nickname for the S and G streams. The nickname and I-SID are used to create a multicast MAC address.

Note:

If you want to run a Layer 2 tracemroute on a VRF, make sure you are in the proper VRF context.

Before you begin

On the source and destination nodes, you must configure a CFM MD, MA, and MEP, and assign a nodal MEP to the B-VLAN.

Procedure

1. In the navigation pane, expand the **Configuration** > **Edit** > **Diagnostics** > **L2Ping/L2Trace Route** folders.
2. Select the **L2MCAST Traceroute** tab.
3. Select **Insert** to insert the L2 MCAST Traceroute.
4. Enter the **SrcIpAddr**.
5. Enter the **GroupIpAddr**.
6. Enter the **ServiceType**. If you want to perform a Layer 2 tracemroute on a VLAN, select **vlan**. If you want to perform a Layer 2 tracemroute on a Layer 3 GRT, select **vrfid**.

   Note:

   If you want to perform a Layer 2 tracemroute on a Layer 2 or a Layer 3 VRF, review the following procedure Configuring Layer 2 tracemroute on a VRF on page 143.
7. In the **ServiceId**field, enter the VLAN ID.
8. Enter the **Priority**.
9. Enter the **Ttl** value.
10. Select **Insert**.
11. Select **Apply** to save your changes.
12. To start the Layer 2 tracemoute, set the Status to transmit and select **Start**.
13. Update the Layer 2 tracemroute by selecting **Refresh** .
14. To stop the Layer 2 tracemroute, select **Stop** .
15. To see the result, select **Result**.

L2 MCAST Traceroute Field Descriptions

Use the data in the following table to use the L2 MCAST Traceroute tab.

Name	Description
SrcIpAddrType	Specifies the source IP address type as IPv4.
SrcIpAddr	Specifies the source IP address of the flow where the multicast trace tree originates.

Table continues…

Name	Description
GroupIpAddrType	Specifies the group IP address type as IPv4.
GroupIpAddr	Specifies the group IP address.
ServiceType	Specifies where you configure the Layer 2 tracemroute. This is either VLAN or VRF.
ServiceId	Specifies the VLAN ID.
VRFName	Specifies the VRF name.
Priority	Specifies the priority value. The value is between 0 and 7.
Ttl	Specifies the returned trace response. The TTL value is between 1 and 255.
SeqNumber	Specifies the transaction identifier/sequence number of the first message to be sent.
Status	Specifies the status of the transmit loopback service: • ready: Specifies the service is available. • transmit: Specifies the service is transmitting, or about to transmit the trace messages. • abort: Specifies the services is aborted or about to abort the trace messages. The column will also be used to avoid concurrency or race condition problems that can occur if two or more management entities try to use the service at the same time.
ResultOK	Specifies the result of the operation: • true: The trace messages will be or have been sent. • false: The trace messages will not be sent.
Flag	Specifies the result flag indicating that the L2 trace status or error code. Each value represents a status or error case. • 1 — No error • 2 — Internal Error • 3 — Mep must be enabled to perform the trace • 4 — No response received • 5 — Trace completed • 6 — On a leaf node in the I-SID tree • 7 — No data I-SID was found for S, G

Configure Layer 2 Tracemroute on a VRF

Note:

This procedure only applies to VSP 4000 Series, VSP 4900 Series, VSP 7200 Series, VSP 7400 Series, VSP 8200 Series, VSP 8400 Series, and VSP 8600 Series.

Use this procedure to configure the Layer 2 tracemroute on the VRF (Layer 3). This procedure queries the SPBM multicast module to determine the B-VLAN, I-SID and nickname for the S and G streams. The nickname and I-SID are used to create a multicast MAC address.

Note:

If you want to run a Layer 2 tracemroute on a VRF, make sure you are in the proper VRF context.

See the following procedure to perform a Layer 3 tracemroute on a VLAN Configuring Layer 2 tracemroute on a VLAN on page 140.

Before you begin

On the source and destination nodes, you must configure a CFM MD, MA, and MEP, and assign a nodal MEP to the B-VLAN.

Procedure

1. In the navigation pane, expand the **Configuration** > **VRF Context View** > **Set VRF Context View** folders.
2. Select a VRF and click the **Launch VRF Context View** tab.
3. In the navigation pane, expand the following folders:**Configuration** > **Edit** > **Diagnostics** > **L2Ping/L2Trace Route**.
4. Select the **L2MCAST Traceroute** tab.
5. Select **Insert** to insert the L2 MCAST traceroute.
6. Enter the **SrcIpAddr**.
7. Enter the **GroupIpAddr**.
8. Enter the **ServiceType**. If you want to perform a Layer 2 tracemroute on a Layer 2 VRF, select **vlan**. If you want to perform a Layer 2 tracemroute on a Layer 3 VRF, select **vrfid**.
9. In the **ServiceId**, enter the VLAN ID.
10. Enter the **Priority**.
11. Enter the **Ttl** value.
12. Select **Insert**.
13. Select **Apply** to save your changes.
14. To start the Layer 2 tracemoute, set the Status to transmit and select **Start**.

15. Update the Layer 2 tracemroute by select **Refresh** .
16. To stop the Layer 2 tracemroute, select **Stop** .
17. To see the result, select **Result**.

L2 MCAST Traceroute Field Descriptions

Use the data in the following table to use the L2 MCAST Traceroute tab.

Name	Description
SrcIpAddrType	Specifies the source IP address type as IPv4.
SrcIpAddr	Specifies the source IP address of the flow where the multicast trace tree originates.
GroupIpAddrType	Specifies the group IP address type as IPv4.
GroupIpAddr	Specifies the group IP address.
ServiceType	Specifies where you configure the Layer 2 tracemroute. This is either VLAN or VRF.
ServiceId	Specifies the VLAN ID.
VRFName	Specifies the VRF name.
Priority	Specifies the priority value. The value is between 0 and 7.
Ttl	Specifies the returned trace response. The TTL value is between 1 and 255.
SeqNumber	Specifies the transaction identifier/sequence number of the first message to be sent.
Status	Specifies the status of the transmit loopback service: • ready: Specifies the service is available. • transmit: Specifies the service is transmitting, or about to transmit the trace messages. • abort: Specifies the services is aborted or about to abort the trace messages. The column will also be used to avoid concurrency or race condition problems that can occur if two or more management entities try to use the service at the same time.
ResultOK	Specifies the result of the operation: • true: The trace messages will be or have been sent. • false: The trace messages will not be sent.
Flag	Specifies the result flag indicating that the L2 trace status or error code. Each value represents a status or error case. • 1 — No error • 2 — Internal Error • 3 — Mep must be enabled to perform the trace

Table continues…

Name	Description
	• 4 — No response received • 5 — Trace completed • 6 — On a leaf node in the I-SID tree • 7 — No data I-SID was found for S, G

Viewing Layer 2 Trace Multicast Route Results

Note:

This procedure only applies to VSP 4000 Series, VSP 4900 Series, VSP 7200 Series, VSP 7400 Series, VSP 8200 Series, VSP 8400 Series, and VSP 8600 Series.

Use this procedure to view Layer 2 tracemroute results.

Note:

If you want to run a Layer 2 tracemroute on a VRF, make sure you are in the proper VRF context.

Procedure

1. In the navigation pane, expand the **Configuration** > **Edit** > **Diagnostics** > **L2Ping/L2Trace Route** folders.
2. Select the **L2 MCAST Traceroute** tab.
3. To view the CFMl2 trace multicast route results, highlight an entry and select **Result**.

L2tracemroute Result Field Descriptions

Use the data in the following table to use the **L2tracemroute Result** tab.

Name	Description
VlanId	Specifies a value that uniquely identifies the C-VLAN.
SeqNumber	Specifies the transaction identifier/sequence number returned by a previous transmit linktrace message command. Indicates which l2 tracemroute response is going to be returned.
Hop	Specifies the number of hops away from the l2 tracemroute initiator.
ReceiveOrder	Specifies an index to distinguish among multiple l2 tracemroute responses with the same transaction identifier field value. This value is assigned sequentially from 1, in the order that the linktrace initiator received the responses.
Ttl	Specifies the TTL value for a returned l2 tracemroute response.

Table continues…

Name	Description
SrcMac	Specifies the MAC address of the MP that responds to the l2 tracemroute request for this l2 tracemrouteReply.
HostName	Specifies the host name of the replying node.
LastSrcMac	Specifies the MAC address of the node that forwarded the l2 tracemroute to the responding node.
LastHostName	Specifies the host name of the node that forwarded the l2 tracemroute to the responding node.

CFM Configuration Example

This section provides a configuration example for Connectivity Fault Management (CFM).

CFM Configuration Example

The following sections show the steps required to configure CFM.

Note:

The following commands are not supported on all hardware platforms:

- cfm maintenance-domain
- cfm maintenance-association
- vlan nodal-mip-level

Switch A

```
MAINTENANCE-DOMAIN CONFIGURATION

cfm maintenance-domain "spbm" index 1 maintenance-level 6

MAINTENANCE-ASSOCIATION CONFIGURATION

cfm maintenance-association "spbm" "2" index 1
cfm maintenance-association "spbm" "3" index 2

MAINTENANCE-ENDPOINT CONFIGURATION

cfm maintenance-endpoint "spbm" "2" 1 state enable
cfm maintenance-endpoint "spbm" "3" 1 state enable

VLAN NODAL MEP/MIP CONFIGURATION

vlan nodal-mep 2 spbm 2 1
vlan nodal-mip-level 2 6
vlan nodal-mep 3 spbm 3 1
vlan nodal-mip-level 3 6
```

Switch B

```
MAINTENANCE-DOMAIN CONFIGURATION

cfm maintenance-domain spbm index 1 maintenance-level 6

MAINTENANCE-ASSOCIATION CONFIGURATION

cfm maintenance-association "spbm" "2" index 1
cfm maintenance-association "spbm" "3" index 2

MAINTENANCE-ENDPOINT CONFIGURATION

cfm maintenance-endpoint "spbm" "2" 2 state enable
cfm maintenance-endpoint "spbm" "3" 2 state enable

VLAN NODAL MEP/MIP CONFIGURATION

vlan nodal-mep 2 spbm 2 2
vlan nodal-mip-level 2 6
vlan nodal-mep 3 spbm 3 2
vlan nodal-mip-level 3 6
```

CFM Sample Output

The following sections show sample CFM output.

L2ping can use the system ID or the router name. The example below shows a case where the VLAN and MAC are given.

show isis adjacencies

```
Switch:1# show isis adjacencies
================================================================================
                                ISIS Adjacencies
================================================================================
INTERFACE     IP ADDR            L STATE        UPTIME         PRI     HOLDTIME     SYSID
--------------------------------------------------------------------------------
Port1/3      192.0.2.33        1 UP         00:37:37         127   19
0014.0dbf.a3df
Port1/19     192.0.2.36        1 UP      1d 05:09:16         127   21
0014.0da2.b3df
--------------------------------------------------------------------------------
2 out of 2 interfaces have formed an adjacency
--------------------------------------------------------------------------------
```

l2 ping with vlan

```
Switch:1# l2 ping vlan 500 mac 00.14.0d.bf.a3.df


Please wait for l2ping to complete or press any key to abort

----00:14:0d:bf:a3:df    L2 PING Statistics----  0(68) bytes of data
1 packets transmitted, 0 packets received,   100.00% packet loss
```

l2 ping with vlan

```
Switch:1# l2 ping vlan 500 routernodename MONTI0


Please wait for l2ping to complete or press any key to abort
```

```
----00:14:0d:a2:b3:df    L2 PING Statistics----  0(68) bytes of data
1 packets transmitted, 1 packets received,  0.00% packet loss
  round-trip (us)        min/max/ave/stdv =  26895/26895/26895.00/ 0.00
```

l2 traceroute with vlan

```
Switch:1# l2 traceroute vlan 500 routernodename MONTI0


Please wait for l2traceroute to complete or press any key to abort

l2traceroute to MONTI0  (00:14:0d:a2:b3:df),  vlan 500
0    PETER4         (00:15:9b:11:33:df)
1    MONTI0          (00:14:0d:a2:b3:df)
```

l2 tracetree with vlan

```
Switch:1# l2 tracetree 500 1

Please wait for l2tracetree to complete or press any key to abort

l2tracetree to 53:55:10:00:00:01, vlan 500 i-sid 1 nickname 5.55.10 hops 64
1   PETER4        00:15:9b:11:33:df -> MONTI0        00:14:0d:a2:b3:df
2   MONTI0        00:14:0d:a2:b3:df -> LEE2          00:15:e8:b8:a3:df
```

L2ping and L2traceroute can also be used with an IP address. The following outputs show examples using an IP address.

l2 ping with IP address

```
Switch:1# l2 ping ip-address 192.0.2.10


Please wait for l2ping to complete or press any key to abort

L2 PING  Statistics : IP 192.0.2.10, paths found 1, paths attempted 1
================================================================================
TX    RX    PERCENT  ROUND TRIP TIME
VLAN NEXT HOP                              PKTS  PKTS  LOSS     MIN/MAX/AVE (us)
================================================================================
500  SHAMIM        (00:1a:8f:08:53:df)  1     0     100.00%  0/0/0.00
```

l2 ping with IPv6 address

```
Switch:1# l2 ping ip-address 49:0:0:0:0:0:0:11

Please wait for l2ping to complete or press any key to abort

L2 PING  Statistics : IP 49:0:0:0:0:0:0:11, paths found 1, paths attempted 1
================================================================================
                       TX    RX    PERCENT  ROUND TRIP TIME
VLAN NEXT HOP                              PKTS  PKTS  LOSS     MIN/MAX/AVE (us)
================================================================================

41   SHAMIM            (00:49:00:01:00:11)  1     1       0.00%  11876/11876/11876.00
```

l2 traceroute with IP address

```
Switch:1# l2 traceroute ip-address 192.0.2.10


Please wait for l2trace to complete or press any key to abort
```

```
L2 Trace  Statistics : IP 192.0.2.10, paths found 1
================================================================
SHAMIM  (00:1a:8f:08:53:df),  vlan 500
0    PETER4          (00:15:9b:11:33:df)
1    MONTI0          (00:14:0d:a2:b3:df)
```

l2 traceroute with IPv6 address

```
Switch:1# l2 traceroute ip-address 49:0:0:0:0:0:0:11

Please wait for l2trace to complete or press any key to abort

L2 Trace  Statistics : IP 49:0:0:0:0:0:0:11, paths found 1

================================================================
SHAMIM  (00:49:00:01:00:11),  vlan 41
0    4K-DUT7             (00:49:00:01:00:17)
1    9k-2              (00:49:00:01:00:92)
2    8K-1              (00:49:00:08:00:81)
3    4K-DUT1             (00:49:00:01:00:11)
```

show cfm maintenance-domain

Note:

The following commands are not supported on all hardware platforms:

- cfm maintenance-domain
- cfm maintenance-association
- vlan nodal-mip-level

```
Switch:1#show cfm maintenance-domain
================================================================================
                          Maintenance Domain
================================================================================
Domain Name              Domain Index    Level Domain Type
--------------------------------------------------------------------------------
md1                      99              3     NONE

 Total number of Maintenance Domain entries: 1.
```

show cfm maintenance-association

```
Switch:1#show cfm maintenance-association
================================================================================
                          Maintenance Association Status

================================================================================
Domain Name            Assn Name              Domain Idx  Assn Idx
--------------------------------------------------------------------------------
md1                    ma1                    1           98

 Total number of Maintenance Association entries: 1.

================================================================================
                          Maintenance Association config
================================================================================
Domain Name            Assn Name
--------------------------------------------------------------------------------
md1                    ma1

 Total number of MA entries: 1.
```

show cfm maintenance-endpoint

```
Switch:1#show cfm maintenance-endpoint
================================================================================
                          Maintenance Endpoint Config
================================================================================
DOMAIN                  ASSOCIATION            MEP  ADMIN
NAME                    NAME                   ID
--------------------------------------------------------------------------------
md1                     ma1                    1    enable

 Total number of MEP entries: 1.

================================================================================
                          Maintenance Endpoint Service
================================================================================
DOMAIN_NAME             ASSN_NAME              MEP_ID TYPE    SERVICE_DESCRIPTION
--------------------------------------------------------------------------------
md1                     ma1                    1      unused

 Total number of MEP entries: 1.
```

show vlan nodal-mep

```
Switch:1#show vlan nodal-mep

================================================================================
                                 Vlan Nodal Mep
================================================================================
VLAN_ID    DOMAIN_NAME.ASSOCIATION_NAME.MEP_ID
--------------------------------------------------------------------------------
1
2
3
4          md1.ma1.1
5
6
7
8
9
10
11
12
13
14
```

show vlan nodal-mip-level

```
Switch:1#show vlan nodal-mip-level

================================================================================
                             Vlan Nodal Mip Level
================================================================================
VLAN_ID    NODAL_MIP_LEVEL_LIST
--------------------------------------------------------------------------------
1
2
3
4          6
5
6
7
8
9
10
11
```

```
12
13
14
```

Chapter 6: Software Troubleshooting tool configuration

Use the information in this chapter to learn about the methods and tools that you can use to troubleshoot and isolate problems in the switch.

This chapter includes the following sections:

- Troubleshooting Tool Fundamentals
- Software Troubleshooting Tool Configuration Using CLI
- Software Troubleshooting Tool Configuration Using EDM

Troubleshooting Tool Fundamentals

This section provides conceptual information about the methods and monitoring tools you can use for troubleshooting problems. This section also contains precautionary notices that you must read for the safe operation of the network.

Troubleshooting Overview

The types of problems that typically occur with networks involve connectivity and performance. This section also contains precautionary notices that you must read for the safe operation of the switch. The switch supports a diverse range of network architectures and protocols, some of which maintain and monitor connectivity and isolate connectivity faults.

In addition, the switch 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 switch network topologies. Other tools are more general in their application and you can use them to diagnose and monitor ingress and egress traffic on the switch.

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 cable and port 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 switch-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 ...... 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.

Debug Files

The switch stores debug files in the intflash directory.

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

- a memory snapshot
- logs
- traces

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

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

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

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

If you want to delete a specific file, you must use the **`remove`** command.

SNMP

The switch does not support SNMP for the **`show debug-file`** or the **`debug-file remove`** commands.

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. Transceivers in various form factors support DDM. Use the CLI 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 transceivers on a variety of switches and routers. 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 supported transceivers, see *Extreme Networks Pluggable Transceivers Installation Guide* and Monitoring Performance for VOSS.

Example

```
Switch: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
```

Flight Recorder

Table 10: Flight Recorder product support

Feature	Product	Release introduced
For configuration details, see Troubleshooting VOSS.		
Flight Recorder for system health monitoring	VSP 4450 Series	VSP 4000 4.0
	VSP 4900 Series	VOSS 8.1
	VSP 7200 Series	VOSS 4.2.1
	VSP 7400 Series	VOSS 8.0
	VSP 8200 Series	VSP 8200 4.0
	VSP 8400 Series	VOSS 4.2
	VSP 8600 Series	VSP 8600 4.5
	XA1400 Series	VOSS 8.0.50

The Flight Recorder is a high level term for the framework in place on the switch 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 outside of any process. 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 process stops responding
- the hardware watchdog activates
- the user initiates a snapshot in the CLI

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. The Always-On Trace feature uses circular logging, and therefore stores the most recent traces of the process.

Flight Recorder functionality is provided only through CLI. The following commands are used to make use of this feature:

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

  For VSP 8600, the valid slots are 1–8, SF1–SF3, all.

  The command creates a flight-recorder snapshot, trace and archive.

- **`flight-recorder archive {slot[-slot][,...]}`**

  For VSP 8600, the valid slots are 1–8, SF1–SF3, all.

  This command creates a tarball of flight-recorder files, log files, and configuration files.

- **`flight-recorder snapshot {slot[-slot][,...]}`**.

  For VSP 8600, the valid slots are 1–8, SF1–SF3, all.

  This command takes a snapshot of PMEM data.

- **`flight-recorder trace {slot[-slot][,...]}`**

  For VSP 8600, the valid slots are 1–8, SF1–SF3, all.

  This command takes snapshot of always-on-trace data.

Port Mirroring

Table 11: Port Mirroring product support

Feature	Product	Release introduced
For configuration details, see Troubleshooting VOSS.		
Ingress mirroring (port and flow-based)	VSP 4450 Series	VSP 4000 4.0
	VSP 4900 Series	VOSS 8.1
	VSP 7200 Series	VOSS 4.2.1
	VSP 7400 Series	VOSS 8.0
	VSP 8200 Series	VSP 8200 4.0
	VSP 8400 Series	VOSS 4.2
	VSP 8600 Series	VSP 8600 4.5
	XA1400 Series	VOSS 8.0.50
Egress mirroring (port-based)	VSP 4450 Series	VSP 4000 4.0
	VSP 4900 Series	VOSS 8.1
	VSP 7200 Series	VOSS 4.2.1
	VSP 7400 Series	VOSS 8.0
	VSP 8200 Series	VSP 8200 4.0

Table continues…

Feature	Product	Release introduced
	VSP 8400 Series	VOSS 4.2
	VSP 8600 Series	VSP 8600 4.5
	XA1400 Series	VOSS 8.0.50

True egress port-based mirroring that produces an identical copy of an outgoing packet is supported only on the VSP 4450 Series. On all other platforms, the mirrored copy does not reflect changes that occur in the switch to the outgoing packet (for example, packet fields that are updated during IP routing). As a result, the mirrored copy is not identical to the outgoing packet.

Use the port mirroring feature to 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. Egress mirroring on the switch is done at the end of the ingress pipeline. Since packet modifications occur in the egress pipeline, some of the changes will not be reflected in the mirrored version of the packet. Changes that occur in the egress pipleline may be reflected in the mirrored packed due to the metadata that is carried with the packet. Metadata notifies the egress pipeline what to change.

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 or a MultiLink Trunking (MLT) group. The switch supports one-to-many, many-to-one, and many-to-many mirroring configurations.

Ingress and egress mirrored ports

You can use all 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 is limited by the hardware. The hardware limitation is 4 ports simultaneously (where each mirroring direction counts as one). For example, if two mirroring ports are designated to mirror both ingress and egress traffic then all 4 mirroring ports are consumed.

The following table describes ingress mirroring functionality. 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. However, you can mirror to multiple physical destinations.

Important:

- Flow or ACL-based based mirroring is not supported for ingress and egress on VSP 8600 Series.
- Mirroring packets from one NNI port to another NNI port is not supported. Mirror to access ports, not NNI ports.

Table 12: Ingress mirroring functionality

Function	Support information
Ingress port mirroring and ingress flow mirroring	Supported. Maximum of 4 mirror-to-ports per box.
One port to one port	Supported
One to MLT group [for threat protection system (TPS applications)]	Supported
One to many (multicast group ID/VLAN)	Not supported
One to one (remote mirrored destination)	Not supported
Many to one (multiple mirrored ports to one mirroring port)	Supported
Many to MLT group	Supported
Many to many (VLAN/multicast group ID) (multiple ports with several different destinations)	Not supported
Many to one (relation between Remote Mirror Source [RMS] and Remote Mirror Termination [RMT])	Not 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

The following table describes egress mirroring functionality.

Table 13: Egress mirroring functionality

Function	Support information
Egress port mirroring	Supported
One port to one port	Supported
One to MLT groups (for TPS applications)	Supported
One to many (multicast group ID/VLAN)	Not supported
Many to one (multiple mirrored ports to one mirroring port)	Supported
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])	Not supported
VLAN and port combination as mirroring destination	Not supported

Table continues…

Function	Support information
Egress flow mirroring	Supported
Allow filter to specify a separate destination for each access control entry	Supported

Port configuration

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

There are two port mirroring modes: rx (ingress, that is, inPort) and tx (egress, that is, outPort). 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.

Note:

Not all hardware platforms support ACL global action to mirror.

To modify a port mirroring instance, first disable the instance. Also, to change a port or MLT entry, first remove whichever parameter is attached to the entry, and then add the required entry.

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 Configuring QoS and ACL-Based Traffic Filtering for VOSS.

Port mirroring considerations and restrictions

Although you can configure the switch to monitor both ingress and egress traffic, some restrictions apply:

- VSP 4900 Series, VSP 7200 Series, VSP 7400 Series, VSP 8000 Series, VSP 8600 Series, and XA1400 Series do not support true egress mirroring because packets are mirrored prior to the completion of packet processing, so egress mirrored packets can differ from the packets egressing the port.

  For the VOSS platforms, only VSP 4000 Series supports true egress mirroring.

  Note:

  To mirror the egress traffic for VSP 4900 Series, VSP 7200 Series, VSP 7400 Series, VSP 8000 Series, and VSP 8600 Series platforms, you can use the NEXT-hop device ingress mirroring to capture the egress packets of the switch.
- 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.
- Certain control packets generated by the CP cannot be egress mirrored, such as those in the following list:
  - BPDU
  - EAPoL
  - IP Directed Broadcast
  - LACP
  - LLDP
  - Multicast routed packets
  - NAAP
  - NLB
  - Nodal CFM
  - TDP
  - VLACP
- Ingress multicast packets appear in egress mirroring.
- On the VSP 7400 Series, if incoming traffic from the same source port is simultaneously ingress mirrored on an incoming port into one I-SID and egress mirrored on another outgoing port into a different I-SID, the mirrored packet carries an I-SID associated with ingress mirroring.

- On the VSP 7400 Series, any incoming traffic that does not contain a VLAN tag is not mirrored into an I-SID if the offset ID is in the range 2-1000. It is mirrored to an I-SID only if the offset ID is 1.
- On the XA1400 Series the global mirrored action is not supported.
- On the XA1400 Series you can configure the ACE to perform the mirroring operation. You can use the ACE mirror parameter, to configure a specific mirror to an MLT ID, port, or port list. Filters are not applied on mirrored packets.
- On the XA1400 Series you can configure a port mirror and a flow based ACE mirror action on the same port.

Port Mirroring Resources

Port mirroring resources are limited to four ports simultaneously (where each mirroring direction counts as one). For example, if two mirroring ports are designated to mirror both ingress and egress traffic then all four mirroring ports are consumed.

Port mirroring shares these four resources with other applications such as port mirroring RSPAN, Fabric Extend, Application Telemetry, IPFIX, and ACL with mirror action. Each one of these applications consumes at least one port mirroring resource. (port mirroring RSPAN consumes two if you configure both Ingress and Egress modes.)

Important:

- To enable any one of the above applications, you must have at least one free mirroring resource. If all four port mirroring resources are already in use, the switch displays a `Resource not available` error message when you try to enable the application.
- The VSP 8600 uses the four reserved resources for port mirroring and ACLs that have a mirroring action. For the other applications, this restriction does not apply because the VSP 8600 uses mirroring resources that do not come out of the four reserved port mirroring resources.

If you receive a `Resource not available` error message, you can use the **`show mirror-resources`** command to view information about mirror resource usage. For more information, see Displaying Mirror Resource Usage on page 184.

The **`show mirror-resources`** command is not available on all platforms. Use CLI command completion Help to determine if the command is available on your switch.

General Diagnostic Tools

The switch has diagnostic features available with Enterprise Device Manager (EDM) and Command Line Interface (CLI). 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, and viewing the address resolution table.

For more information about statistics, see Monitoring Performance for VOSS.

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 the switch. 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.

Caution:

Risk of traffic loss

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

Tip:

While these occurrences are uncommon, when using the trace level tool, minimize this risk. The following actions are recommended:

- 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.
- 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.

  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.

Fabric RSPAN (Mirror to I-SID)

Table 14: Fabric RSPAN product support

Feature	Product	Release introduced
For configuration details, see Troubleshooting VOSS.		
Fabric RSPAN (Mirror to I-SID)	VSP 4450 Series	VOSS 6.0 Flow-based mirroring into single I-SID only
	VSP 4900 Series	VOSS 8.1
	VSP 7200 Series	VOSS 6.0
	VSP 7400 Series	VOSS 8.0
	VSP 8200 Series	VOSS 6.0
	VSP 8400 Series	VOSS 6.0
	VSP 8600 Series	Not Supported
	XA1400 Series	Not Supported

Remote mirroring is an important functionality that helps in:

- Intrusion Detection or Intrusion Prevention Systems
- Network Port debugging and packet capture
- Mirror and Monitor traffic to central collector or analyzers
- Mirror and Monitor traffic to distributed collectors or analyzers

With the Fabric RSPAN (Mirror to I-SID) feature, mirrored traffic captured from any switch in the network is sent to a remote switch over an SPB cloud for traffic analysis. With this feature, you can monitor traffic on ports from different switches connected in the network, using just one network analyzer connected to a remote switch which acts as a collector. The source device where the traffic is mirrored to an I-SID, is known as Mirroring BEB (Backbone Edge Bridge), and the remote device where the traffic analyzer is connected for mirrored traffic analysis is known as Monitoring BEB.

The traffic source on the mirroring BEB is supported in the following ways:

- Port based mirroring — Any packet incoming or outgoing through a port is mirrored to a monitoring I-SID configured for that port.
- Flow based mirroring — Any particular packet flow configured in the system using filter based ACLs is mirrored to a monitoring I-SID configured for that flow.

Fabric RSPAN (Mirror to I-SID) limitations

- Remote mirroring of traffic is not supported on NNI ports, Fabric Extend Layer 2 core ports, and Open Networking Adapter (ONA) devices and ports.
- Mirroring resources will be shared between Fabric RSPAN and regular port mirroring. Fabric RSPAN uses one out of four resources for mirroring if the mode is configured as Rx (Ingress) mirroring. In case of mode Tx (Egress) mirroring, it uses one more entry with same TX-LB port. Hence if mode Rx and Tx are configured for Fabric RSPAN, then only two unique destination

ports can be used for regular port mirroring. For example, if you configure Fabric RSPAN on the switch, the regular port mirroring functionality can use only three unique destination ports. And, if all the four unique ports are used by the port mirroring functionality, you cannot configure Fabric RSPAN functionality on that node.

- When the monitor I-SID used for mirroring Fabric RSPAN traffic ingressing to I-SID is used to mirror regular traffic into SPB network, it will remove the customer tag in the mirrored packets. Hence, it is recommended that you use different monitor I-SIDs for mirroring regular traffic and Fabric RSPAN traffic.
- Monitoring egress-ports and egress-mlt will not support regular production network traffic.
- The QoS value must be same for all mirror entries having common monitor I-SID, as the BMAC QoS is mapped to a monitor I-SID. QoS value configured for a specific monitor I-SID offset overrides the existing value for all mirroring entries sharing the same monitor I-SID.
- Do not configure the source of mirrored traffic (mirroring to an I-SID) and the analyzer (monitoring an I-SID) on the same local device with the same I-SID offset. If you require mirroring and monitoring on the same local device, use standard port-based mirroring instead of Fabric RSPAN. Fabric RSPAN mirrors traffic into an I-SID of the SPB Fabric network and monitors traffic on the remote device; the network analyzer resides on the remote monitoring device and not on the same local device.
- Egress flow-based I-SID mirroring is not supported.

Software Troubleshooting Tool Configuration Using CLI

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

Using CLI for Troubleshooting

You can use CLI to provide diagnostic information.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Disable scrolling of the output display:

   ```
   terminal more disable
   ```

3. View configuration file information:

   ```
   more WORD<1-99>
   ```

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

   ```
   show running-config [verbose] [module {app-telemetry | boot | cfm |
   cli | diag | dvr | eap | endpoint-tracking | energy-saver | fa | fhs
   ```

```
| filter | ike | ip | ipfix | ipsec | ipv6 | iqagent | isis | i-sid
| lacp | license | lldp | lst | macsec | mlt | naap | nls | ntp |
ovsdb | port | qos | radius | restconf | rmon | sflow | security |
slamon | slpp | smtp | spbm | stg | sys | tacacs | virtualservice |
vlan | web | vxlan}]
```

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

   ```
   show tech
   ```

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

   Note:

   The **`show interfaces gigabitEthernet statistics rmon`** command displays information only if you previously configured **`rmon stats`** or **`rmon history`**.

   ```
   show interfaces gigabitEthernet statistics <dhcp-relay [vrf WORD<0-
   16>][vrfids WORD<0-512>] [{slot/port[/sub-port] [-slot/port[/sub-
   port]] [,...]}]
   ```

   ```
   show interfaces gigabitEthernet statistics lacp [{slot/port[/sub-
   port] [-slot/port[/sub-port]] [,...]}]
   ```

   ```
   show interfaces gigabitEthernet statistics rate—limiting [{slot/
   port[/sub-port] [-slot/port[/sub-port]] [,...]}]
   ```

   ```
   show interfaces gigabitEthernet statistics policer [{slot/port[/sub-
   port] [-slot/port[/sub-port]] [,...]}]
   ```

   ```
   show interfaces gigabitEthernet statistics rmon [{slot/port[/sub-
   port] [-slot/port[/sub-port]] [,...]}
   ```

   ```
   show interfaces gigabitEthernet statistics vlacp [{slot/port[/sub-
   port] [-slot/port[/sub-port]] [,...]}]
   ```

   ```
   show interfaces gigabitEthernet statistics verbose [{slot/port[/sub-
   port] [-slot/port[/sub-port]] [,...]}]
   ```

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

   Note:

   Product Notice: If you run this command on an XA1400 Series switch, most of the values report as zero (0). To see port-level error counters for XA1400 Series, use the `show io nic-counters` command instead.

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

Examples

Due to the length of command output, the following examples are truncated.

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

```
Switch:1#show running-config module cli
Preparing to Display Configuration...
#
# Thu Dec 19 18:38:02 2016 UTC
# box type             : Switch
# software version     : 4.5.0.1_B008 (PRIVATE)
# cli mode             : ECLI
#

config terminal

#
# CLI CONFIGURATION
#

cli timeout 65535
prompt "Switch-1"
```

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

```
Switch:1>show tech

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


General Info :

        SysDescr     : Switch (4.5.0.1_B008) (PRIVATE)
        SysName      : Switch-1
        SysUpTime    : 7 day(s), 18:23:36
        SysContact   : support@extremenetworks.com
        SysLocation  :

Chassis Info:

        Chassis            : Switch
        Serial#            : SDNIV84Q0010
        H/W Revision       : 1
        H/W Config         :
        NumSlots           : 2
        NumPorts           : 97
        BaseMacAddr        : b0:ad:aa:40:04:00
        MacAddrCapacity    : 1024

        MgmtMacAddr        : b0:ad:aa:40:04:81
        System MTU         : 1950
--More-- (q = quit)
```

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

```
Switch:1#show interfaces gigabitethernet statistics

================================================================================
                              Port Stats Interface
================================================================================
PORT     IN                   OUT                   IN                   OUT
NUM    OCTETS               OCTETS               PACKET               PACKET
--------------------------------------------------------------------------------
1/1    1215232              1852156              18988                25083
1/2    11866260             3650340              128847               51849
1/3    0                    0                    0                    0
1/4    0                    0                    0                    0
```

```
1/5    0                    0                    0                    0
1/6    2606433776           2605569408           40718802             40712022
1/7    2383797478           2368788480           37189478             37012320
1/8    2639779622           2624836140           41201664             40945760
1/9    0                    0                    0                    0
1/10   0                    0                    0                    0
1/11   0                    0                    0                    0
1/12   0                    6776546              0                    62572
1/13   1215232              997632               18988                15588
1/14   7459408              1396224              69625                18702
--More-- (q = quit)
```

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

```
Switch: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
--------------------------------------------------------------------------------
1/1   0       0       0       0       0       0       0       0       0
1/2   0       0       0       0       0       0       0       0       0
1/3   0       0       0       0       0       0       0       0       0
1/4   0       0       0       0       0       0       0       0       0
1/5   0       0       0       0       0       0       0       0       0
1/6   0       0       0       0       0       0       0       0       0
1/7   0       0       0       0       0       0       0       0       0
1/8   0       0       0       0       0       0       0       0       0
1/9   0       0       0       0       0       0       0       0       0
1/10  0       0       0       0       0       0       0       0       0
1/11  0       0       0       0       0       0       0       0       0
1/12  0       0       0       0       0       0       0       0       0
1/13  0       0       0       0       0       0       0       0       0
1/14  0       0       0       0       0       0       0       0       0

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

```
Switch:1#show io nic-counters

======== PORT 1 NIC STATS ========

bash-4.4# ethtool -S eth1 | grep -v queue
NIC statistics:
     rx_packets: 49378
     tx_packets: 1801
     rx_bytes: 3614051
     tx_bytes: 153217
     rx_broadcast: 11142
     tx_broadcast: 61
     rx_multicast: 37666
     tx_multicast: 1283
     multicast: 37666
     collisions: 0
     rx_crc_errors: 0
     rx_no_buffer_count: 0
     rx_missed_errors: 0
     tx_aborted_errors: 0
     tx_carrier_errors: 0
     tx_window_errors: 0
     tx_abort_late_coll: 0
     tx_deferred_ok: 0
     tx_single_coll_ok: 0
```

```
     tx_multi_coll_ok: 0
     tx_timeout_count: 0
     rx_long_length_errors: 0
     rx_short_length_errors: 0
     rx_align_errors: 0
     tx_tcp_seg_good: 0
     tx_tcp_seg_failed: 0
     rx_flow_control_xon: 0
     rx_flow_control_xoff: 0
     tx_flow_control_xon: 0
     tx_flow_control_xoff: 0
     rx_long_byte_count: 3614051
     tx_dma_out_of_sync: 0
     tx_smbus: 0
     rx_smbus: 0
     dropped_smbus: 0
     os2bmc_rx_by_bmc: 0
     os2bmc_tx_by_bmc: 0
     os2bmc_tx_by_host: 0
     os2bmc_rx_by_host:
     tx_hwtstamp_timeouts: 0
     tx_hwtstamp_skipped: 0
     rx_hwtstamp_cleared: 0
     rx_errors: 0
     tx_errors: 0
     tx_dropped: 0
     rx_length_errors: 0
     rx_over_errors: 0
     rx_frame_errors: 0
     rx_fifo_errors: 0
     tx_fifo_errors: 0
     tx_heartbeat_errors: 0
bash-4.4#
bash-4.4# exit
exit
--More-- (q = quit)
```

Variable Definitions

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

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>, /intflash/<file>.

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

Variable	Value
module {app-telemetry | boot | cfm | cli | diag | dvr | eap | endpoint-tracking | energy-saver | fa | fhs | filter | ike | ip | ipfix | ipsec | ipv6 | iqagent | isis | i-sid | lacp | license | lldp | lst | macsec | mlt | naap | nls | ntp | ovsdb | port | qos | radius | restconf | rmon | sflow | security | slamon | slpp | smtp | spbm | stg | sys | tacacs | virtualservice | vlan | web | vxlan}	Specifies the command group for which you request configuration settings. **Note:** All configuration modules are not supported on all hardware platforms. For information about feature support, see VOSS Feature Support Matrix.

Table continues…

Variable	Value
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`** command.

Variable	Value
dhcp-relay [vrf *WORD<1–16>*][vrfids *WORD<0–512> {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}*]	Displays port Dynamic Host Configuration Protocol (DHCP) statistics.
lacp *{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}*	Displays Link Aggregation Control Protocol (LACP) statistics.
rate-limiting *{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}*	Displays port ingress rate-limiting statistics.
rmon *{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]} [history]*	Displays Remote Network Monitoring (RMON) statistics.
verbose	Displays a complete list of all statistics.
vlacp *[history] {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}*	Displays port Virtual Link Aggregation Control Protocol (VLACP) statistics. • history—Displays the VLACP port counter profile. • *{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}* Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

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

Variable	Value
collision	Displays port collision error information.
ospf	Displays ports Open Shortest Path First (OSPF) error information.
verbose	Displays all port error information.
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Using Software Record Dumps

About this task

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

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Dump software record information:

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

Example

```
Switch:1> enable
Switch:1# dump ar 1 vlan 1
```

Variable Definitions

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

Variable	Value
<1>	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, protocol, all.
<0-3>	Specifies the verbosity from 0–3. Higher numbers specify more verbosity.

Using Trace to Diagnose Problems

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

About this task

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. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Clear the trace:

```
clear trace
```

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

   ```
   show trace modid-list
   ```

4. Begin the trace operation:

   ```
   trace level [<Module_ID>] [<0-4>]
   ```

5. Wait approximately 30 seconds.

   The default trace settings for CPU utilization are:

   - High CPU Utilization: 90%
   - High Track Duration: 5 seconds
   - Low CPU Utilization: 75%
   - Low Track Duration: 5 seconds

6. Stop tracing:

   ```
   trace shutdown
   ```

7. View the trace results:

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

8. 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

```
Switch:1> enable
```

Clear the trace:

```
Switch:1# clear trace
```

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

```
Switch:1# show trace modid-list
0   - COMMON
1   - SNMP
2   - RMON
3   - PORT_MGR
4   - CHAS_MGR
5   - BRIDGE
6   - HWIF
7   - SIM
8   - CPP
9  - NETDRV
10  - VLAN_MGR
```

```
11  - CLI
12  - MAIN
12  - P2IP
12  - RCIP
15  - WEBSRV
16  - ACIF
17  - GBIF
18  - WDT
19  - TDP
20  - MAN_DIAG
21  - MAN_TEST

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

Begin the trace operation:

```
Switch:1# trace level 2 3
```

Stop tracing:

```
Switch:1# trace shutdown
```

Save the trace file to the internal flash card for retrieval:

```
Switch:1# save trace
```

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

```
Switch:1# trace grep error
```

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

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

Variable Definitions

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

Variable	Value
grep *[WORD<0-128>]*	Search trace results for a specific string value, for example, the word error. Performs a comparison of trace messages.
level [<*Module_ID*>] *[<0-4>]*	Starts the trace by specifying the module ID and level. • <*Module_ID*> specifies the module for the trace. Different hardware platforms support different ID ranges because of feature support differences. To see which module IDs are available on the switch, use the **`show trace modid-list`** command or CLI command completion Help. • <0-4> specifies the trace level from 0–4, where 0 is disabled; 1 is very terse; 2 is terse; 3 is verbose, 4 is very 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.

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

Using Trace to Diagnose IPv6 Problems

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

Before you begin

- Confirm that trace level 99 is set to a value of 1 before you use trace to diagnose IPv6 problems. Trace level 1 is very terse.

Caution:

Risk of traffic loss

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

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

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

   ```
   trace ipv6 base <disable|enable> <all|debug|error|icmp|info|
   ipclient|nbr|pkt|warn> [vrf WORD <1-16>]
   ```

3. Activate or deactivate the trace for IPv6 forwarding:

   ```
   trace ipv6 forwarding <disable|enable> <all|debug|error|info|pkt|
   warn> [vrf WORD <1-16>]
   ```

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

   ```
   trace ipv6 nd <disable|enable> <all|debug|error|info|nbr|pkt|
   redirect|warn> [vrf WORD <1-16>]
   ```

5. Activate or deactivate the trace for IPv6 OSPF:

   ```
   trace ipv6 ospf <disable|enable> <adj|all|config|error|import|info|
   lsa|pkt|spf|warn> [vrf WORD <1-16>]
   ```

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

   ```
   trace ipv6 rtm <disable|enable> <all|change-list|debug|error|fib|
   info|redist|update|warn> [vrf WORD <1-16>]
   ```

7. Activate or deactivate the trace for IPv6 transport:

```
trace ipv6 transport <disable|enable> <all|common|tcp|udp> [vrf WORD
<1-16>]
```

8. Deactivate the trace to prevent service degradation:

```
trace shutdown
clear trace
```

Example

```
Switch:1>enable
```

Activate the trace for all the IPv6 base categories on the Management Router:

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

Activate the trace for all the IPv6 forwarding categories:

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

Activate the trace for all the IPv6 neighbor discovery categories:

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

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

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

Activate the trace for all the IPv6 transport caterories:

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

Deactivate the trace:

```
Switch:1#trace shutdown
Switch:1#clear trace
Removed 5 files.
```

Variable Definitions

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

Table 15: Variable definitions

Variable	Value
base <disable|enable> <all|debug|error|icmp|info|ipclient|nbr|pkt|warn>	Enables or disables a specific trace category for IPv6 base.
forwarding <disable|enable> <all|debug|error|info|pkt|warn>	Enables or disables a specific trace category for IPv6 forwarding.
nd <disable|enable> <all|debug|error|info|nbr|pkt|redirect|warn>	Enables or disables a specific trace category for IPv6 neighbor discovery.
ospf <disable|enable> <adj|all|config|error|import|info|lsa|pkt|spf|warn>	Enables or disables a specific trace category for IPv6 OSPF.

Table continues…

Variable	Value
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.
vrf WORD<1–16>	Specifies the VRF name.

View and Delete Debug Files

Use this procedure to view and delete debug files.

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

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

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

If you want to delete a specific file, you must use the **`remove`** command. For more information, see Configuring User Interfaces and Operating Systems for VOSS.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View debug files:

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

3. Delete debug files:

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

4. Enter Privileged EXEC mode:

   ```
   enable
   ```

5. View core files:

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

Example

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

```
Switch>show debug-file

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

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


================================================================================
                                  Archive Files
================================================================================
Directory: /intflash/archive/1
1.  File:    archive.20120611083021.1.tar
    Size:    34296320 bytes
    Created: Mon Jun 11 08:30:22 2012
2.  File:    archive.20120611163454.1.tar
    Size:    31108096 bytes
    Created: Mon Jun 11 16:34:54 2012
3.  File:    archive.20120611164354.1.tar
    Size:    31792128 bytes
    Created: Mon Jun 11 16:43:55 2012
4.  File:    archive.20120611164507.1.tar
     Size:    31881216 bytes
    Created: Mon Jun 11 16:45:08 2012

Remote CP Directory: /intflash//archive/2
1.  File:    archive.20120611163507.2.tar
    Size:    30903296 bytes
```

```
    Created: Mon Jun 11 16:35:08 2012
2.  File:    archive.20120611164408.2.tar
    Size:    31314432 bytes
    Created: Mon Jun 11 16:44:09 2012
3.  File:    archive.20120611164521.2.tar
    Size:    31367168 bytes
    Created: Mon Jun 11 16:45:21 2012

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

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


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

================================================================================
                                 DMalloc Files
================================================================================

================================================================================
                                  Flrec Files
================================================================================

================================================================================
                                 WdStats Files
================================================================================
Directory: /intflash/wd_stats/4
1.  File:    wd_stats.log.backup
    Size:    2311 bytes
    Created: Mon Jun 11 09:25:07 2012

Switch>debug-file remove 1
Switch>show debug-file

================================================================================
                                   Core Files
================================================================================
Remote CP Directory: /intflash/coreFiles/2
1.  File:    core.coreManager.x.20120612085548.2.tar
    Size:    1162240 bytes
    Created: Tue Jun 12 08:55:49 2012
```

```
2.  File:    core.coreManager.x.20120612085602.2.tar
    Size:    478208 bytes
    Created: Tue Jun 12 08:56:02 2012
3.  File:    core.coreManager.x.20120612085553.2.tar
    Size:    1170432 bytes
    Created: Tue Jun 12 08:55:56 2012
4.  File:    core.coreManager.x.20120612085558.2.tar
    Size:    1883136 bytes
    Created: Tue Jun 12 08:56:00 2012


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

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

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


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

================================================================================
                                  DMalloc Files
================================================================================

================================================================================
                                   Flrec Files
================================================================================

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

```
                                  WdStats Files
================================================================================
Directory: /intflash/wd_stats/4
1.  File:    wd_stats.log.backup
    Size:    2311 bytes
    Created: Mon Jun 11 09:25:07 2012
```

The following example shows how to view only core files on the switch.

```
Switch#show core-files
================================================================================
                                    Core Files
================================================================================
Directory: /intflash/coreFiles/1
1.  File:    core.1353113115.lifecycle.CP.1.gz
    Size:    139406 bytes
    Created: Fri Nov 16 19:45:15 2012
2.  File:    core.cbcp-main.x.20121114043335.1.tar
    Size:    14059520 bytes
    Created: Wed Nov 14 04:35:36 2012
3.  File:    core.cbcp-main.x.20121114045202.1.tar
    Size:    12809728 bytes
    Created: Wed Nov 14 04:54:03 2012
4.  File:    core.cbcp-main.x.20121114050825.1.tar
    Size:    12638720 bytes
    Created: Wed Nov 14 05:10:26 2012
5.  File:    core.cbcp-main.x.20121114122506.1.tar
    Size:    13020160 bytes
    Created: Wed Nov 14 12:27:07 2012
6.  File:    core.1353336274.lifecycle.CP.1.gz
    Size:    139390 bytes
    Created: Mon Nov 19 09:44:34 2012
7.  File:    core.1353319337.lifecycle.CP.1.gz
    Size:    139404 bytes
    Created: Mon Nov 19 05:02:17 2012
8.  File:    core.cbcp-main.x.20130122182946.1.tar
    Size:    13683712 bytes
    Created: Tue Jan 22 18:32:08 2013
9.  File:    core.cbcp-main.x.20130220143809.1.tar
    Size:    13969920 bytes
    Created: Wed Feb 20 14:38:10 2013
10.  File:    core.cbcp-main.x.20130225155025.1.tar
    Size:    13526016 bytes
    Created: Mon Feb 25 15:50:25 2013
11.  File:    core.cbcp-main.x.20130225155407.1.tar
    Size:    12674560 bytes
    Created: Mon Feb 25 15:54:07 2013
```

Variable Definitions

Use the data in the following table to use the **`show core-files`** command.

Variable	Value
{slot[-slot][,...]}	Displays the core files for the slot that you select.

Use the data in the following table to use the **`show debug-file`** command.

Variable	Value
all	Displays all types of debug files

Table continues…

Variable	Value
{slot[-slot][,...]}	Displays debug files for the slot that you select. If you do not select a slot number, the device displays all types of the archived debug files present in a slot by slot basis. If you select a slot number, the device only displays archived files for the slot you select.

Use the data in the following table to use the **debug-file remove** command.

Variable	Value
all	Removes all types of debug files in all slots. If you use the option **all** with the **remove debug-file** command, then the device deletes all types of debug files, including the latest debug files.
{slot[-slot][,...]} **Note:** Exception: only supported on VSP 8600 Series.	Removes debug files for the slot that you select. When you clear archived files, if you do not select a slot number, the device deletes all types of archived debug files except the latest file in each slot. Valid slots are 1–12, SF1–SF6, and all.

Configuring Port Mirroring

Use port mirroring to aid in diagnostic and security operations.

About this task

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.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Create a port mirroring instance:

   ```
   mirror-by-port <1-479> in-port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]} {monitor-mlt <1-512>| out-port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}
   ```

3. Create an I-SID mirroring instance:

```
mirror-by-port <1-479> [in-port {slot/port[/sub-port] [-slot/port[/
sub-port]] [,...]} monitor-isid-offset <1-1000> [mode <rx|tx|both>]
[qos <qos-level>]]
```

4. Configure the mode:

   ```
   mirror-by-port <1-479> mode <both|rx|tx>
   ```

   Note:

   - When you configure `tx` mode port mirroring on T-UNI and SPBM NNI ports, unknown unicast, broadcast and multicast traffic packets that ingress these ports appear on the mirror destination port, although they do not egress the mirror source port. This is because `tx` mode port mirroring happens on the mirror source port *before* the source port squelching logic drops the packets at the egress port.
   - The available four mirroring resources are shared between Fabric RSPAN and regular port mirroring, and are allocated based on the mode configured, Ingress (`rx`) or Egress (`tx`). Each configured mode occupies one mirroring resource, but when you configure the mode as `both`, it occupies two mirroring resources (one for Rx and one for Tx).
   - Do not configure the source of mirrored traffic (mirroring to an I-SID) and the analyzer (monitoring an I-SID) on the same local device with the same I-SID offset. If you require mirroring and monitoring on the same local device, use standard port-based mirroring instead of Fabric RSPAN. Fabric RSPAN mirrors traffic into an I-SID of the SPB Fabric network and monitors traffic on the remote device; the network analyzer resides on the remote monitoring device and not on the same local device.

5. Enable the mirroring instance:

   ```
   mirror-by-port <1-479> enable
   ```

6. Modify existing mirroring entries as required:

   ```
   mirror-by-port mirror-port <1-479> {slot/port[/sub-port] [-slot/
   port[/sub-port]] [,...]}
   ```

   OR

   ```
   mirror-by-port monitor-mlt <1-479> <1-512>
   ```

   OR

   ```
   mirror-by-port monitor-port <1-479> {slot/port[/sub-port] [-slot/
   port[/sub-port]] [,...]}
   ```

   Note:

   Before you can modify an existing entry, you must disable the entry: `no mirror-by-port <1-479> enable`.

7. Modify QoS value for Fabric RSPAN mirroring session:

   ```
   mirror-by-port <1-479> qos <0-5>
   ```

8. Verify the configuration:

   ```
   show mirror-by-port
   ```

Example

Port mirroring configuration:

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

Create the port mirroring instance:

```
Switch:1(config)# mirror-by-port 8 in-port 1/15 out-port 1/1
```

The analyzer connects to port 1/1.

Disable the entry:

```
Switch:1(config)# no mirror-by-port 8 enable
```

Mirror both ingress and egress traffic passing through port 1/16:

```
Switch:1(config)# mirror-by-port 8 mode both
```

Enable mirroring for the instance:

```
Switch:1(config)# mirror-by-port 8 enable
```

Fabric RSPAN configuration:

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

Create the Fabric RSPAN mirroring instance:

```
Switch:1(config)#mirror-by-port 3 in-port 1/3 monitor-isid-offset 3 mode
both qos 3
```

Disable the entry:

```
Switch:1(config)# no mirror-by-port 3 enable
```

Mirror the egress traffic passing through port 1/3:

```
Switch:1(config)# mirror-by-port 3 mode tx
```

Enable Fabric RSPAN for the instance:

```
Switch:1(config)# mirror-by-port 3 enable
```

The sample command output in the following example does not necessarily reflect the preceding examples.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch:1(config)#show mirror-by-port

================================================================================
                              Diag Mirror-By-Port
```

```
================================================================================
ID   MIRRORED_PORT   MIRRORING_DEST  ENABLE   MODE  REMOTE-MIRROR   DSCP TTL
                                                        VLAN-ID
--------------------------------------------------------------------------------
1     1/1              2/1           true     both    0   0             64
2     1/2              2/2           true     rx      0   0             64
3     1/3              2/3           true     tx      0   0             64
4     1/4              2/4           true     both    0   0             64
```

Variable Definitions

Use the data in the following table to use the **mirror-by-port** command.

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[/sub-port] [-slot/port[/sub-port]] [,...]}*{|monitor-mlt *<1-512>* |out-port *{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}*	Creates a new mirror-by-port table entry. • in-port *{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}* specifies the mirrored port. • monitor-mlt *<1-512>* specifies the mirroring MLT ID from 1–512. • out-port *{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}* specifies the mirroring port.
mirror-port <1-479> *{slot/port[/sub-port] [-slot/port[/sub-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 mirroring. The destination must be an IP address {A.B.C.D}. The default DSCP is 0 and the default TTL is 255.
monitor-mlt <1-479> <1-512>	Modifies the monitoring MLT.*<1-512>* specifies the mirroring MLT ID. 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[/sub-port] [-slot/port[/sub-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> <1-4059>*	Modifies the monitoring VLAN. Before you can modify an existing entry, you must disable the entry: `no mirror-by-port <1-479> enable`. Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags,

Table continues…

Variable	Value
	the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
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.
monitor-isid-offset <1-1000>	Specifies the offset ID that is mapped to the actual monitor I-SID where packets are mirrored. `Monitor I-SID = base monitor I-SID + offset ID.` The base monitor I-SID is 16776000.
qos <0-5>	Specifies the Quality of Service (QoS) profiles for the system. Monitoring I-SID supports six different QoS levels, each QoS level can be configured individually. Default value is 1.

Displaying Mirror Resource Usage

About this task

Mirror resources can be consumed by internal processes that are not easily identified, which can lead to unexpected resource shortages. Use the following procedure to display mirror resource usage.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display mirror resource usage:

   ```
   show mirror-resources
   ```

Example

```
Switch:1#show mirror-resources
    <Switch> Mirror Resource Manager stats:
            2 unique mirror destinations allocated, system max is 4
    ------------------------------------------------------------------------------
     Mirror      | Mirror    | Total | Mirror
     Destination | Direction | Refs  | Users
    ------------------------------------------------------------------------------
     Port 1/2    | Egress    |   1   | MBP ID 1
     Port 1/2    | Ingress   |   1   | MBP ID 1
    ------------------------------------------------------------------------------

    Function execution completed successfully
```

Configuring Global Mirroring Actions with an ACL

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

Before you begin

- The ACL exists.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure the global action for an ACL:

   ```
   filter acl set <1-2048> global-action {monitor-dst-mlt <1-512>|
   monitor-dst-ports {slot/port[/sub-port][-slot/port[/sub-port]]
   [,...]}
   ```

Example

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

Configure the global action for an ACL:

```
Switch: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.

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[/sub-port] [-slot/port[/sub-port]] [,...]}*	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.

Configure ACE Actions to Mirror

Before you begin

- The access control entry (ACE) exists.

About this task

If you use the mirror action, ensure that you specify the mirroring destination: MLTs or ports.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure actions for an ACE:

   ```
   filter acl ace action <1-2048> <1-2000> {permit|deny} monitor-dst-mlt <1-512>
   ```

   OR

   ```
   filter acl ace action <1-2048> <1-2000> {permit|deny} monitor-dst-ports {slot/port[/sub-port][-slot/port[/sub-port]][,...]}
   ```

3. Ensure the configuration is correct:

   ```
   show filter acl action [<1-2048>] [<1-2000>]
   ```

Example

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# filter acl ace action 901 1 permit monitor-dst-mlt 5
```

Variable Definitions

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

Variable	Value
1-2048	Specifies the ACL ID from 1–2048
1-2000	Specifies the ACE ID from 1–2000.
monitor-dst-mlt <*1–512*>	Configures mirroring to a destination MLT group.
monitor-dst-ports *{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}*	Configures mirroring to a destination port or ports.
{permit|deny}	Configures the action mode for security ACEs. The default value is permit.

Clearing ARP Information for an Interface

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

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Clear ARP information:

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

OR

```
clear ip arp interface vlan <1-4059>
```

Example

```
Switch:1> enable
Switch:1# clear ip arp interface gigabithethernet 1/1
```

Variable Definitions

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

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Flushing Routing, MAC, and ARP Tables for a Port

Flush or clear the routing tables for administrative and troubleshooting purposes. The clear and flush commands perform the same function; they remove the contents of the table.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

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

Note:

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

2. Flush IP routing tables by port:

   ```
   action flushIp
   ```

3. Flush the MAC address tables by port:

   ```
   action flushMacFdb
   ```

4. Flush ARP tables by port:

   ```
   action flushArp
   ```

5. Flush all tables with one command:

   ```
   action flushAll
   ```

6. Exit to Global Configuration mode:

   ```
   exit
   ```

7. Clear a routing table for a port:

   ```
   clear ip route gigabitethernet {slot/port[sub-port]}
   ```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface gigabitEthernet 1/1
Switch:1(config-if)#action flushAll
Switch:1(config-if)#exit
Switch:1(config)#clear ip route gigabitethernet 1/1
```

Variable Definitions

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

Variable	Value
{slot/port[/sub-port]}	Identifies a single slot and port. If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Flushing Routing, MAC, and ARP Tables for a VLAN

Flush or clear the routing tables for administrative and troubleshooting purposes. The clear and flush commands perform the same function; they remove the contents of the table.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Flush IP routing tables by VLAN:

   ```
   vlan action <1-4059> flushIp
   ```

3. Flush the MAC address tables by VLAN:

   ```
   vlan action <1-4059> flushMacFdb
   ```

4. Flush ARP tables by VLAN:

   ```
   vlan action <1-4059> flushArp
   ```

5. Flush all tables with one command:

   ```
   vlan action <1-4059> all
   ```

6. Clear a routing table for a VLAN:

   ```
   clear ip route vlan <1-4059>
   ```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config-if)#vlan action 123 all
Switch:1(config)#clear ip route vlan 123
```

Variable Definitions

Use the data in the following table to use the **`vlan action`** and **`clear ip route vlan`** commands.

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.

Pinging an IP Device

About this task

Ping a device to test the connection between the switch 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 can fail for VRF routes if you use large packet sizes for the operation. Do not use packet sizes larger than the following:

- Ping for VRF: 1480 bytes
- Traceroute for VRF: 1444 bytes

A management instance ID can be specified to allow the OS to use the correct source for the outgoing ICMP ECHO request packet.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Ping an IP network connection:

   ```
   ping WORD<0-256> [-d] [-I <1-60>] [-s] [-t <1-120>] [count <1-9999>]
   [datasize <28-9216|28-51200>] [interface gigabitEthernet {slot/
   port[sub-port]}|mgmtEthernet mgmt | tunnel <1-2000> | vlan <1-4059>]
   [scopeid <1-9999>] [source WORD<1-256>] [vrf WORD<1-16>]
   ```

   Note:

   The mgmtEthernet and mgmt interface only applies to hardware with a dedicated, physical management interface.

3. Ping a network connection using a Segmented Management Instance:

   ```
   ping WORD<0-256> [-s] [-t <1-120>] [count <1-9999>] [datasize
   <28-9216|28-51200>] mgmt [clip | vlan]
   ```

   Note:

   If you do not use the mgmt parameter, the **ping** command uses the VOSS IP routing stack to initiate the ping request.

Example

Ping an IP network connection through the management interface for IPv4, and for IPv6:

```
Switch:1>ping 192.0.2.2 vrf mgmtrouter
Switch:1>ping 2001:0db8:0000:0000:0000:0000:0000:0001 vrf mgmtrouter
```

Ping an IP device from a GRT VLAN IP interface:

```
Switch:1#ping 192.0.2.16
192.0.2.16 is alive
```

Ping a device using the management routing table:

```
Switch:1#ping 192.0.2.12 mgmt
```

Ping a device using a management CLIP:

```
Switch:1#ping 192.0.2.12 mgmt clip
```

Ping an IP device using a management VLAN:

```
Switch:1#ping 192.0.2.12 mgmt vlan
```

Variable Definitions

The following table defines parameters for the **`ping`** command.

Variable	Value
count <1–9999>	Specifies the number of times to ping. The default is 1.
-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). This parameter does not apply if you use the mgmt [clip | vlan] parameter.
datasize *<28-9216|28-51200>*	Specifies the size of ping data sent in bytes. The datasize for IPv4 addresses is 28-9216. The datasize for IPv6 addresses is 28-51200. The default is 64.
-I <1–60>	Specifies the interval between transmissions in seconds. This parameter does not apply if you use the mgmt [clip | vlan] parameter.
interface *gigabitEthernet {slot/port[sub-port]}| mgmtEthernet mgmt| tunnel <1–2000> | vlan <1-4059>*	Specifies the outgoing interface. Additional ping interface parameters: • gigabitEthernet {slot/port[sub-port]}: gigabitethernet port • mgmtEthernet mgmt: identifies the physical management port • tunnel: tunnel ID as a value from 1 to 2000 • vlan: Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. This parameter does not apply if you use the mgmt [clip | vlan] parameter.

Table continues…

Variable	Value
mgmt [clip | vlan]	Specifies the Segmented Management Instance as the source for the outgoing ICMP ECHO packet. The packet goes out this specific interface only. If you do not specify the management interface type, the **ping** command uses the management routing table to determine the best management interface and selects the source IP based on the egress management interface.
-s	Configures the continuous ping at the interval rate defined by the [-I] parameter or until you enter a **Ctrl + C** keystroke.
scopeid *<1–9999>*	Specifies the circuit ID for IPv6. This parameter does not apply if you use the mgmt [clip | vlan] parameter.
source *WORD<1-256>*	Specifies the source IP address for the ping command. This parameter does not apply if you use the mgmt [clip | vlan] parameter.
-t <1–120>	Specifies the no-answer timeout value in seconds. The default is 5.
WORD<0–256>	Specifies the host name or IPv4 (a.b.c.d) or IPv6 (x:x:x:x:x:x:x:x) address.
vrf WORD<1–16>	Specifies the virtual router and forwarder (VRF) name. This parameter does not apply if you use the mgmt [clip | vlan] parameter.

Running a Traceroute Test

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

About this task

Ping and traceroute can fail for VRF routes if you use large packet sizes for the operation. Do not use packet sizes larger than the following:

- Ping for VRF: 1480 bytes
- Traceroute for VRF: 1444 bytes

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Run a traceroute test:

```
traceroute WORD<0-256> [<1-1176>] [-m <1-255>] [-p <1-65535>] [-q
<1-255>] [-v] [-w <1-255>] [source <WORD 1-256>] [vrf <WORD 1-16>]
```

3. Run a traceroute test using a Segmented Management Instance:

```
traceroute WORD<0-256> [<1-1176>] [-m <1-255>] [-p <1-65535>] [-q
<1-255>] [-w <1-255>] mgmt [<clip | vlan>]
```

Note:

If you do not use the mgmt parameter, the **traceroute** command uses the VOSS IP routing stack to initiate the traceroute request.

Example

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

```
Switch:1>enable
Switch:1#traceroute 192.0.2.33 200 –m 60
```

Run a traceroute test for an IPv6 address:

```
Switch:1#traceroute 2001:db8::1
```

Run a traceroute test using the management routing table:

```
Switch:1#traceroute 192.0.2.12 mgmt
```

Run a traceroute test using a management CLIP:

```
Switch:1#traceroute 192.0.2.12 mgmt clip
```

Run a traceroute test using a management VLAN:

```
Switch:1#traceroute 192.0.2.12 mgmt vlan
```

Variable Definitions

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

Variable	Value
-m *<1-255>*	Specifies the maximum time-to-live (TTL).
-p *<1-65535>*	Specifies the base UDP port number. The default is 33434. **Note:** If the traceroute action is directed to an IPv6 host address, Linux increments the UDP port on a per-TTL basis. For an IPv4 host address, Linux increments the UDP port on a per-probe basis. Because the **traceroute** command sends a default of three probes, three incrementing ports will be sent for an IPv4 host address. If you use the -p parameter with a value greater than 65533, the **traceroute** command fails for an IPv4 host address because the maximum port number, 65535, is exceeded.

Table continues…

Variable	Value
	If you send a traceroute probe into the device through the Segmented Management Instance, you must use the default UDP port range of 33434 to 33464. Using other ports will fail.
-q <*1-255*>	Specifies the number of probes per TTL.
-v	Specifies verbose mode (detailed output). This parameter does not apply if you use the mgmt [clip | vlan] parameter.
-w <*1-255*>	Specifies the wait time for each probe.
<*1-1176*>	Specifies the size of the probe packet.
mgmt [<clip | vlan>]	Specifies the Segmented Management Instance as the source for the outgoing packet. The packet goes out this specific interface only. If you do not specify the management interface type, the **`traceroute`** command uses the management routing table to determine the best management interface and selects the source IP based on the egress management interface.
source <*WORD 1-256*>	Specifies the source IP address. This parameter does not apply if you use the mgmt [clip | vlan] parameter.
WORD<0-256>	Specifies the destination IPv4 or IPv6 address, or hostname.
vrf <*WORD 1-16*>	Specifies the VRF instance by VRF name. This parameter does not apply if you use the mgmt [clip | vlan] parameter.

Showing SNMP Logs

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

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Show the logs:

   ```
   show fulltech file WORD<1-99>
   ```

Variable Definitions

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

Variable	Value
file *WORD<1-99>*	This variable represents the log file to be opened and displayed. It is displayed in the following format: • /intflash/<file>

Using Trace to Examine IS-IS Control Packets

Use trace as a debug tool to examine the code flow and Intermediate-System-to-Intermediate-System (IS-IS) control packets. When you enable IS-IS trace flags, only trace information about the set flag appears.

Before you begin

- You must know what you want to trace before you enable trace.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Enable the Intermediate-System-to-Intermediate-System trace flags:

   ```
   trace flags isis set { none | tx-hello | rx-hello | tx-pkt | rx-pkt
   | adj | opt | tx-lsack | rx-lsack | tx-lsp | rx-lsp | pkt-err | nbr-
   mismatch | flood | prefix | nbr-change | intf-change | decide | fdb
   | dr | dd-masterslave | auth-fail | config | purge | policy | redist
   | tx-snp | rx-snp | timer | global | perf | ucast-fib | node |
   mcast-fib | isid | ip-shortcut }
   ```

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

   ```
   show trace modid-list
   ```

4. Clear the trace:

   ```
   clear trace
   ```

5. Begin the trace operation:

   ```
   trace level [<Module_ID>] [<0-4>]
   ```

   OR

   ```
   trace spbm isis level [<0-4>]
   ```

   OR

   ```
   trace cfm level [<0-4>]
   ```

Note:

- Module ID 119 represents the IS-IS module.

6. 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

7. Stop tracing:

   ```
   trace shutdown
   ```

8. View the trace results:

   ```
   trace screen enable
   ```

   Important:

   If you use trace level 3 (verbose) or trace level 4 (very verbose), do not use the screen to view commands due to the volume of information the system generates and the effect on the system.

9. Save the trace file.

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

   If you do not specify a file name, the file name is systrace.txt.

10. Display trace lines saved to a file:

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

11. Search trace results for a specific string value:

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

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

12. Stop tracing:

    ```
    trace shutdown
    ```

13. Disable the Intermediate-System-to-Intermediate-System trace flags:

    ```
    trace flags isis remove { none | tx-hello | rx-hello | tx-pkt | rx-
    pkt | adj | opt | tx-lsack | rx-lsack | tx-lsp | rx-lsp | pkt-err |
    nbr-mismatch | flood | spf-intra | spf-inter | spf-extern | prefix |
    nbr-change | intf-change | decide | fdb | dr | dd-masterslave |
    auth-fail | config | purge | policy | redist | tx-snp | rx-snp |
    timer | spbm-decide | global | perf | ucast-fib | node | mcast-fib |
    isid | ip-shortcut }
    ```

Example

```
Switch:1> enable
```

Clear prior trace information:

```
Switch:1# clear trace
```

Enable IS-IS trace flags for received IS-IS hello packets:

```
Switch:1# trace flags isis set rx-hello
```

Enable IS-IS trace flags for transmitted IS-IS hello packets:

```
Switch:1# trace flags isis set tx-hello
```

Configure the module ID to 119 (IS-IS module) and the trace to 4 (very verbose):

```
Switch:1# trace level 119 4
```

Enable the display of trace output to the screen:

```
Switch:1# trace screen enable
Switch:1# Screen tracing is on
```

Disable the display of trace output to the screen:

```
Switch:1# trace screen disable
Switch:1# Screen tracing is off
```

Variable Definitions

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

Variable	Value
remove *{ none | tx-hello | rx-hello | tx-pkt | rx-pkt | adj | opt | tx-lsack | rx-lsack | tx-lsp | rx-lsp | pkt-err | nbr-mismatch | flood | prefix | nbr-change | intf-change | decide | fdb | dr | auth-fail | config | purge | policy | redist | tx-snp | rx-snp | timer | global | perf | ucast-fib | node | isid | ip-shortcut }*	Removes the Intermediate-System-to-Intermediate-System (IS-IS) trace flags for the specified option.
set *{ none | tx-hello | rx-hello | tx-pkt | rx-pkt | adj | opt | tx-lsack | rx-lsack | tx-lsp | rx-lsp | pkt-err | nbr-mismatch | flood | prefix | nbr-change | intf-change | decide | fdb | dr | auth-fail | config | purge | policy | redist | tx-snp | rx-snp | timer | global | perf | ucast-fib | node | isid | ip-shortcut }*	Enables the Intermediate-System-to-Intermediate-System (IS-IS) trace flags for the specified option. • none — none • tx-hello — Transmitted IS-IS hello packet • rx-hello — Received IS-IS hello packet • tx-pkt — Transmitted packet • rx-pkt — Received packet • adj — Adjacencies • opt — IS-IS TLVs

Table continues…

Variable	Value
	• tx-lsack — Transmitted LSP acknowledgement • rx-lsack — Received LSP acknowledgement • tx-lsp — Transmitted Link State Packet • rx-lsp — Received Link State Packet • pkt-err — Packet Error • nbr-mismatch — Neighbor mismatch • flood — Flood • prefix — Prefix • nbr-change — Neighbor change • intf-change — IS-IS circuit (interface) events • decide — Shortest Path First computation • fdb — Filtering Database • dr — Designated Router • auth-fail — Authorization failed • config — Configuration • purge — Link State Packet purge • redist — Redistribute • tx-snp — Transmitted Sequence Number PDU (CSNP and PSNP) • rx-snp — Received Sequence Number Packet (CSNP and PSNP) • timer — Timer • perf — SPBM performance • ucast-fib — Unicast Forwarding Information Base • node — Node • isid — I-SID • ip-shortcut — IP Shortcut

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

Variable	Value
auto	Displays the current configuration for the automatic trace function.
file [tail]	Displays the trace results saved to a file.
level	Displays the current trace level for all modules.
modid-list	Specifies the module ID list.

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

Variable	Value
grep *[WORD<0-128>]*	Specifies the search keyword. You can use a specific MAC address. You can search for errors, using the command, **trace grep error**.
cfm level *[<0-4>]*	Starts tracing by CFM. • <0-4> specifies the trace level from 0–4, where 0 is disabled; 1 is very terse; 2 is terse; 3 is verbose, 4 is very verbose.
spbm isis level *[<0-4>]*	Specifies exactly which IS-IS component to display. • <0-4> specifies the trace level from 0–4, where 0 is disabled; 1 is very terse; 2 is terse; 3 is verbose, 4 is very verbose.
level [<*Module_ID*>] *[<0-4>]*	Starts the trace by specifying the module ID and level. <*Module_ID*> specifies the module for the trace. Different hardware platforms support different ID ranges because of feature support differences. To see which module IDs are available on the switch, use the **show trace modid-list** command or CLI command completion Help. 0–4 specifies the trace level: • 0 — Disabled • 1 — Very terse • 2 — Terse • 3 — Verbose • 4 — Very verbose
shutdown	Stops the trace operation.
screen {disable|enable}	Enables or disables the display of trace output to the screen. **Important:** Avoid using the screen to view commands if you use trace level 3 (verbose) or trace level 4 (very verbose) due to the volume of information generated and the effect on the system.

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

Variable	Value
file *WORD<1–99>*	Specifies the file name in one of the following formats: • a.b.c.d: <file> WORD<1–99> is a string of 1–99 characters. **Note:** If you do not specify a file name, the file name is systrace.txt.

Viewing the Metric Type of IS-IS Route in TLVs – Detailed

About this task

Use the following procedure to view the detailed information about metric type of IS-IS routes in TLVs in Link State Packets (LSP).

Procedure

1. Display the detail view of TLV 135:

   ```
   show isis lsdb tlv 135 detail
   ```

2. Display the detail view of TLV 184:

   ```
   show isis lsdb tlv 184 detail
   ```

Example

Viewing the metric type of IS-IS route in TLV 135

```
Switch:1#show isis lsdb tlv 135 detail

================================================================================
                               ISIS LSDB (DETAIL)
================================================================================
--------------------------------------------------------------------------------

Level-1 LspID: 4072.0000.0000.00-02     SeqNum: 0x00000009      Lifetime:  1110
        Chksum: 0x31ce  PDU Length: 46
        Host_name: evp4k
        Attributes:     IS-Type 1
TLV:135 TE IP Reachability: 2
                Metric: 1 Metric Type:Internal  Prefix Length: 32
                UP/Down Bit: FALSE              Sub TLV Bit: FALSE
                IP Address: 15.15.15.72
                Metric: 1 Metric Type:External  Prefix Length: 24
                UP/Down Bit: FALSE              Sub TLV Bit: FALSE
                IP Address: 192.0.2.5
```

Viewing the metric type of IS-IS route in TLV 184

```
Switch:1#show isis lsdb tlv 184 detail

================================================================================
                               ISIS LSDB (DETAIL)
================================================================================
--------------------------------------------------------------------------------

Level-1 LspID: 4072.0000.0000.00-03     SeqNum: 0x00000008      Lifetime:  1103
        Chksum: 0x3ce6  PDU Length: 72
        Host_name: evp4k
        Attributes:     IS-Type 1
TLV:184 SPBM IPVPN Reachability:
                Vrf ISID:100
                        Metric:1 Metric Type:External   Prefix Length:32
                        IP Address: 192.0.2.3
                        Metric:1 Metric Type:Internal   Prefix Length:32
                        IP Address: 192.0.2.72
```

Viewing the Metric Type of IS-IS Route in TLVs – Summarized

About this task

Use the following procedure to view the summarized information about metric type of IS-IS routes in TLVs. You can also view the metric type of the prefix.

Procedure

Display the summarized view of TLVs 135 and 184:

```
show isis lsdb ip-unicast
```

Example

Display the summarized view of TLVs.

```
evp4k:1#show isis lsdb ip-unicast

================================================================================
                          ISIS IP-UNICAST-ROUTE SUMMARY
================================================================================
                            PREFIX          METRIC    TLV   LSP   HOST
I-SID     ADDRESS           LENGTH  METRIC  TYPE      TYPE  FRAG  NAME
--------------------------------------------------------------------------------
-         192.0.2.5         32      1       Internal  135   0x2   evp4k
-         192.0.2.20        24      1       External  135   0x2   evp4k
100       192.0.2.15        32      1       External  184   0x3   evp4k
100       192.0.2.5         32      1       Internal  184   0x3   evp4k
-         192.0.2.4         32      1       Internal  135   0x2   esp1
--------------------------------------------------------------------------------
5 out of 5 Total Num of Entries
```

Configuring I-SID Monitoring

Use the following procedure to configure Fabric RSPAN (Mirror to I-SID) on the Backbone Edge Bridge (BEB) connected to the monitor station.

Note:

Do not configure the source of mirrored traffic (mirroring to an I-SID) and the analyzer (monitoring an I-SID) on the same local device with the same I-SID offset. If you require mirroring and monitoring on the same local device, use standard port-based mirroring instead of Fabric RSPAN. Fabric RSPAN mirrors traffic into an I-SID of the SPB Fabric network and monitors traffic on the remote device; the network analyzer resides on the remote monitoring device and not on the same local device.

Note:

If you change the egress port or egress MLT for a particular session using a separate CLI command, it overwrites the existing egress port list or egress MLT.

Procedure

1. Enter Global Configuration mode:

   ```
   enable

   configure terminal
   ```

2. Create a monitor by I-SID entry:

   Note:

   If you use the Insight port 1/s1 as the analyzer port on the monitoring BEB for remote mirroring, you must associate it to VLAN ID 4091.

   ```
   monitor-by-isid <1-1000> [monitor-isid-offset <1-1000> {egress-mlt
   <1-512> | egress-ports {slot/port[/sub-port][-slot/port[/sub-port]]
   [,...]}} [map-to-vid <1-4093>]]
   ```

3. Configure map to VLAN ID:

   ```
   monitor-by-isid <1-1000> map-to-vid <1-4093>
   ```

4. Configure egress MLT:

   ```
   monitor-by-isid <1-1000> egress-mlt <1-512>
   ```

5. Configure egress port:

   ```
   monitor-by-isid <1-1000> egress-ports {slot/port[/sub-port][-slot/
   port[/sub-port]][,...]}
   ```

6. Enable monitoring by I-SID entry:

   ```
   monitor-by-isid <1-1000> enable
   ```

   Note:

   Disable the entries (egress ports, MLT, and VLAN ID) to modify or remove parameters in the existing configuration.

Example

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# monitor-by-isid 1 monitor-isid-offset 1 egress-port 1/6
Switch:1(config)# monitor-by-isid 2 monitor-isid-offset 2 egress-port 1/7 map-to-vid 200
Switch:1(config)# monitor-by-isid 3 monitor-isid-offset 3 egress-port 1/7 map-to-vid 201
Switch:1(config)# monitor-by-isid 2 egress-port 1/8
Switch:1(config)# monitor-by-isid 1 monitor-isid-offset 1000 egress-ports 1/1 egress-mlt
16 map-to-vid 1000
Switch:1(config)# monitor-by-isid 7 monitor-isid-offset 7 egress-mlt 2 map-to-vid 203
Switch:1(config)# monitor-by-isid 2 egress-mlt 3
```

Variable Definitions

Use the data in the following table to use the **monitor-by-isid** command.

Variable	Value
<1–1000>	Specifies the monitoring session.
monitor-isid-offset <1-1000>	Specifies the offset ID that is mapped to the actual monitor I-SID where packets are mirrored. `Monitor I-SID = Base monitor I-SID + Offset ID.` The base monitor I-SID is 16776000.
egress-ports {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}	Specifies the port to which the analyzers connect.
egress-mlt <1–512>	Specifies the MLT to which the analyzers connect.
map-to-vid <1–4093>	Maps the mirrored packet to a specified VLAN ID for analysis. This parameter is optional. **Note:** If you use the Insight port 1/s1 as the analyzer port on the monitoring BEB for remote mirroring, you must associate it to VLAN ID 4091.

Displaying I-SID Monitoring Diagnostics

Use the following procedure to display the **`monitor-by-isid`** entries on the monitoring BEB.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enter the following command:

   ```
   show monitor-by-isid WORD<1-1024>
   ```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)# show monitor-by-isid 3
================================================================================
            Diag Monitor-By-ISID
================================================================================
ID MONITOR_ISID ISID_OFFSET   EGRESS_PORTS   EGRESS_MLT  MAP_TO_VLAN  ENABLE
--------------------------------------------------------------------------------
3    16776000       1           1/4, 1/5         1          999         true
--------------------------------------------------------------------------------
```

Variable Definitions

Use the data in the following table to use the **`show monitor-by-isid`** command.

Variable	Value
WORD<1-1024>	Specifies the session ID list ranging from 1 to 1000.

Displaying I-SID Mirroring Statistics

Use the following procedure to display the statistics of the number of packets mirrored into I-SID on the mirroring BEB.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enter the following command:

   ```
   show isid-mirroring stats [monitor-isid-offset WORD<1-1024>]
   ```

Example

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# show isid-mirroring stats monitor-isid-offset 1
==================================
            Mirror Statistics Info
==================================
ISID        ISID_OFFSET      PACKETS
----------------------------------
16776000        1              100
----------------------------------
```

Variable Definitions

Use the data in the following table to use the **`show isid-mirroring stats`** command.

Variable	Value
monitor-isid-offset WORD<1–1024>	Specifies the offset ID mapped to monitor the I-SID. The offset ID ranges from 1 to 1000.

Clearing Fabric RSPAN (Mirror to I-SID) Statistics

Use the following procedure to clear Fabric RSPAN (Mirror to I-SID) statistics of packets mirrored into the specified mirroring I-SID or all mirroring I-SIDs on the BEB.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   ```

```
configure terminal
```

2. Enter the following command:

```
clear isid-mirroring stats monitor-isid-offset WORD<1-1024>
```

Note:

You must use this command on the Mirroring BEB to clear the statistics of packets mirrored into I-SID.

Example

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

Clear all Fabric RSPAN statistics:

```
Switch:1(config)# clear isid-mirroring stats
```

Clear all Fabric RSPAN (Mirror to I-SID) statistics of packets mirrored into the specified mirroring I-SID

```
Switch:1(config)# clear isid-mirroring stats monitor-isid-offset 1
```

Variable Definitions

Use the data in the following table to use the **`clear isid-mirroring stats`** command.

Variable	Value
monitor-isid-offset *WORD<1–1024>*	Specifies the offset ID that is mapped to the actual monitor I-SID where packets are mirrored. `Monitor I-SID = base monitor I-SID + offset ID.` The range of the *monitor-isid-offset* is 1 to 1000. The base monitor I-SID is 16776000.

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, expand the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Click the **Advanced** tab.
4. In the **Vlan Operation Action** box for the VLAN you want to flush, double-click, and then select a flush option from the list.

   In a VLAN context, 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. On the Device Physical View, select a port.
2. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **Port**.
3. Click **General**.
4. Click the Interface tab.
5. In the **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.

Procedure

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

   Note:

   - When you configure `tx` mode port mirroring on T-UNI and SPBM NNI ports, unknown unicast, broadcast and multicast traffic packets that ingress these ports appear on the mirror destination port, although they do not egress the mirror source port. This is because `tx` mode port mirroring happens on the mirror source port *before* the source port squelching logic drops the packets at the egress port.
   - The available four mirroring resources are shared between Fabric RSPAN and regular port mirroring, and are allocated based on the mode configured, Ingress (`rx`) or Egress (`tx`). Each configured mode occupies one mirroring resource, but when you configure the mode as `both`, it occupies two mirroring resources (one for Rx and one for Tx).
   - Do not configure the source of mirrored traffic (mirroring to an I-SID) and the analyzer (monitoring an I-SID) on the same local device with the same I-SID offset. If you require mirroring and monitoring on the same local device, use standard port-based mirroring instead of Fabric RSPAN. Fabric RSPAN mirrors traffic into an I-SID of the SPB Fabric network and monitors traffic on the remote device; the network analyzer resides on the remote monitoring device and not on the same local device.

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.

Table continues…

Name	Description
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). Configures the mirroring port.
MirroringMltId	Specifies the destination MultiLink trunking ID.
MonitoringIsidOffset	Used to configure the monitoring I-SID offset value. The offset ID is mapped to the actual monitor I-SID value to which the packets are mirrored.
MonitoringIsid	Specifies the actual monitor I-SID value to which the packets are mirrored.
MirroringQos	Used to define the Quality of Service (QoS) profiles for the mirrored packet into monitoring I-SID.

Configuring ACLs for Mirroring

Configure the access control list (ACL) to mirror packets for an access control entry (ACE) that matches a particular packet.

Before you begin

- The ACL exists.

About this task

To modify an ACL parameter, double-click the parameter you wish to change. Change the value, and then click Apply. You cannot change a parameter that appears dimmed; in this case, delete the ACL, and then configure a new one.

Procedure

1. In the navigation pane, expand the **Configuration** > **Security** > **Data Path** folders.
2. Click **Advanced Filters (ACE/ACLs)**.
3. Click the **ACL** tab.
4. Double-click the parameter**MirrorMltId** to configure mirroring to a destination MLT group.
5. Double-click the parameter **MirrorDstPortList** to configure mirroring to a destination port or ports.

ACL Field Descriptions

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

Name	Description
AclId	Specifies a unique identifier for the ACL from 1–2048.
Type	Specifies whether the ACL is VLAN- or port-based. Valid options are • inVlan • inPort • outPort **Important:** The inVlan ACLs drop packets if you add a VLAN after ACE creation.
Name	Specifies a descriptive user-defined name for the ACL.
VlanList	For inVlan type, specifies all VLANs to associate with the ACL.
PortList	For inPort and outPort ACL types, specifies the ports to associate with the ACL.
DefaultAction	Specifies the action taken when no ACEs in the ACL match. Valid options are deny and permit. Deny means the system drops the packets; permit means the system forwards packets. The default is permit.
ControlPktAction	Specifies the action for control packets, if you configure DefaultAction to deny. If DefaultAction is permit, this value is ignored.
State	Enables or disables all of the ACEs in the ACL. The default value is enable.
PktType	Indicates the packet type that this ACL is applicable to. The default is IPv4.
MirrorMltId	Configures mirroring to a destination MLT group.
MirrorDstPortList	Configures mirroring to a destination port or ports.

Configuring ACEs for Mirroring

Before you begin

- The ACL exists.
- The ACE exists.

About this task

Configure actions to use filters for flow mirroring. Use an ACE to define the mirroring actions the filter performs.

If you use the mirror action, ensure that you specify the mirroring destination: IP address, MLTs, ports, or VLANs.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Security** > **Data Path**.

2. Click **Advanced Filters (ACE/ACLs)**.
3. Click the **ACL** tab.
4. Select the ACL for which to modify an ACE.
5. Click **ACE**.
6. Select an ACE, and then click **Action**.
7. Configure one of: **DstPortList**, **DstMltId**, or **DstIp**.
8. Click **Apply**.

Action Field Descriptions

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

Note:

The table lists the options for both Security ACEs and QoS ACEs. Dependent upon the ACE different options appear on the EDM interface.

Name	Description
AclId	Specifies the ACL ID.
AceId	Specifies a unique identifier and priority for the ACE.
Mode	Indicates the operating mode associated with this ACE. Valid options are deny, permit and none. The default is none.
RemarkDscp	Specifies the new Per-Hop Behavior (PHB) for matching packets: phbcs0, phbcs1, phbaf11, phbaf12, phbaf13, phbcs2, phbaf21, phbaf22, phbaf23, phbcs3, phbaf31, phbaf32, phbaf33, phbcs4, phbaf41, phbaf42, phbaf43, phbcs5, phbef, phbcs6, phbcs7. This action is a QoS action. The ACE ID must be in the range of 1001–2000.
RemarkDot1Priority	Specifies the new 802.1 priority bit for matching packets: zero, one, two, three, four, five, six, or seven. This action is a QoS action. The ACE ID must be in the range of 1001–2000. The default is disable.
InternalQoS	This variable is a QoS action. The ACE ID must be in the range of 1001–2000. The default value is 1.
RedirectNextHop	Redirects matching IP traffic to the next hop. The default is 0.0.0.0.
RedirectUnreach	Configures the desired behavior for redirected traffic when the specified next-hop is not reachable. The default value is deny.
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 switch. Use this parameter with either a security or QoS ACE. The default is disabled.

Table continues…

Name	Description
DstPortList	Specifies the ports 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.

Running a Ping Test

About this task

Use ping to determine if an entity is reachable.

Note:

Troubleshooting using ping and traceroute (including Layer 2 ping and Layer 2 traceroute) is not supported on EDM. For more information, see Release Notes for VOSS. As an alternative, use CLI.

Procedure

1. From the navigation tree, expand the following folders: **Configuration** > **Edit** > **Diagnostics**.
2. Click **Ping/Trace Route**.
3. Click the **Ping Control** tab.
4. Click **Insert**.
5. In the **OwnerIndex** box, type the owner index.
6. In the **TestName** box, type the name of the test.
7. In the **TargetAddress** box, type the host IP address.
8. From the **AdminStatus** options, select **enabled**.
9. In the remainder of the option boxes, type the desired values.
10. Click **Insert**.
11. Select 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.
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 seconds.
ProbeCount	Specifies the number of times to perform a ping operation at a remote host. The default is 1.
AdminStatus	Specifies the state of the ping control entry: enabled or disabled. The default is disabled.
DataFill	Determines the data portion of a probe packet.
Frequency	Specifies the number of seconds to wait before repeating a ping test. The default is 0.
MaxRows	Specifies the maximum number of entries allowed in the PingProbeHistory table. The default is 50.
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. The value of this object defaults to zero, indicating that none of the above options have been selected.
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.
Descr	Describes the remote ping test. The default is 0x00.
SourceAddress	Specifies the IP address (a.b.c.d) as the source address in outgoing probe packets.
ByPassRouteTable	Enables (optionally) the bypassing of the route table. The default is disabled.

Viewing Ping Results

About this task

View ping results to view performance-related data.

Note:

Troubleshooting using ping and traceroute (including Layer 2 ping and Layer 2 traceroute) is not supported on EDM. For more information, see Release Notes for VOSS. As an alternative, use CLI.

Procedure

1. From the navigation tree, expand the following folders: **Configuration** > **Edit** > **Diagnostics**.
2. Click **Ping/Trace Route**.
3. Click the **Ping Control** tab.
4. Select a ping test entry.
5. Click **Ping Result**.

Ping Result Field Descriptions

Use the data in the following table to use the **Ping Result** tab.

Name	Description
OwnerIndex	Specifies the ping test owner.
TestName	Specifies the test name.
OperStatus	Indicates the operational status of the test. The default is disabled.
IpTargetAddressType	Specifies the IP address type of the target. The default is unknown.
IpTargetAddress	Specifies the IP address of the target.
MinRtt	Specifies the minimum ping round-trip-time (RTT) received. A value of 0 means that no RTT is received.
MaxRtt	Specifies the maximum ping RTT received. A value of 0 means that no RTT is received.
AverageRtt	Specifies the current average ping RTT.
ProbeResponses	Specifies the number of responses to probes.
SentProbes	Specifies the number of sent probes.
RttSumOfSquares	Specifies the sum of squares of RTT for all probes received.
LastGoodProbe	Specifies the date and time when the last response is received for a probe.

Viewing Ping Probe History

About this task

View the ping probe history to view the history of ping tests performed by the switch.

Note:

Troubleshooting using ping and traceroute (including Layer 2 ping and Layer 2 traceroute) is not supported on EDM. For more information, see Release Notes for VOSS. As an alternative, use CLI.

Procedure

1. From the navigation tree, expand the following folders: **Configuration** > **Edit** > **Diagnostics**.
2. Click **Ping/Trace Route**.
3. Select a ping entry.
4. Click **Ping Probe History**.

Ping Probe History Field Descriptions

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

Name	Description
OwnerIndex	Specifies the owner index
TestName	Indicates the name given to the test.
Index	Specifies the index number.
Response	Indicates the amount of time, measured in milliseconds, between request (probe) and response, or when the request timed out. Response is reported as 0 when it is not possible to transmit a probe.
Status	Indicates the status of the response; the result of a particular probe done by a remote host.
LastRC	Indicates the last implementation-method-specific reply code (RC) received. If ICMP Echo is used, then a successful probe ends when an ICMP response is received that contains the code ICMP_ECHOREPLY(0).
Time	Indicates the timestamp for this probe result.

Running a Traceroute Test

About this task

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

Note:

Troubleshooting using ping and traceroute (including Layer 2 ping and Layer 2 traceroute) is not supported on EDM. For more information, see Release Notes for VOSS. As an alternative, use CLI.

Procedure

1. From the navigation tree, expand the following folders: **Configuration** > **Edit** > **Diagnostics**.
2. Click **Ping/Trace Route**.
3. Click the **Trace Route Control** tab.
4. Click **Insert**.
5. Configure the instance as required.
6. Click **Insert**.
7. Select an entry, and then click **Start**.

   Let the test run for several seconds.
8. Select an entry, and then click **Stop**.
9. View the traceroute test results.

Trace Route Control Field Descriptions

Use the data in the following table to use the Trace Route Control tab.

Name	Description
OwnerIndex	Provides access control by a security administrator using the VACM for tables in which multiple users need to independently create or modify entries.
TestName	Specifies the name of the traceroute test.
TargetAddressType	Specifies the type of host address to use on the traceroute request at the remote host. The default is IPv4.
TargetAddress	Specifies the host address used on the traceroute request at the remote host.
ByPassRouteTable	Enables bypassing of the route table. If you enable this variable, the remote host bypasses the normal routing tables and sends directly to a host on an attached network. If the host is not on a directly-attached network, an error is returned. You can use this variable to perform the traceroute operation to a local host through an interface that has no route defined. The default is disabled.
DataSize	Specifies the size of the data portion of a traceroute request in octets. The default is 1.
TimeOut	Specifies the timeout value, in seconds, for a traceroute request. The default is 3.

Table continues…

Name	Description
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. **Note:** If the traceroute action is directed to an IPv6 host address, Linux increments the UDP port on a per-TTL basis. For an IPv4 host address, Linux increments the UDP port on a per-probe basis. Because the **traceroute** command sends a default of three probes, three incrementing ports will be sent for an IPv4 host address. If you use the -p parameter with a value greater than 65533, the **traceroute** command fails for an IPv4 host address because the maximum port number, 65535, is exceeded.
MaxTtl	Specifies the maximum time-to-live from 1–255. The default is 30.
DSField	Specifies the value to store in the Differentiated Services (DS) field in the IP packet used to encapsulate the traceroute probe. The default is 0.
SourceAddressType	Specifies the type of the source address to use 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. The default is 0.
MiscOptions	Enables an application to specify implementation-dependent options.
MaxFailures	Indicates the maximum number of consecutive timeouts allowed before terminating a remote traceroute request. The default is 5.
DontFragment	Enables setting of the do not fragment (DF) flag in the IP header for a probe. The default is disabled.
InitialTtl	Specifies the initial time-to-live (TTL) value to use. The default is 1.
Frequency	Specifies the number of seconds to wait before repeating a traceroute test as defined by the value of the various objects in the corresponding row. The default is 0.
StorageType	Specifies the storage type for this row.
AdminStatus	Specifies the desired state for TraceRouteCtlEntry. The options are enabled or disabled. The default is disabled.
MaxRows	Specifies the maximum number of entries allowed in the TraceRouteProbeHistoryTable. The default is 50.

Table continues…

Name	Description
TrapGeneration	Determines when to generate a notification for this entry. The options are • **pathChange** —Generates a TraceRoutePathChange notification after the current path varies from a previously determined path. • **testFailure** —Generates a TraceRouteTestFailed notification after the full path to a target cannot be determined. • **testCompletion** —Generates a TraceRouteTestCompleted notification after the path to a target has been determined.
Descr	Describes the remote traceroute test.
CreateHopsEntries	Stores the current path for a traceroute test in the TraceRouteHopsTable on an individual hop basis when the value of this object is true. The default is false.
Type	Reports or selects the implementation method to use for performing a traceroute operation.

Viewing Traceroute Results

About this task

View traceroute results to view performance-related data.

Note:

Troubleshooting using ping and traceroute (including Layer 2 ping and Layer 2 traceroute) is not supported on EDM. For more information, see Release Notes for VOSS. As an alternative, use CLI.

Procedure

1. From the navigation tree, expand the following folders: **Configuration** > **Edit** > **Diagnostics**.
2. Click **Ping/Trace Route**.
3. Click the **Trace Route Control** tab.
4. Select a traceroute entry.
5. Click **Trace Route Result**.

Trace Route Result Field Descriptions

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

Name	Description
OwnerIndex	Specifies the index of the owner.
TestName	Specifies the name of the test.

Table continues…

Name	Description
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.

Note:

Troubleshooting using ping and traceroute (including Layer 2 ping and Layer 2 traceroute) is not supported on EDM. For more information, see Release Notes for VOSS. As an alternative, use CLI.

Procedure

1. From the navigation tree, expand the following folders: **Configuration** > **Edit** > **Diagnostics**.
2. Click **Ping/Trace Route**.
3. Click the **Trace Route Control** tab.
4. Select an entry.
5. Click **Trace Route Probe History**.

Route Probe History Field Descriptions

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

Name	Description
OwnerIndex	Identifies the Trace Route entry to which a probe result belongs.
TestName	Specifies the test name.
Index	Specifies the Index.
HopIndex	Indicates for which hop in a traceroute path the probe results are intended.

Table continues…

Name	Description
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.

Configuring I-SID Monitoring

Use the following procedure to configure Fabric RSPAN on the Backbone Edge Bridge (BEB) connected to the monitor station.

Note:

Do not configure the source of mirrored traffic (mirroring to an I-SID) and the analyzer (monitoring an I-SID) on the same local device with the same I-SID offset. If you require mirroring and monitoring on the same local device, use standard port-based mirroring instead of Fabric RSPAN. Fabric RSPAN mirrors traffic into an I-SID of the SPB Fabric network and monitors traffic on the remote device; the network analyzer resides on the remote monitoring device and not on the same local device.

Note:

If you change the egress port or egress MLT for a particular session using a CLI command, it overwrites the existing egress port list or egress MLT.

Procedure

1. In the navigation pane, expand the **Configuration** > **Edit** > **Diagnostics** folders.
2. Click **General**.
3. Click the **Monitor-By-ISID** 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**.

Monitor-By-ISID Field Descriptions

Use the data in the following table to use the Monitor-By-ISID tab.

Name	Description
Index	Specifies the entry that contains monitor by I-SID information.
MonitorIsidOffset	Configures the monitoring I-SID offset value. The offset ID is mapped to the actual monitor I-SID value to which the packets are mirrored.
MonitorIsid	Specifies the actual monitor I-SID value to which packets are mirrored.
EgressPortList	Specifies the egress ports to which traffic analyzers connect.
EgressMltId	Specifies the egress MLT ID to which traffic analyzers connect.
MapToVlanId	Specifies the VLAN ID to map with mirrored traffic on the monitoring node.
Enable	Enables or disables monitoring by I-SID.

Viewing and Clearing Fabric RSPAN (Mirror to I-SID) Statistics

Use the following procedure to view or clear statistics of the number of packets mirrored into I-SID on the mirroring BEB.

Procedure

1. In the navigation pane, expand the **Configuration** > **Edit** > **Diagnostics** folders.
2. Click **General**.
3. Click the **Isid-Mirroring Stats** tab.

Isid-Mirroring Stats Field Descriptions

Use the data in the following table to use the Isid-Mirroring Stats tab.

Name	Description
Index	Specifies the entry that contains Fabric RSPAN statistics information.
MonitorIsid	Specifies the actual monitor I-SID value to which the packets are mirrored.
MirroredPackets	Specifies the number of packets mirrored into I-SID on the mirroring BEB.
ClearStats	Clears the Fabric RSPAN statistics.

Chapter 7: General Troubleshooting

Use the information in this chapter to learn about the troubleshooting guidelines for hardware and the switch software.

This chapter includes the following sections:

- Hardware Troubleshooting
- Software Troubleshooting

Hardware Troubleshooting

The following sections provide troubleshooting information for common hardware problems.

Using Trace to Diagnose Hardware Problems

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

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Begin the trace operation:

   ```
   line-card 1 trace level [<Module_ID>]{<0-4>}
   ```

3. Search the trace for a specific string value:

   ```
   line-card 1 trace grep {WORD<0-1024>}
   ```

Example

```
Switch:1>enable
```

Begin the trace operation:

```
Switch:1#line-card 1 trace level 67 1
```

Search the trace for a specific string value:

```
Switch:1#line-card 1 trace grep 00-1A-4B-8A-FB-6B
```

Variable Definitions

Use the data in the following table to use the **`line-card 1`** command.

Variable	Value
<Module_ID> {*<0–4>*}	Starts the trace by specifying the module ID and level. *<Module_ID>* specifies the module for the trace. Different hardware platforms support different ID ranges because of feature support differences. To see which module IDs are available on the switch, use the **`show trace modid-list`** command or CLI command completion Help. *<0-4>* specifies the trace level from 0–4, where 0 is disabled; 1 is very terse; 2 is terse; 3 is verbose, 4 is very verbose.
WORD<0–1024>	Performs a string search in the trace.

Troubleshooting USB Viewing Problems

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

Note:

Not all hardware platforms can use the USB device for additional file storage. Some platforms use the USB as part of the system operation and must never be removed. For more information, see your hardware documentation.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Check the file system:

   ```
   ls /usb/
   ```

3. Remove a USB device:
   a. Unmount the USB device:

      ```
      usb-stop
      ```

   b. Wait for the response that indicates it is safe to remove the device.
   c. Physically remove the device.
4. Remove and then reinsert the device.
5. Check the device for errors:

```
dos-chkdsk /usb
```

Run the **`dos-chkdsk /usb repair`** command, if at the end of the **`dos-chkdsk /usb`** command output you see:

1) Correct

2) Don’t correct

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

```
dos-format /usb
```

⚠ Caution:

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

Example

Check the file system:

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

Check the device for errors:

```
Switch:1#usb-stop
It is now safe to remove the USB device.
Switch:1#dos-chkdsk /usb
/usr/sbin/fsck.vfat /dev/usb1 -v >& /dev/console dosfsck 2.11a
(05 Mar 2010)
dosfsck 2.11a, 05 Mar 2010, FAT32, LFN
Checking we can access the last sector of the filesystem
Boot sector contents:
System ID "mkdosfs"
Media byte 0xf8 (hard disk)
512 bytes per logical sector
4096 bytes per cluster
32 reserved sectors
First FAT starts at byte 16384 (sector 32)
2 FATs, 32 bit entries
3897344 bytes per FAT (= 7612 sectors)
Root directory start at cluster 2 (arbitrary size)
Data area starts at byte 7811072 (sector 15256)
974240 data clusters (3990487040 bytes)
62 sectors/track, 124 heads
0 hidden sectors
7809178 sectors total
Checking for unused clusters.
Checking free cluster summary.
/dev/usb1: 17 files, 174804/974240 clusters
```

If errors are detected, reformat the disk:

```
Switch:1#dos-format /usb
```

Software Troubleshooting

This section contains general troubleshooting information for the switch software.

Failure to Read Failed 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

```
Switch:1> enable
Switch:1# configure terminal
Switch: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.

CPU Queues

This section provides information about CPU traffic and queues for different hardware platforms.

Note:

To view the CPU traffic and queue information for the VSP 8600 Series use the command `show qos cosq-stats cpu-port`.

The following table outlines the CPU traffic, and which queue it uses, for the VSP 4000 Series.

Queue description	Queue ID	Queue pps, burst size	Traffic type
Other	**0**	125 pps, 1500 kbps	Data Exception
BC_Other	**1**	4000 pps, 1400 kbps	Broadcast

Table continues…

Queue description	Queue ID	Queue pps, burst size	Traffic type
HI_CPU looper	**2**	4000 pps, 1000 kbps	RIP
HI_CPU	**3**	4000 pps, 1000 kbps	PIM Unicast, IGMP
HI_CPU	**4**	4000 pps, 1000 kbps	OSPF Unicast, BGP
HI_CPU	**5**	4000 pps, 1000 kbps	Multicast Data
HI_CPU	**6**	4000 pps, 12000 kbps	ISIS
HI_CPU	**7**	4000 pps, 6000 kbps	OSPF Multicast, PIM Multicast
mgmt	**8**	4000 pps, 2000 kbps	Telnet, RSH, SSH, RLOGIN, DNS, HTTP, TFTP, FTP, SNMP
Low rate looper	**9**	4000 pps, 5000 kbps	ICMP Broacast, DHCP, ARP, ND Multicast and Unicast
Low rate 1	**10**	4000 pps, 1000 kbps	ISIS Hello, RADIUS, EAP, NTP, LLDP, ICMP Unicast, IST Control
Low latency	**11**	250 pps, 20 kbps	CFM, VRRP
Low latency	**12**	250 pps, 20 kbps	LACP
Low latency	**13**	250 pps, 20 kbps	BPDU, SLPP
Low latency	**14**	300 pps, 1000 kpbs	ISIS Hello
Low latency	**15**	250 pps, 20 kbps	VLACP

The following table outlines the CPU traffic, and which queue it uses, for the VSP 4900 Series, VSP 7200 Series, VSP 7400 Series, and VSP 8000 Series.

Queue description	Queue ID	Queue pps, burst size	Traffic type
Other	**0**	1400 pps, 100 kbps	Others
BC_Other	**1**	1400 pps, 100 kbps	Broadcast
HI_CPU looper	**2**	4000 pps, 100 kbps	IP_COS01
HI_CPU	**3**	4000 pps, 100 kbps	IP_COS23, Multicast Data
HI_CPU	**4**	4000 pps, 100 kbps	IP_COS4
HI_CPU	**5**	4000 pps, 100 kbps	IP_COS5, IGMP, PIM Unicast
HI_CPU	**6**	4000 pps, 100 kbps	ARP, RARP, ND Multicast and Unicast, MLD
HI_CPU	**7**	4000 pps, 100 kbps	IP_COS6, Telnet, SSH
mgmt	**8**	4000 pps, 200 kbps	IP_COS7, OSPF

Table continues…

Queue description	Queue ID	Queue pps, burst size	Traffic type
Low rate looper	**9**	4000 pps, 6000 kbps	OSPF Multicast, PIM Multicast, RIP_v1, RIP_v2
Low rate 1	**10**	4000 pps,12000 kbps	ISIS
Low latency	**11**	4000 pps, 1000 kbps	ISIS Hello, LLDP, EAP, IST Control
Low latency	**12**	500 pps, 32 kbps	VRRP, CFM
Low latency	**13**	256 pps, 32 kbps	SLPP
Low latency	**14**	100 pps, 32 kbps	BPDU
Low latency	**15**	250 pps, 32 kbps	LACP, VLACP, TUNI-extract

The XA1400 Series uses filters to rate limit packets to the CPU.

Packet Type	Rate (PPS)	Packet Size (Bytes)
CFM, VLACP, RIPv2, SONMP, RIPv1, VRRP, ARP.	100	64
DHCP	50	1400
ISIS, OSPF	100	1400
Other Packets	Not rate limited	Not Applicable

Chapter 8: Layer 1 Troubleshooting

Use the information in this chapter to 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 transceivers, see *Extreme Networks Pluggable Transceivers Installation Guide*.

Important:

Extreme Networks recommends using transceivers documented in *Extreme Networks Pluggable Transceivers Installation Guide*, as they have been through extensive qualification and testing. Extreme Networks is not responsible for issues related to third party transceivers.

Procedure

1. Measure the 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.

Resetting a QSFP+ or QSFP28 Transceiver

Reset a transceiver to simulate removal and reinsertion of the transceiver, which can be helpful in troubleshooting. For example, if authentication of the transceiver fails but you believe the transceiver is a qualified Extreme Networks part, you can reset the transceiver to begin the authentication process again.

About this task

Resetting the transceiver stops traffic and triggers log messages similar to the removal and insertion of the transceiver.

Before you begin

- Before you use the **`pluggable-optical-module reset`** command, ensure the port is administratively down to avoid link flaps.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Reset the transceiver:

   ```
   pluggable-optical-module reset {slot/port[/sub-port]}
   ```

Important:

Not all hardware platforms support these port types. For more information, see your hardware documentation.

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#pluggable-optical-module reset 1/41
Switch:1(config)#
CP1 [06/25/14 22:15:09.644] 0x0000c5e7 00300001.232 DYNAMIC SET GlobalRouter HW INFO Link
Down(1/41)
CP1  [06/25/14 22:15:10.267] 0x000e0597 00000000 GlobalRouter HAL INFO GBIC removed from
slot 1 Port 41 Type:40GbSR4 Vendor:Extreme Networks
CP1  [06/25/14 22:15:13.015] 0x000e0598 00000000 GlobalRouter HAL INFO GBIC inserted in
slot 1 Port 41 Type:40GbSR4 Vendor:Extreme Networks
CP1 [06/25/14 22:15:14.562] 0x0000c5ec 00300001.232 DYNAMIC CLEAR GlobalRouter HW INFO
Link Up(1/41)
```

```
Switch:1(config)#pluggable-optical-module reset 1/1
Switch:1(config)#CP1  [03/31/16 10:48:24.492:UTC] 0x0000c5e7 00300001.384 DYNAMIC SET
GlobalRouter HW INFO Link Down(1/1)
CP1  [03/31/16 10:48:24.601:UTC] 0x000e0597 00000000 GlobalRouter HAL INFO GBIC removed
from slot 1 Port 1 Type:100GbCR4 Vendor:Extreme Networks
CP1  [03/31/16 10:48:24.710:UTC] 0x0000c5e7 00300001.385 DYNAMIC SET GlobalRouter HW INFO
Link Down(1/2)
CP1  [03/31/16 10:48:26.668:UTC] 0x000e0598 00000000 GlobalRouter HAL INFO GBIC inserted
in slot 1 Port 1 Type:100GbCR4 Vendor:Extreme Networks
CP1  [03/31/16 10:48:26.988:UTC] 0x0000c5ec 00300001.385 DYNAMIC CLEAR GlobalRouter HW
INFO Link Up(1/2)
CP1  [03/31/16 10:48:27.099:UTC] 0x0000c5ec 00300001.384 DYNAMIC CLEAR GlobalRouter HW
INFO Link Up(1/1)
```

Variable Definitions

Use the data in the following table to use the **`pluggable-optical-module reset`** command.

Variable	Value
{slot/port[/sub-port]}	Specifies location of the transceiver to reset. Identifies a single slot and port. If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Chapter 9: Layer 2 and 3 Troubleshooting

Use the information in this chapter to learn about Layer 2 and Layer 3 troubleshooting.

This chapter includes the following sections:

- Troubleshooting BPDU Guard
- Troubleshooting IPv6 VRRP
- Troubleshooting RSMLT
- Troubleshooting vIST Failure
- Troubleshooting Transparent Port UNI
- Troubleshooting Multicast feature
- Troubleshooting MACsec
- Troubleshooting Fabric Attach

Troubleshooting BPDU Guard

The following procedures provide information to troubleshoot issues with Bridge Protocol Data Unit (BPDU) Guard.

No Packets Received on the Port

For BPDU Guard to work on a port, the port must receive BPDU packets. Perform the following procedure to troubleshoot cases when the port does not receive packets.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Show the BPDU Guard status for the port:

   ```
   show spanning-tree bpduguard {slot/port[/sub-port][-slot/port[/sub-port]][,...]}
   ```

3. Use the following command to verify that the port receives packets:

```
show interface gigabitEthernet statistics verbose {slot/port[/sub-
port][-slot/port[/sub-port]][,...]}
```

4. Verify that the remote port is sending packets:

```
show spanning-tree {mstp|rstp} port role [{slot/port[/sub-port][-
slot/port[/sub-port]][,...]}]
```

```
show spanning-tree {mstp|rstp} port statistics [{slot/port[/sub-
port][-slot/port[/sub-port]][,...]}]
```

Example

Port 1/8 receives packets. The remote port is disabled and does not send BPDU packets.

The following example shows that BPDU Guard is enabled for port 1/8. The BPDU Guard administrative state for the port is enabled but the timer counter is 0.

```
Switch:1>enable
Switch:1#show spanning-tree bpduguard 1/8
================================================================
                                  Bpdu Guard
================================================================
Port      PORT         PORT               TIMER   BPDUGUARD
NUM MLTID ADMIN_STATE  OPER_STATE TIMEOUT  COUNT   ADMIN_STATE
----------------------------------------------------------------
1/8        Up           Up         120      0       Enabled
Switch:1#show interface gigabitEthernet statistics verbose 1/8
================================================================================================
=======
                                          Port Stats Interface Extended
================================================================================================
=======
PORT_NUM IN_UNICST  OUT_UNICST IN_MULTICST  OUT_MULTICST IN_BRDCST  OUT_BRDCST   IN_LSM
OUT_LSM
------------------------------------------------------------------------------------------------
-------
1/8    201          0          160062       60943        4          72
0         0
Switch:1#show spanning-tree mstp port role 1/8
================================================================================
                           CIST Port Roles and States
================================================================================
Port-Index  Port-Role    Port-State   PortSTPStatus  PortOperStatus
--------------------------------------------------------------------------------
1/8       Disabled     Forwarding   Disabled       Disabled
Switch:1#show spanning-tree mstp port statistics 1/8
================================================================================
                           MSTP Cist Port Statistics
================================================================================
Port Number                       : 1/8
Cist Port Fwd Transitions         : 0
Cist Port Rx MST BPDUs Count      : 0
Cist Port Rx RST BPDUs Count      : 0
Cist Port Rx Config BPDUs Count   : 0
Cist Port Rx TCN BPDUs Count      : 0
Cist Port Tx MST BPDUs Count      : 0
Cist Port Tx RST BPDUs Count      : 0
Cist Port Tx Config BPDUs Count   : 0
Cist Port Tx TCN BPDUs Count      : 0
Cist Port Invalid MSTP BPDUs Rx   : 0
Cist Port Invalid RST BPDUs Rx    : 0
Cist Port Invalid Config BPDUs Rx  : 0
```

```
Cist Port Invalid TCN BPDUs Rx      : 0
Cist Port Proto Migr Count          : 0
```

Variable Definitions

Use the data in the following table to use the **show spanning-tree bpduguard** command.

Variable	Value
{slot/port[/sub-port][-slot/port[/sub-port]][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Use the data in the following table to use the **show interface gigabitEthernet statistics verbose** command.

Variable	Value
{slot/port[/sub-port][-slot/port[/sub-port]][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Use the data in the following table to use the **show spanning-tree** command.

Variable	Value
{mstp|rstp}	Specifies the spanning tree protocol.
{slot/port[/sub-port][-slot/port[/sub-port]][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Troubleshooting IPv6 VRRP

The following sections describe troubleshooting information for IPv6 Virtual Router Redundancy Protocol (VRRP).

VRRP Transitions

When a VRRP transition takes place with the backup taking over as the master, look for the following message in the syslog on the new master, as well as the old master. This message provides information to allow you to determine the cause of the transition.

```
IPv6 Vrrp State Transition Trap(Port/Vlan=200, Type=masterToInitialize,
Cause=shutdownReceived, VrId=20,VrIpAddr=fe80:0:0:0:0:0:0:200,
Addr=fe80:0:0:0:224:7fff:fe9d:1a03)
```

In this message, see the Type and Cause fields.

Note:

Although all of the possible causes and types are listed below, not all of the listed causes and types appear in the trap/log message.

The following table describes the VRRP transition types.

Table 16: 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 17: 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

Table continues…

Cause value	Cause definition
7	Critical IP goes down
8	User disabling VRRP
9	VRRP status synced from primary
10	IPv6 interface on which VRRP is configured goes down
11	Lower priority advertisement received
12	Advertisement received from higher interface IP address with equal priority
13	Advertisement received from lower interface IP address with equal priority
14	User enabled VRRP
15	Transition because of any other cause

Enabling Trace Messages for IPv6 VRRP Troubleshooting

Use this procedure to enable trace messages for IPv6 VRRP.

When VRRP is enabled on two routing switches, the master-backup relationship forms with one router taking the responsibility of routing. If the master-backup relationship is not formed between the VRRP virtual routers, look for the following trace messages to ensure that the master is sending the advertisements correctly and the backup is processing them.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. To troubleshoot IPv6 VRRP, you can enable RCIP6 trace messages with the command:

   ```
   trace level 66 3
   ```

3. And to provide additional trace information, you can also enable the following traces:

   ```
   trace ipv6 nd enable
   trace ipv6 base enable all
   trace ipv6 forwarding enable all
   trace ipv6 rtm enable all
   trace ipv6 transport enable all
   ```

4. When VRRP is enabled on two routing switches, the master-backup relationship forms with one router taking the responsibility of routing. If the master-backup relationship is not formed between the VRRP virtual routers, look for the following trace messages to ensure that the master is sending the advertisements correctly and the backup is processing them. On the master router, look for the following RCIP6 trace messages.

   - `tMainTask RCIP6: rcip6_vrrp.c: 5118: VRF name: GlobalRouter (VRF id 0): ipv6VrrpTic: Am Master for Vrid 200 on IfIndex 2053 Timer 1`

If VRRP is enabled on the interface, this timer kicks off every second and shows the state for the VRID.

- `[11/18/09 15:08:20:383] tMainTask RCIP6: rcip6_vrrp.c: 5924: ipv6VrrpSendAdvertisement: for Vrid 200 on IfIndex 2053`

  `[11/18/09 15:08:20:583] tMainTask RCIP6: rcip6_vrrp.c: 5175: VRF name: GlobalRouter (VRF id 0): ipv6VrrpTic: ipv6VrrpSendAdvertisement`

The preceding trace messages show that the VRRP master is sending the advertisements correctly at the end of advertisement interval for a VRID.

5. On the backup router, look for the following RCIP6 trace messages.

   - `tMainTask RCIP6: rcip6_vrrp.c: 5236: VRF name: GlobalRouter (VRF id 0): ipv6VrrpTic: Am Backup for VrId 200 on IfIndex 2052 Timer 1`
   - `tMainTask RCIP6: rcip6_vrrp.c: 4854: ipv6VrrpIn: Vrid 200 on IfIndex 2052`
   - `tMainTask RCIP6: rcip6_vrrp.c: 5545: VRF name: GlobalRouter (VRF id 0): rcIpVrrpProcessAdvt: Am backup for Vrid 200 on IfIndex 2052`

   The preceding trace messages show that the backup router is receiving the advertisements sent by the master and correctly processing them.

Risks Associated with Enabling Trace Messages

When traces are enabled on VRRP master, VrrpTic messages are logged for every second and any other configured traces keep displaying, so there is no guarantee that the backup will receive the advertisement from the master within 3 seconds, so it can transit to master also. There is also the risk of toggling of VRRP states (from backup to master and back again).

Enable the limited traces based on whichever is required.

VRRP with Higher Priority Running as Backup

The VRRP router with the higher priority can display as the backup for the following reasons

- Hold-down timer is running.
- The configured Critical IP is not reachable or does not exist.

If the critical-IP is configured for VRRP master, and the critical interface goes down or is deleted, the master transitions to the backup state. In this case, the log shows the transition cause as 1 like many other cases.

If the holddown timer is configured for VRRP master, the holddown timer delays the preemption, giving the device, which is becoming the master enough time to construct routing tables.

Procedure

1. To determine that the issue is with the critical interface, look for the following trace message.

   ```
   tMainTask RCIP6: rcip6_vrrp.c: 5152: VRF name: GlobalRouter (VRF id
   0): ipv6VrrpTic: Becoming backup for Vrid 200 on IfIndex 2052
   because of invalid critical IP
   ```

2. If the holddown Timer is configured for VRRP master, the holddown timer delays the preemption, giving the device, which is becoming the master enough time to construct routing tables.

   ```
   tMainTask RCIP6: rcip6_vrrp.c: Enter in HoldDown processing,Vrid 200
   LastRecvd 0 MasterDown 3, Holddown time remaining 970, Holddownstate
   2
   ```

Troubleshooting RSMLT

The following sections provide information for troubleshooting IPv4 Split Multi-Link Trunking (RSMLT).

RSMLT Peers Not Up

If, after a series of reconfigurations, RSMLT peers do not transition to the up state, use the following procedure to troubleshoot the issue. You can observe this issue on dual-stack VLANs after multiple delete and re-adds of IPv4 interfaces.

Procedure

1. Display the RSMLT configuration. This command shows whether the peers are up:

   ```
   show ip rsmlt peer
   ```

2. Enter VLAN Interface Configuration mode:

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

3. To recover the peers if they are down, disable and reenable RSMLT on both IST peers:

   ```
   no ip rsmlt
   ip rsmlt
   ```

4. If the problem persists, boot from a saved configuration.

Example

Display the RSMLT configuration:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface vlan 1
Switch:1(config-if)#show ip rsmlt peer

==============================================================================

                        Ip Rsmlt Peer Info - GlobalRouter
==============================================================================


VID    IP              MAC               ADMIN    OPER   HDTMR   HUTMR
------------------------------------------------------------------------------
1      192.0.2.1     00:1f:ca:1e:d3:1e       Enable      Up      60   180
2      198.51.100.1  00:1b:ca:1d:e3:1d       Enable      Up      60   180


VID    HDT REMAIN   HUT REMAIN   SMLT ID
------------------------------------------------------------------------------
1      60    180      10
2      60    180      10, 16

VID    IPv6            MAC               ADMIN    OPER   HDTMR   HUTMR
------------------------------------------------------------------------------


VID    HDT REMAIN   HUT REMAIN   SMLT ID
------------------------------------------------------------------------------
Switch:1(config-if)#no ip rsmlt
Switch:1(config-if)#ip rsmlt
```

Enabling Trace Messages for RSMLT Troubleshooting

Use the following procedure to obtain additional RSMLT-related information.

Procedure

If the preceding information does not resolve the issue, you can use the following command to obtain additional RSMLT-related information:

```
trace level 173 4
```

Important:

Enabling this trace on a loaded system can slow down the CPU, especially if executed through the console. Use Telnet if possible.

Troubleshooting IPv6 Connectivity Loss

If the switch experiences loss of IPv6 connectivity, use the following procedure to troubleshoot the issue.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Through the command line interface, make sure the required routes are in place and the corresponding neighbor entries are resolved (that is, in REACHABLE, PROBE, DELAY or STALE state).
3. INCOMPLETE neighbor state indicates a problem if the corresponding neighbor is used by some of the IPv6 routes. This applies to neighbor entries with link-local addresses.

   Note:

   Global addresses are not normally used as next hops. Having a global IPv6 neighbor entry as INCOMPLETE does not usually lead to a connectivity issue.
4. If the corresponding route is not in place then this is a routing issue. If the neighbor is not present or is INCOMPLETE, then further debugging is needed on the network level (that is, the state of other nodes needs to be examined).
5. Disabling and re-enabling IPv6 on the VLAN often recovers connectivity.
6. Display the RSMLT and MLT status:

   ```
   show ip rsmlt
   show mlt
   ```

   Make sure the RSMLT peer MAC is learned and the IST state is `ist`.

Troubleshooting vIST Failure

About this task

When you use Virtual Inter-Switch Trunk (vIST), all critical network traffic runs on this link. If vIST fails, network protocols such as RIP, VRRP, OSPF, and VLACP go down and eventually cause a network outage.

vIST uses an SPBM tunnel to virtually connect two nodes that can be anywhere in the SPBM cloud. Even if the two vIST nodes are directly connected by an MLT link, the vIST VLAN does not have MLT ports as members. Instead, it is configured to be an SPBM C-VLAN.

Note:

For more information on vIST and a configuration example, see Configuring Link Aggregation, MLT, SMLT and vIST for VOSS.

The vIST tunnel is up as long as there is SPBM connectivity between the IST peers. If there is a vIST failure, check the following procedure for some possible reasons:

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Verify that the vIST VLAN is configured on the vIST switch:

   ```
   show virtual-ist
   ```

3. Verify that an I-SID is associated with the vIST VLAN:

   ```
   show isis spbm i-sid discover
   ```

   Important:

   The I-SID associated with the vIST VLAN should be the same on the vIST peer, and this I-SID should not be used anywhere else in the network.

4. Verify that the vIST peers are on the same subnet.
5. If peer ARP is not resolved, enable **`trace level 14`** to see if ARP request/response are being sent/received.
6. If vIST is not up, check the mac fdb table and verify that the peer MAC is synchronized:

   ```
   show vlan mac-address-entry <1-4059>
   ```

7. If the vIST peer MAC is learned, check to see if the peer IP address is reachable.
   a. Use **`show virtual-ist`** to obtain the vIST peer IP address.
   b. Ping the peer IP address.
8. If unable to **`ping`** the peer IP address, check to see if ARP is resolved.

   ```
   show ip arp vlan <vid>
   ```

Troubleshooting Transparent Port UNI

Use the information in this section to troubleshoot problems with Transparent Port UNI (T-UNI), using the CLI.

Viewing all Configured I-SIDs

Perform this procedure to view all the configured I-SIDs including their types, ports, and MLTs.

About this task

View all configured I-SIDs (both CVLAN and T-UNI). View also the I-SID types and the ports or MLTs that are assigned to each I-SID.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. View all configured I-SIDs. This command displays both CVLAN and T-UNI based I-SIDs.

   ```
   show i-sid
   ```

3. View all T-UNI (Elan-Transparent) I-SIDs.

   ```
   show i-sid [elan-transparent]
   ```

4. View information for a particular T-UNI I-SID.

   ```
   show i-sid [<1-16777215>]
   ```

5. View all IS-IS SPBM I-SID information by I-SID ID:

   ```
   show isis spbm i-sid {all|config|discover} [vlan <2-4059>] [id <1-16777215>] [nick-name <x.xx.xx>]
   ```

Example

View all configured I-SIDs.

```
Switch:1(config)#show i-sid
================================================================================
                                Isid Info
================================================================================

ISID        ISID                    PORT               MLT
ID          TYPE       VLANID       INTERFACES         INTERFACES        ORIGIN
--------------------------------------------------------------------------------
2           CVLAN      2            -                  -                  CONFIG
111         ELAN_TR    N/A          1/2-1/8,8/11       111                CONFIG
1000        ELAN_TR    N/A          2/19,8/23          1,5-6,20,30        CONFIG
10001       ELAN       1000         -                  -                  CONFIG

c: customer vid    u: untagged-traffic

All 4 out of 4 Total Num of i-sids displayed
```

View T-UNI (ELAN Transparent) I-SIDs.

```
Switch:1(config)#show i-sid elan-transparent
================================================================================
                                Isid Info
================================================================================
ISID        ISID                        PORT               MLT
ID          TYPE           VLANID       INTERFACES         INTERFACES
--------------------------------------------------------------------------------
100         ELAN_TR        N/A          1/12
6           ELAN_TR        N/A          1/2

All 2 out of 2 Total Num of elan-tp i-sids displayed
```

View MLT or port information for a particular T-UNI I-SID.

```
Switch:1(config)#show i-sid 111

================================================================================
                                Isid Info
```

```
================================================================================
          PORT               MLT
ISID        ISID                    PORT               MLT
ID          TYPE       VLANID       INTERFACES         INTERFACES        ORIGIN
--------------------------------------------------------------------------------
111         ELAN_TR    N/A          1/2-1/8,8/11       111                CONFIG
```

View all IS-IS SPBM I-SID information:

```
Switch:1#show isis spbm i-sid all

================================================================================
                                SPBM ISID INFO
================================================================================
ISID      SOURCE NAME   VLAN   SYSID                TYPE         HOST_NAME
--------------------------------------------------------------------------------
100       1.11.16       20    0014.c7e1.33df       config       Switch1
6         1.11.20       10    0014.c723.67df       discover     Switch2
--------------------------------------------------------------------------------
 Total number of SPBM ISID entries configured: 1
--------------------------------------------------------------------------------
 Total number of SPBM ISID entries discovered: 1
--------------------------------------------------------------------------------
 Total number of SPBM ISID entries: 2
--------------------------------------------------------------------------------
```

View all IS-IS SPBM I-SID information by I-SID ID:

```
Switch:1#show isis spbm i-sid all id 300

================================================================================
                                SPBM ISID INFO
================================================================================
ISID   SOURCE NAME   VLAN   SYSID                TYPE         HOST_NAME
--------------------------------------------------------------------------------
300    7.15.16       20    a425.1b51.9484       config       Switch1
300    4.01.18       10    b4a9.5a2a.d065       discover     Switch2

--------------------------------------------------------------------------------
 Total number of SPBM ISID entries configured: 1
--------------------------------------------------------------------------------
 Total number of SPBM ISID entries discovered: 1
--------------------------------------------------------------------------------
 Total number of SPBM ISID entries: 2
--------------------------------------------------------------------------------
```

Variable Definitions

Use the data in the following table to use the **`show i-sid`** command.

Note:

When SPB is enabled, I-SID IDs 16777216 and greater are reserved for dynamic data I-SIDs, used to carry Multicast traffic over SPB.

Variable	Value
<1–16777215>	Specifies the service interface identifier (ISID).
elan-transparent	Displays only all the Elan-Transparent (T-UNI based) ISIDs.

Table continues…

Variable	Value
spbm i-sid {all|config|discover}	• all: displays all I-SID entries • config: displays configured I-SID entries • discover: displayes discovered I-SID entries
vlan <*2-4059*>	Displays I-SID information for the specified SPBM VLAN.
id <1–16777215>	Displays I-SID information for the specified I-SID.
nick-name <x.xx.xx>	Displays I-SID information for the specified nickname.

Job Aid

The following table describes the fields in the output for the **`show i-sid`** command.

Table 18: show i-sid

Field	Description
ISID ID	Specifies the service interface identifier (I-SID)
ISID TYPE	Specifies the type of I-SID
VLANID	Specifies the backbone VLAN
PORT INTERFACES	Specifies the port that is assigned to the I-SID
MLT INTERFACES	Specifies the mlt that is assigned to the I-SID

The following describes the fields in the output for the **`show isis spbm i-sid`** command.

Table 19: show isis spbm i-sid

Field	Description
ISID	Indicates the IS-IS SPBM I-SID identifier.
SOURCE NAME	Indicates the nickname of the node where this I-SID was configured or discovered. **Note:** SOURCE NAME is equivalent to nickname.
VLAN	Indicates the B-VLAN where this I-SID was configured or discovered.
SYSID	Indicates the system identifier.
TYPE	Indicates the SPBM I-SID type as either configured or discovered.
HOST_NAME	Indicates the host name of the multicast FIB entry.

Viewing C-MACs Learned on T-UNI Ports for an I-SID

Perform this procedure to view the I-SID bridge forwarding database.

About this task

The **`show i-sid mac-address-entry`** command displays the C-MACs learned on T-UNI I-SIDs. It also displays the C-MACs learned on T-UNI I-SIDs for a specific I-SID, MAC address, port or port list or remote MAC address.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View C-MACs learned on the T-UNI I-SIDs:

   `show i-sid mac-address-entry [<1-16777215>] [mac <0x00:0x00:0x00:0x00:0x00:0x00>] [port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}] [remote]`

Example

View C-MACs learned on all T-UNI I-SIDs.

```
Switch:1#show i-sid mac-address-entry

================================================================================================
                                        I-SID Fdb Table
================================================================================================
I-SID   STATUS     MAC-ADDRESS         INTERFACE  TYPE      DEST-MAC          BVLAN  DEST-SYSNAME
------------------------------------------------------------------------------------------------
100     learned    cc:f9:54:ae:28:81   Port-1/16  LOCAL     00:00:00:00:00:00      0
4       learned    cc:f9:54:ae:2c:18   mlt-6      LOCAL     00:00:00:00:00:00      0
252     learned    cc:f9:54:ae:38:64   Port-1/15  REMOTE    00:13:0a:0c:d3:e0  128   DIST-1B


All 3 out of 3 Total Num of i-sid FDB Entries displayed
```

View C-MACs learned on a specific T-UNI I-SID.

```
Switch:1#show i-sid mac-address-entry 100

================================================================================================
                                        I-SID Fdb Table
================================================================================================
I-SID STATUS   MAC-ADDRESS        INTERFACE  TYPE    DEST-MAC          BVLAN DEST-SYSNAME
------------------------------------------------------------------------------------------------
100   learned  cc:f9:54:ae:28:81  Port-1/16  LOCAL  00:00:00:00:00:00      0

All 1 out of 1 Total Num of i-sid FDB Entries displayed
```

```
Switch:1#show i-sid mac-address-entry 252

================================================================================================
                                        I-SID Fdb Table
================================================================================================
I-SID STATUS   MAC-ADDRESS        INTERFACE  TYPE    DEST-MAC          BVLAN DEST-SYSNAME
------------------------------------------------------------------------------------------------
252   learned  cc:f9:54:ae:38:64  Port-1/15  REMOTE  00:13:0a:0c:d3:e0 128   DIST-1B

All 1 out of 1 Total Num of i-sid FDB Entries displayed
```

View C-MACs learned on a T-UNI I-SID for a specific MAC address.

```
Switch:1#show i-sid mac-address-entry mac cc:f9:54:ae:38:64

================================================================================================
                                        I-SID Fdb Table
================================================================================================
I-SID STATUS   MAC-ADDRESS        INTERFACE  TYPE    DEST-MAC          BVLAN DEST-SYSNAME
------------------------------------------------------------------------------------------------
252   learned  cc:f9:54:ae:38:64  Port-1/15  REMOTE  00:13:0a:0c:d3:e0 128   DIST-1B

All 1 out of 1 Total Num of i-sid FDB Entries displayed
```

View C-MACs learned on aT-UNI I-SID for a specific port.

```
Switch:1#show i-sid mac-address-entry port 1/15

================================================================================================
                                        I-SID Fdb Table
================================================================================================
I-SID STATUS   MAC-ADDRESS        INTERFACE  TYPE    DEST-MAC          BVLAN DEST-SYSNAME
------------------------------------------------------------------------------------------------
252   learned  cc:f9:54:ae:38:64  Port-1/15  REMOTE  00:13:0a:0c:d3:e0 128   DIST-1B

All 1 out of 1 Total Num of i-sid FDB Entries displayed
```

View C-MACs learned on a T-UNI I-SID as a remote MAC address.

```
Switch:1#show i-sid mac-address-entry remote

================================================================================================
                                        I-SID Fdb Table
================================================================================================
I-SID STATUS   MAC-ADDRESS        INTERFACE  TYPE    DEST-MAC          BVLAN DEST-SYSNAME
------------------------------------------------------------------------------------------------
252   learned  cc:f9:54:ae:38:64  Port-1/15  REMOTE  00:13:0a:0c:d3:e0 128   DIST-1B

All 1 out of 1 Total Num of i-sid FDB Entries displayed
```

Variable Definitions

Use the data in the following table to use the **`show i-sid mac-address-entry`** command.

Variable	Value
<1-16777215>	Displays the MAC address learned on the service interface identifier (ISID).
mac *<0x00:0x00:0x00:0x00:0x00:0x00>*	Displays the I-SID FDB details for the specified MAC address.
port *{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}*	Displays the MAC address learned on the specified port or port list.
remote	Displays the remote MAC address learned on the I-SID.

Job Aid

The following table describes the fields in the output for the **`show i-sid mac-address-entry`** command.

Table 20: show i-sid

Field	Description
I-SID	Specifies the service interface identifier (I-SID).
STATUS	Specifies the learning status of the associated MAC.
MAC-ADDRESS	Specifies the MAC address of the port assigned to the specific I-SID or MAC learned on the specific I-SID.
INTERFACE	Specifies the port or MLT on which the MAC is learned for the specific I-SID.
TYPE	Specifies whether the MAC is a Local or IST PEER or a Remote MAC.
DEST-MAC	Specifies the virtual BMAC address or system ID, in MAC format, of the destination node.
BVLAN	Specifies the BVLAN on which the destination node is discovered for the I-SID.
DEST-SYSNAME	Specifies the destination system name.

Multicast Troubleshooting

Use the following information to troubleshoot multicast features and multicast routing.

Multicast Feature Troubleshooting

Use the information in this section to troubleshoot multicast feature problems.

Troubleshooting IGMP Layer 2 Querier

The following sections provide troubleshooting information for the IGMP Layer 2 Querier feature.

Querier Not Elected

If a Querier is not elected, use the following procedure to troubleshoot the issue.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. As the IGMP Layer 2 Querier is based on IGMP snoop, check whether IGMP snoop is enabled on the VLAN:

   ```
   show ip igmp interface vlan
   ```

If IGMP snoop is disabled, the Layer 2 Querier cannot work until IGMP snoop and IGMP Layer 2 Querier are reenabled.

Example

Check whether IGMP snoop is enabled on the VLAN:

```
Switch:1>enable
Switch:1#show ip igmp interface vlan

================================================================================
                                  Vlan Ip Igmp
================================================================================
VLAN QUERY QUERY ROBUST VERSION LAST  PROXY  SNOOP  SSM    FAST   FAST
ID   INTVL MAX                 MEMB  SNOOP  ENABLE SNOOP  LEAVE  LEAVE
           RESP                QUERY ENABLE        ENABLE ENABLE PORTS
--------------------------------------------------------------------------------
1    125   100   2      2       10    false  false  false  false
2    125   100   2      2       10    false  false  false  false
3    125   100   2      2       10    false  false  false  false
4    125   100   2      2       10    false  false  false  false
5    125   100   2      2       10    false  false  false  false
10   125   100   2      2       10    false  false  false  false
100  125   100   2      2       10    false  false  false  false
200  125   100   2      2       10    false  false  false  false
300  125   100   2      2       10    false  false  false  false
444  125   100   2      2       10    false  false  false  false

All 10 out of 10 Total Num of Igmp entries displayed

VLAN SNOOP   SNOOP           DYNAMIC   COMPATIBILITY EXPLICIT
ID   QUERIER QUERIER         DOWNGRADE MODE          HOST

     ENABLE  ADDRESS         VERSION                 TRACKING
--------------------------------------------------------------------------------
1    false   0.0.0.0         enable    disable       disable
2    false   0.0.0.0         enable    disable       disable
3    false   0.0.0.0         enable    disable       disable
4    false   0.0.0.0         enable    disable       disable
5    false   0.0.0.0         enable    disable       disable
10   false   0.0.0.0         enable    disable       disable
100  false   0.0.0.0         enable    disable       disable
200  false   0.0.0.0         enable    disable       disable
300  false   0.0.0.0         enable    disable       disable
444  false   0.0.0.0         enable    disable       disable

All 10 out of 10 Total Num of Igmp entries displayed
```

Job Aid

The following table describes the fields in the output for the **`show ip igmp interface vlan`** command.

Note:

The following table shows the field descriptions for this command if you use the optional parameter **`vlan`**. If you do not the output is different.

Field	Description
VLAN ID	Identifies the VLAN or port where IGMP is configured.
QUERY INTVL	Indicates the frequency at which IGMP host query packets transmit on this interface.
QUERY MAX RESP	Indicates the maximum query response time (in tenths of a second) advertised in IGMPv2 queries on this interface.
ROBUST	Indicates the robustness variable, which you can configure for the expected packet loss on a subnet. If you expect packet loss on a subnet, increase the robustness variable.
VERSION	Indicates the version of IGMP that runs on this interface. This object configures a router capable of running either version. For IGMP to function correctly, you must configure all routers on a LAN to run the same version of IGMP.
LAST MEMB QUERY	Indicates the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages, and is also the amount of time between group specific query messages. Use this value to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group. This variable does not apply to IGMPv1.
PROXY SNOOP ENABLE	Indicates if proxy snoop is enabled on the interface.
SNOOP ENABLE	Indicates if snoop is enabled on the interface.
SSM SNOOP ENABLE	Indicates if SSM snoop is enabled on the interface.
FAST LEAVE ENABLE	Indicates if fast leave mode is enabled on the interface.
FAST LEAVE PORTS (VLAN parameter only)	Indicates the set of ports that are enabled for fast leave.
VLAN ID	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
SNOOP QUERIER ENABLE	Specifies whether the snoop querier is enabled.
SNOOP QUERIER ADDRESS	Specifies the pseudo address of the IGMP snoop querier.

Table continues…

Field	Description
DYNAMIC DOWNGRADE VERSION	Indicates if the dynamic downgrade feature is enabled.
COMPATIBILITY MODE	Indicates whether compatibility mode is enabled.
EXPLICIT HOST TRACKING	Specifies whether the IGMP protocol version 3 is enabled to track hosts for each channel or groups.

Enabling Trace Messages for IGMP Layer 2 Querier Troubleshooting

If the preceding information does not address your issue, you can also use the following trace command to view additional information related to Layer 2 querier.

Caution:

Using the trace tool inappropriately can cause primary CPU lockup conditions, loss of access to the device, loss of protocols, and service degradation. If you use trace level 3 (verbose) or trace level 4 (very verbose), do not use the screen to view commands due to the volume of information the system generates and the effect on the system.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Use the following trace command to begin the trace operation for additional information related to Layer 2 querier:

   ```
   trace level 23 <1-4>
   ```

3. Stop tracing:

   ```
   trace shutdown
   ```

4. View the trace results:

   ```
   trace screen enable
   ```

5. View trace saved to a file:

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

Variable Definitions

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

Variable	Value
level [<*Module_ID*>] *[<1–4>]*	Starts the trace by specifying the module ID and level. Module ID 23 represents the IGMP module <*Module_ID*> specifies the module for the trace. Different hardware platforms support different ID ranges because of feature support differences. To see which module IDs are available on the switch,

Table continues…

Variable	Value
	use the **show trace modid-list** command or CLI command completion Help. <0-4> specifies the trace level: • 0 — Disabled • 1 — Very terse • 2 — Terse • 3 — Verbose • 4 — Very verbose
shutdown	Stops the trace operation.
screen {disable|enable}	Enables or disables the display of trace output to the screen. **Important:** Avoid using the screen to view commands if you use trace level 3 (verbose) or trace level 4 (very verbose) due to the volume of information generated and the effect on the system.

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

Variable	Value
file [tail]	Displays the trace results saved to a file.
level	Displays the current trace level for all modules.
modid-list	Specifies the module ID list.

Troubleshooting IGMPv3 Backwards Compatibility

If you configure the switch to operate in v2-v3 compatibility mode, the switch supports all IGMPv2 and v3 messages. The switch parses the group address of the messages. If the group address is out of SSM range and it is a v3 message, the switch drops the message. If it is a v2 message, IGMP snoop processes handle the message.

To troubleshoot issues with the IGMPv3 backwards compatibility feature, perform the following procedure.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Verify that the SSM static channel is configured for the v1/v2 joins received. Display the configured SSM static channels:

   ```
   show ip igmp ssm-map
   ```

3. Verify that the SSM group range is configured for the v1/v2 joins received. Display the configured SSM group range:

   ```
   show ip igmp ssm
   ```

Example

Display the configured SSM static channels and display the configured SSM group range:

```
Switch:>enable
Switch:1#show ip igmp ssm-map

================================================================================
                                Igmp Ssm Channel
================================================================================
GROUP           SOURCE          MODE       ACTIVE     STATUS
--------------------------------------------------------------------------------
233.252.0.1       192.0.2.200 dynamic    false      enabled
233.252.0.2       192.0.2.200 dynamic    false      enabled
233.252.0.3       192.0.2.200 dynamic    false      enabled
233.252.0.4       192.0.2.200 dynamic    false      enabled
233.252.0.5       192.0.2.200 dynamic    false      enabled
233.252.0.6       192.0.2.200 dynamic    false      enabled
233.252.0.7       192.0.2.200 dynamic    false      enabled
233.252.0.8       192.0.2.200 dynamic    false      enabled
233.252.0.9       192.0.2.200 dynamic    false      enabled
233.252.0.10      192.0.2.200 dynamic    false      enabled

10 out of 10 entries displayed

Switch:1(config)#show ip igmp ssm

================================================================================
                          Igmp Ssm Global - GlobalRouter
================================================================================
DYNAMIC LEARNING     SSM GROUP RANGE
--------------------------------------------------------------------------------
enable               233.252.0.0/255.0.0.0
```

Job Aid

The following table shows the field descriptions for the **show ip igmp ssm-map** command.

Table 21: show ip igmp ssm-map command

Field	Description
GROUP	Indicates the IP multicast group address that uses the default range of 232/8.
SOURCE	Indicates the IP address of the source that sends traffic to the group source.
MODE	Indicates that the entry is a statically configured entry (static) or a dynamically learned entry from IGMPv3 (dynamic).
ACTIVE	Indicates the activity on the corresponding source and group. If the source is active and traffic is flowing to the switch, this status is active; otherwise, it is nonactive.
STATUS	Indicates the administrative state and whether to use the entry. If the status is enabled (default), the entry is used. If the status is disabled, the entry is not used but is saved for future use.

The following table shows the field descriptions for the **show ip igmp ssm** command.

Table 22: show ip igmp ssm command

Field	Description
DYNAMIC LEARNING	Indicates whether dynamic learning is enabled at a global level.
SSM GROUP RANGE	Indicates the IP address range for the SSM group.

Multicast Routing Troubleshooting Using CLI

Use the information in this section to help you troubleshoot multicast routing problems.

Viewing IGMP Interface Information

Perform this procedure to view the IGMP interface table.

About this task

If an interface does not use an IP address, it does not appear in the IGMP table. One exception is an IGMP snooping interface, which does not require an interface IP address.

If an interface uses an IP address, but neither IGMP snoop or PIM is enabled, the interface appears as inactive in the Status field.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. View IGMP interfaces:

   ```
   show ip igmp interface [gigabitethernet {slot/port[/sub-port][-slot/port[/sub-port]][,...]}|vlan <1-4059>] [vrf WORD<1-16>][vrfids WORD<0-512>]
   ```

Example

View IGMP interfaces:

```
Switch:1#show ip igmp interface

================================================================================
                          Igmp Interface - GlobalRouter
================================================================================
       QUERY              OPER            QUERY   WRONG
LASTMEM

IF     INTVL STATUS VERS. VERS QUERIER    MAXRSPT QUERY JOINS ROBUST QUERY
MODE

--------------------------------------------------------------------------------
V300   125    active 3     3   21.0.0.12  100     0     674   2      10
pim
V400   125    active 3     3   41.0.0.12  100     0     0     2      10
pim
V500   125    active 3     3   31.0.0.12  100     0     3707  2      10
pim
V700   125    active 2     2   62.0.0.206 100     0     336   2      10
```

```
pim
V701    125     active 1      1    62.0.1.206 100      0      336    2       10     pim

5 out of 5 entries displayed
```

Variable Definitions

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

Variable	Value
gigabitethernet *{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}*	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port. If you do not specify a slot and port, the command output includes all IGMP interfaces.
vlan *<1-4059>*	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. If you do not specify a VLAN ID, the command output includes all IGMP interfaces.
vrf *WORD <1–16>*	Optionally, identifies the VRF name. If you do not specify a VRF name, the results display information for the Global Router. If you specify a VRF name, the results display information only for the VRF you specify.
vrfids *WORD <0–512>*	Optionally, identifies the VRF ID. If you do not specify a range of VRF IDs, the results display information for the Global Router. If you specify a VRF ID or range of VRF IDs, the results display information only for the VRF you specify.

Job Aid

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

Table continues…

Field	Description
STATUS	Indicates the activation of a row, which activates IGMP on the interface. The destruction of a row disables IGMP on the interface.
VERS	Indicates the version of IGMP that runs on this interface. This object configures a router capable of running either version. For IGMP to function correctly, you must configure all routers on a LAN to run the same version of IGMP.
OPER VERS	Indicates the operational version of IGMP.
QUERIER	Indicates the address of the IGMP querier on the IP subnet to which this interface attaches.
QUERY MAXRSPT	Indicates the maximum query response time (in tenths of a second) advertised in IGMPv2 queries on this interface.
WRONG QUERY	Indicates the number of queries received whose IGMP version does not match the interface version. You must configure all routers on a LAN to run the same version of IGMP. If queries are received with the wrong version, a configuration error occurs.
JOINS	Indicates the number of times this interface added a group membership.
ROBUST	Indicates the robustness variable, which you can configure for the expected packet loss on a subnet. If you expect packet loss on a subnet, increase the robustness variable.
LASTMEM QUERY	Indicates the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages, and is also the amount of time between group specific query messages. Use this value to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group. This variable does not apply to IGMPv1.

The following table shows the field descriptions for the command output if you use the interface parameters.

Table 23: show ip igmp interface command output with interface parameters

Field	Description
VLAN ID or PORT NUM	Identifies the VLAN or port where IGMP is configured.
QUERY INTVL	Indicates the frequency at which IGMP host query packets transmit on this interface.

Table continues…

Field	Description
QUERY MAX RESP	Indicates the maximum query response time (in tenths of a second) advertised in IGMPv2 queries on this interface.
ROBUST	Indicates the robustness variable, which you can configure for the expected packet loss on a subnet. If you expect packet loss on a subnet, increase the robustness variable.
VERSION	Indicates the version of IGMP that runs on this interface. This object configures a router capable of running either version. For IGMP to function correctly, you must configure all routers on a LAN to run the same version of IGMP.
LAST MEMB QUERY	Indicates the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages, and is also the amount of time between group-specific query messages. Use this value to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group. This variable does not apply to IGMPv1.
PROXY SNOOP ENABLE	Indicates if proxy snoop is enabled on the interface.
SNOOP ENABLE	Indicates if snoop is enabled on the interface.
SSM SNOOP ENABLE	Indicates if SSM snoop is enabled on the interface.
FAST LEAVE ENABLE	Indicates if fast leave mode is enabled on the interface.
FAST LEAVE PORTS (VLAN parameter only)	Indicates the set of ports that are enabled for fast leave.
DYNAMIC DOWNGRADE VERSION	Indicates if the dynamic downgrade feature is enabled.
COMPATIBILITY MODE	Indicates if compatibility mode is enabled.
EXPLICIT HOST TRACKING	Indicates if explicit host tracking is enabled for IGMPv3. Explicit host tracking enables the IGMP to track all source and group members.

Viewing Multicast Group Trace Information for IGMP Snoop

About this task

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

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Display the multicast group trace for an IGMP snoop-enabled interface:

```
show ip igmp snoop-trace [source {A.B.C.D}] [group {A.B.C.D}]
```

Example

Display the multicast group trace for an IGMP snoop-enabled interface:

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

================================================================================
                            Snoop Trace - GlobalRouter
================================================================================
GROUP               SOURCE              IN     IN          OUT    OUT         TYPE
ADDRESS             ADDRESS             VLAN   PORT        VLAN   PORT
--------------------------------------------------------------------------------
203.0.113.17        192.0.2.7          1015   1/3         1015   1/35-1/40   ACCESS
203.0.113.20        192.0.2.7          1015   1/3         1015   1/35-1/40   ACCESS
203.0.113.19        192.0.2.7          1015   1/3         1015   1/35-1/40   ACCESS
203.0.113.18        192.0.2.7          1015   1/3         1015   1/35-1/40   ACCESS
```

Variable Definitions

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

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.

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.
TYPE	Indicates where the stream is learned. ACCESS indicates the stream is learned locally.

Viewing IGMP Group Information

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

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. View IGMP group information:

```
show ip igmp group group <A.B.C.D> detail [port {slot/port[/sub-port][-slot/port[/sub-port]][,...]}][vlan <1-4059>] [vrf WORD <1-16>][vrfids WORD <0-512>]

show ip igmp group group <A.B.C.D> tracked-members [member-subnet <A.B.C.D./X>][port {slot/port[/sub-port][-slot/port[/sub-port]][,...]}] [source-subnet <A.B.C.D/X>] [vlan <1-4059>][vrf WORD <1-16>][vrfids WORD <0-512>]
```

Example

View IGMP group information:

```
Switch:1>enable
Switch:1#show ip igmp group group 232.0.0.0

================================================================================
                            Igmp Group - GlobalRouter
================================================================================
GRPADDR         INPORT          MEMBER          EXPIRATION TYPE
--------------------------------------------------------------------------------
232.0.0.0       V1015-1/2       200.0.15.53     258        Dynamic

1 out of 271 group Receivers displayed

Total number of unique groups 271
```

Variable Definitions

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

Variable	Value
count	Displays the number of entries in the IGMP group.
group <*A.B.C.D*>	Specifies the address of the IGMP group.
member-subnet {default| <*A.B.C.D*>}]	Specifies the IP address and mask of the IGMP member.
vrf *WORD<1–16>*	Specifies the VRF name.
vrfids *WORD<0–512>*	Specifies the VRF ID.

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

Variable	Value
detail [port *{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}*|vlan <*1-4059*>|vrf*WORD <1–16>*| vrfids*WORD <0–255>*]	Use the detail parameter to show IGMPv3–specific data. For data related to a specific interface use the following: • port*{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}* — Specifies the port list. • vlan <*1-4059*>— Specifies the VLAN. Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-

Table continues…

Variable	Value
	config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. • vrf *WORD<1–16>* — Specifies the VRF name. • vrfids *WORD<0–255>* — Specifies the VRF ID.
tracked-members	Use the tracked-members parameter to view all the tracked members for a specific group.
vrf *WORD<1–16>*	Specifies the VRF name.
vrfids *WORD<0–512>*	Specifies the VRF ID.

Use the data in the following table to use the **`show ip igmp group group <A.B.C.D> tracked-members`** command.

Variable	Value
member-subnet {default| <*A.B.C.D*>}]	Specifies the IP address and mask of the IGMP member.
port *{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}*	Specifies the port list.
source-subnet <*A.B.C.D/X*>	Specifies the source IP address and the subnet mask.
vlan <*1-4059*>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
vrf *WORD<1–16>*	Specifies the VRF name.
vrfids *WORD<0–512>*	Specifies the VRF ID.

Job Aid

The following table shows the field descriptions for the show **`ip igmp group group`** command output.

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.

Table continues…

Field	Description
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.

Displaying the SPBM Multicast Database

You can determine the database used by the SPBM multicast module by using the following procedure.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Show the SPBM multicast database:

   ```
   show isis spbm ip-multicast-route [all][detail][group {A.B.C.D}]
   [vlan <2-4059>][vrf WORD<0-16>][vsn-isid <1-16777215>]
   ```

   Important:

   When you use this command without parameters or use the detail or group optional parameters without specifying a VLAN ID or VSN-ISID, the command output displays Layer 3 context only. No Layer 2 context is displayed.

Example

Show the SPBM multicast database:

```
Switch(config)#show isis spbm ip-multicast-route

================================================================================
                        SPBM IP-MULTICAST FIB ENTRY INFO
================================================================================
Source          Group           Data ISID  BVLAN Source-BEB
--------------------------------------------------------------------------------
192.2.0.1       233.252.0.246        16000001  101   EVP
--------------------------------------------------------------------------------
 Total Number of SPBM IP MULTICAST ROUTE Entries: 1
--------------------------------------------------------------------------------
```

Variable Definitions

Use the data in the following table to use the **`show isis spbm ip-multicast-route`** command.

Variable	Value
all	Displays all IP Multicast over Fabric Connect route information.

Table continues…

Variable	Value
detail	Displays detailed IP Multicast over Fabric Connect route information.
group {A.B.C.D} source {A.B.C.D}	Displays information on the group IP address for the IP Multicast over Fabric Connect route. If you select source it will also display the source IP address.
vlan *<2–4059>*	Displays IP Multicast over Fabric Connect route information by VLAN.
vrf *WORD<0–16>*	Displays IP Multicast over Fabric Connect route information by VRF.
vsn-isid *<1–16777215>*	Displays IP Multicast over Fabric Connect route information by I-SID.

Job Aid

The following table describes fields for the **`show isis spbm ip-multicast-route`** command.

Table 24: show isis spbm ip-multicast-route command

Field	Description
Source	Specifies the IP address of the Global Routing Table.
Group	Specifies the IP multicast group for which this entry specifies a next hop on an outgoing interface.
Data ISID	Specifies the VRF ID for the multicast route.
BVLAN	Specifies the Backbone VLAN (B-VLAN).
Source-BEB	Specifies the source Backbone Edge Bridge (BEB).
Total number of SPBM IP_MULTICAST Route entries	Specifies the number of SPBM IP multicast route entries.

Troubleshooting IP Multicast over Fabric Connect for Layer 2 VSNs

If traffic is not moving properly, use the following procedure to determine the issue.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Ensure that all switch nodes in the network operate with the most recent software release to support IP Multicast over Fabric Connect:

   ```
   show software
   ```

3. If any ERS 8800 nodes exist in the network, ensure you upgrade them to the current release:

   ```
   show software
   ```

4. Ensure that you create and enable SPBM infrastructure globally.
   a. Ensure that SPBM is enabled globally:

      ```
      show spbm
      ```

   b. Ensure that IS-IS is enabled globally:

      ```
      show isis
      ```

   c. Ensure an SPBM instance exists and at least one Backbone VLAN exists (B-VID). Also ensure multicast is enabled:

      ```
      show isis spbm
      ```

   For more information about infrastructure and services configuration, see Configuring Fabric Basics and Layer 2 Services for VOSS.
5. Ensure that you enable the CFM configuration.
   a. Ensure a CFM maintenance-association exists:

      ```
      show cfm maintenance-association
      ```

   b. Ensure a CFM maintenance-domain exists:

      ```
      show cfm maintenance-domain
      ```

   c. Ensure a maintenance-endpoint exists in the MEP ID column and is enabled in the ADMIN column:

      ```
      show cfm maintenance-endpoint
      ```

6. Ensure a Customer VLAN (C-VLAN) exists and ensure you add UNI ports to the C-VLAN.
   a. Display C-VLAN information:

      ```
      show vlan i-sid
      ```

   b. Display ports for the C-VLAN:

      ```
      show vlan members port {slot/port[/sub-port][-slot/port[/sub-port]][,...]}
      ```

   c. Display NNI and UNI receivers:

      ```
      show isis spbm ip-multicast-route detail
      ```

7. Ensure that you assign the same I-SID to the C-VLAN on all of the BEBs where you configure the C-VLAN:

   ```
   show vlan i-sid
   ```

8. Ensure that you enable IP Multicast over Fabric Connect globally:

   ```
   show isis spbm
   ```

9. Ensure the you enable IGMP Snooping on the C-VLAN on all of the Backbone Edge Bridges (BEBs). Ensure the protocol configured on the VLAN added is snoop-spb in the MODE column, which indicates IGMP is enabled on a VLAN with an associated I-SID (IP Multicast over Fabric Connect for a Layer 2 VSN):

```
show ip igmp interface
```

10. Ensure that you enable IGMP Snooping on access Layer 2 switches to prevent flooding of multicast traffic to non-receiver ports:

    ```
    show ip igmp snoop-trace
    ```

    ```
    show ip igmp interface
    ```

11. Ensure that the IGMP version used by the multicast hosts and the Layer 2 switches outside the SPBM network is the same as the IGMP version configured on the C-VLAN:

    ```
    show ip igmp interface
    ```

Troubleshooting IP Multicast over Fabric Connect for Layer 3 VSNs

If traffic is not moving properly, use the following procedure to determine the issue.

Procedure

1. Ensure that all switch nodes in the network operate with the most recent software release to support IP Multicast over Fabric Connect:

   ```
   show software
   ```

2. If ERS 8800 nodes exist in the network, ensure you upgrade them to the current release:

   ```
   show software
   ```

3. Ensure that you create and enable SPBM infrastructure globally.
   a. Ensure that SPBM is enabled globally:

      ```
      show spbm
      ```

   b. Ensure that IS-IS is enabled globally:

      ```
      show isis
      ```

   c. Ensure an SPBM instance exists and at least one Backbone VLAN exists (B-VID). Also ensure multicast is enabled:

      ```
      show isis spbm
      ```

   For more information on infrastructure and services configuration, see Configuring Fabric Basics and Layer 2 Services for VOSS.

4. Ensure that you enable the CFM configuration.
   a. Ensure a CFM maintenance-association exists:

      ```
      show cfm maintenance-association
      ```

   b. Ensure a CFM maintenance-domain exists:

      ```
      show cfm maintenance-domain
      ```

   c. Ensure a maintenance-endpoint exists in the MEP ID column and is enabled in the ADMIN column:

      ```
      show cfm maintenance-endpoint
      ```

5. Ensure the following on all the Backbone Edge Bridges (BEBs) where the Layer 3 VSN is present.
   a. Ensure that you enable IP multicast globally:

      ```
      show isis spbm
      ```

   b. Ensure that you create an IPVPN for the VRF:

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

   c. Ensure that you assign an I-SID to the VRF:

      ```
      show isis spbm ip-multicast-route all
      ```

   d. Ensure that you enable the MVPN:

      ```
      show ip vrf mvpn
      ```

6. On the VLANs that need Layer 3 VSN IP Multicast over Fabric Connect routing, create an IP interface on the VLAN if one does not exist. The address should be on the same subnet as the IGMP hosts connected to the VLAN. Also, ensure that you enable IP Multicast over Fabric Connect.
7. Enter VLAN Interface Configuration mode:

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

8. Create an IP interface on the VLAN and enable IP Multicast over Fabric Connect:

   ```
   ip address <A.B.C.D>
   ip spb-multicast enable
   ```

9. Ensure that you enable IGMP Snooping on access Layer 2 switches to prevent flooding of multicast traffic to non-receiver ports:

   ```
   show ip igmp snoop-trace
   show ip igmp interface
   ```

10. Ensure that the IGMP version used by the multicast hosts and the Layer 2 switches outside the SPBM network is the same as the IGMP version configured on the C-VLAN:

    ```
    show ip igmp interface
    ```

Troubleshooting IP Multicast over Fabric Connect for IP Shortcuts

If traffic is not moving properly, use the following procedure to determine the issue.

Procedure

1. Ensure that all switch nodes in the network operate with the most recent software release to support IP Multicast over Fabric Connect:

   ```
   show software
   ```

2. Ensure that all ERS 8800 nodes in the network have the current release:

   ```
   show software
   ```

3. Ensure that you create and enable SPBM infrastructure globally.
   a. Ensure that SPBM is enabled globally:

      ```
      show spbm
      ```

   b. Ensure that IS-IS is enabled globally:

      ```
      show isis
      ```

   c. Ensure an SPBM instance exists and at least one Backbone VLAN exists (B-VID). Also ensure multicast is enabled:

      ```
      show isis spbm
      ```

   For more information on infrastructure and services configuration, see Configuring Fabric Basics and Layer 2 Services for VOSS.

4. Ensure that you enable the CFM configuration.
   a. Ensure a CFM maintenance-association exists:

      ```
      show cfm maintenance-association
      ```

   b. Ensure a CFM maintenance-domain exists:

      ```
      show cfm maintenance-domain
      ```

   c. Ensure a maintenance-endpoint exists in the MEP ID column and is enabled in the ADMIN column:

      ```
      show cfm maintenance-endpoint
      ```

5. Ensure the following on all BEBs where you want IP Multicast over Fabric Connect. Ensure that you enable IP Multicast over Fabric Connect globally:

   ```
   show isis spbm
   ```

6. On the VLANs that need Layer 3 VSN IP Multicast over Fabric Connect routing, create an IP interface on the VLAN if one does not exist. The address should be on the same subnet as the IGMP hosts connected to the VLAN. Also, ensure that you enable IP Multicast over Fabric Connect. Create an IP interface on the VLAN and enable IP Multicast over Fabric Connect.
7. Enter VLAN Interface Configuration mode:

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

8. Create an IP interface on the VLAN and enable IP Multicast over Fabric Connect:

   ```
   ip address <A.B.C.D>
   ip spb-multicast enable
   ```

9. Ensure that you enable IGMP Snooping on access Layer 2 switches to prevent flooding of multicast traffic to non-receiver ports:

   ```
   show ip igmp snoop-trace
   ```

   ```
   show ip igmp interface
   ```

10. Ensure that the IGMP version used by the multicast hosts and the Layer 2 switches outside the SPBM network is the same as the IGMP version configured on the C-VLAN:

    ```
    show ip igmp interface
    ```

Showing the Hardware Resource Usage

About this task

The switch can query the number of ingress and egress IP multicast streams traversing the switch. After you configure the thresholds for ingress and egress records, if the record-usage goes beyond the threshold, the device notifies you by way of a trap on the console, logged message, or both.

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

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Show the hardware resource usage:

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

Example

Show the hardware resource usage:

```
Switch:1>show ip mroute hw-resource-usage
================================================================================
                        Multicast Hardware Resource Usage
================================================================================
EGRESS      INGRESS     EGRESS      INGRESS     LOG MSG   SEND TRAP  SEND TRAP

REC IN-USE  REC IN-USE  THRESHOLD   THRESHOLD   ONLY      ONLY       AND LOG

--------------------------------------------------------------------------------
0           0           0           0           false     false      false
```

Job Aid

The following table shows the field descriptions for the **`show ip mroute hw-resource-usage`** command.

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.

Table continues…

Field	Description
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

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

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Start debug trace message output:

   ```
   debug ip pim pimdbgtrace
   ```

3. Stop debug trace message output:

   ```
   no debug ip pim pimdbgtrace
   default debug ip pim pimdbgtrace
   ```

4. Configure the system to display trace messages forwarded by the device:

   ```
   debug ip pim send-dbg-trace
   ```

5. Stop the system from displaying trace messages forwarded by the device:

   ```
   no debug ip pim send-dbg-trace
   default debug ip pim send-dbg-trace
   ```

6. Configure the system to display trace messages received by the device:

   ```
   debug ip pim rcv-dbg-trace
   ```

7. Stop the system from displaying trace messages received by the device:

   ```
   no debug ip pim rcv-dbg-trace
   default debug ip pim rcv-dbg-trace
   ```

8. Configure the system to display hello messages forwarded or received by the device:

   ```
   debug ip pim hello
   ```

9. Stop the system from displaying hello messages forwarded or received by the device:

   ```
   no debug ip pim hello
   default debug ip pim hello
   ```

10. Configure the system to display and log debug trace messages:

    ```
    debug ip pim pimdbglog
    ```

11. Stop the system from displaying and logging debug trace messages:

    ```
    no debug ip pim pimdbglog
    ```

    ```
    default debug ip pim pimdbglog
    ```

12. Configure the system to display register messages forwarded or received by the device:

    ```
    debug ip pim register
    ```

13. Stop the system from displaying register messages forwarded or received by the device:

    ```
    no debug ip pim register
    ```

    ```
    default debug ip pim register
    ```

14. Configure the system to display debug trace messages after an enabled message type, for example, hello or register, is received from a specific sender IP address:

    ```
    debug ip pim source {A.B.C.D}
    ```

Variable Definitions

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

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).

Determining the Protocol Configured on the Added VLAN

Use this procedure to determine the protocol configured on the added VLAN.

The protocol configured on the added VLAN can be one of the following values:

- snoop
- snoop-spb
- route-spb
- pim

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Determine the protocol configured on the added VLAN:

   ```
   show ip igmp interface [gigabitethernet {slot/port[/sub-port][-slot/port[/sub-port]][,...]}][vlan <1-4059>][vrf WORD<1-16>] [vrfids WORD<0-512>]
   ```

   The protocol displays under the Mode column of the command output.

Example

Determine the protocol configured on the added VLAN:

```
Switch:1enable
Switch:1#show ip igmp interface

================================================================================
                          IGMP Interface - GlobalRouter
================================================================================
       QUERY              OPER            QUERY   WRONG            LASTMEM
IF     INTVL STATUS VERS. VERS QUERIER    MAXRSPT QUERY JOINS ROBUST QUERY MODE
--------------------------------------------------------------------------------
V300   125    activ  3     3   21.0.0.12    100     0     116     2      10    pim

1 out of 1 entries displayed
```

Variable Definitions

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

Variable	Value
gigabitethernet{*slot/port[/sub-port] [-slot/port[/sub-port]] [,...]*}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Table continues…

Variable	Value
vlan*<1-4059>*	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
vrf*WORD<1–16>*	Specifies the VRF instance by the VRF name.
vrfids*WORD<0–512>*	Specifies the VRF ID for which to display statistics.

Job Aid

The following table shows the field descriptions for the **`show ip igmp interface`** command.

Field	Description
IF	Indicates the interfaces where IGMP is configured.
QUERY INTVL	Indicates the frequency at which the interface transmits IGMP host query packets.
STATUS	Indicates the activation of a row that enables IGMP on the interface. The destruction of a row disables IGMP on the interface.
VERS	Indicates the version of IGMP that runs on this interface. This object configures a router capable of running either version. For IGMP to function correctly, you must configure all routers on a LAN to run the same version of IGMP.
OPER VERS	Indicates the operational version of IGMP.
QUERIER	Indicates the address of the IGMP querier on the IP subnet to which this interface is attached.
QUERY MAXRSPT	Indicates the maximum query response time (in tenths of a second) advertised in IGMPv2 queries on this interface.
WRONG QUERY	Indicates the number of queries received whose IGMP version does not match the IGMP Interface version. You must configure all routers on a LAN to run the same version of IGMP. Therefore, if the interface receives queries with the wrong version, a configuration error occurs.
JOINS	Indicates the number of times IGMP added a group membership on this interface.
ROBUST	Indicates the robustness variable, which you configure for the expected packet loss on a subnet. If packet loss is expected on a subnet, increase the robustness variable.
LAST MEM QUERY	Indicates the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages, and is also the amount of time between

Table continues…

Field	Description
	group-specific query messages. Use this value to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group. This does not apply if igmpInterface version is 1.
MODE	Indicates the protocol configured on the VLAN added. • snoop — Indicates IGMP snooping is enabled on a VLAN. • snoop-spb — Indicates IGMP is enabled on a VLAN with an associated I-SID (IP Multicast over Fabric Connect for a Layer 2 VSN.) • routed-spb — Indicates IP Multicast over Fabric Connect is enabled on the Layer 3 VSN or for IP Shortcuts. • pim — Indicates PIM is enabled.

The following table shows the field descriptions for the **`show ip igmp interface`** command output if you use the optional parameters to specify a port, VLAN, or VRF.

Table 25: show ip igmp interface command with optional parameters

Field	Description
VLAN ID or PORT NUM	Identifies the VLAN or port where IGMP is configured.
QUERY INTVL	Indicates the frequency at which IGMP host query packets transmit on this interface.
QUERY MAX RESP	Indicates the maximum query response time (in tenths of a second) advertised in IGMPv2 queries on this interface.
ROBUST	Indicates the robustness variable, which you can configure for the expected packet loss on a subnet. If you expect packet loss on a subnet, increase the robustness variable.
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.

Table continues…

Field	Description
FAST LEAVE PORTS (VLAN parameter only)	Indicates the set of ports that are enabled for fast leave.
VLAN ID or PORT NUM	Identifies the VLAN or port where IGMP is configured.
SNOOP QUERIER ENABLE (VLAN parameter only)	Indicates if the IGMP Layer 2 Querier feature is enabled.
SNOOP QUERIER ADDRESS (VLAN parameter only)	Indicates the IP address of the IGMP Layer 2 Querier.
DYNAMIC DOWNGRADE VERSION	Indicates if the dynamic downgrade feature is enabled.
COMPATIBILITY MODE	Indicates if compatibility mode is enabled.
EXPLICIT HOST TRACKING	Indicates if explicit host tracking is enabled to track all the source and group members.

Multicast routing troubleshooting using EDM

Use the information in this section to help you troubleshoot multicast routing problems using Enterprise Device Manager (EDM).

Viewing IGMP Interface Information

Use the Interface tab to view the IGMP interface table. You can use this procedure to determine the protocol configured on the added VLAN.

The protocol configured on the added VLAN can be one of the following values:

- snoop
- pim

About this task

If an interface does not use an IP address, it does not appear in the IGMP table. If an interface uses an IP address, but neither IGMP snoop or PIM is enabled, the interface appears as inactive in the Status field.

Procedure

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

Interface Field Descriptions

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

Name	Description
IfIndex	Shows the interface where IGMP is enabled.
QueryInterval	Configures the frequency (in seconds) at which the interface transmits IGMP host query packets. 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.
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 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, which is the same as the number of times an entry for this interface is added to the cache table. This number gives an indication of the amount of IGMP activity over time.
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 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. It is recommended that you configure this parameter to values greater than 3. If you do not need a fast leave process, you can configure values greater than 10. (The value 3 is equal to 0.3 seconds and 10 is equal to 1 second.)

Table continues…

Name	Description
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, configure this parameter according to the version of IGMP currently in use. • IGMPv1—Disable • IGMPv2—Enable • IGMPv3—Enable
SsmSnoopEnable	Enables SSM snoop.
SnoopQuerierEnable	Enables IGMP Layer 2 Querier.
SnoopQuerierAddr	Specifies the pseudo address of the IGMP snoop querier.
ExplicitHostTrackingEnable	Enables or disables IGMPv3 to track hosts for each channel or group. The default is disabled. You must select this field if you want to use fast leave for IGMPv3.
McastMode	Indicates the protocol configured on the VLAN. • snoop — Indicates IGMP snooping is enabled on a VLAN. • snoop-spb — Indicates IGMP is enabled on a VLAN with an associated I-SID (IP multicast over Fabric Connect for a Layer 2 VSN). • pim — Indicates PIM is enabled. • routed-spb — Indicates IP multicast over Fabric Connect is enabled on the Layer 3 VSN or for IP Shortcuts.

Viewing IGMP Group Information

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

About this task

Note:

The following procedure displays the dynamically learned IGMP groups. **IP** > **IGMP** > **Static** displays statically configured IGMP groups. This is in contrast to the CLI command **`show ip igmp group`**, which displays both dynamically learned and statically configured IGMP groups, and the CLI command **`show ip igmp static`**, which displays only the statically configured groups.

You can view IGMP information on a VRF instance the same way you view the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **Groups** tab.

Groups Field Descriptions

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

Name	Description
IpAddress	Shows the multicast group address (Class D). A group address can be the same for many incoming ports.
Members	Shows the IP address of the host that issues the membership report to this group.
InPort	Shows the port that receives the group membership report.
IfIndex	Shows a unique value that identifies a physical interface or a logical interface (VLAN) that receives the membership report.
Expiration	Shows the time left before the group report expires on this port. This variable is updated after the port receives a group report.

Viewing IGMP Snoop Trace Information

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

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **Snoop Trace** tab.

Snoop Trace Field Descriptions

Use the data in the following table to use the **Snoop Trace** tab.

Name	Description
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.
Type	Displays the port type on which the snoop entry is learned.

Determining the Data Stream Learned when IP Multicast over Fabric Connect is Configured on the VLAN

Use the following procedure to determine the data stream learned when IP multicast is configured on the VLAN.

Procedure

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

Multicast Field Descriptions

Use the information in the following table to help you use the **Multicast** tab.

Field	Description
Group	Indicates the IP multicast group for which this entry specifies a next hop on an outgoing interface.
Source	Indicates the network address that, when combined with the corresponding value of ipMRouteNextHopSourceMask, identifies the sources for which this entry specifies a next hop on an outgoing interface.
SourceMask	Indicates the network mask, when combined with the corresponding value of ipMRouteNextHopSource, identifies the sources for which this entry specifies a next hop on an outgoing interface.
UpstreamNeighbor	Indicates the address of the upstream neighbor from which IP datagrams from these sources to this multicast address are received, or 0.0.0.0 if the upstream neighbor is known.
Interface	Indicates the value of ifIndex for the interface on which IP datagrams sent by these sources to this multicast address are received. A value of 0 indicates that datagrams are not subject

Table continues…

Field	Description
	to an incoming interface check, but can be accepted on multiple interfaces (for example, in CBT).
ExpiryTime	Indicates the minimum amount of time remaining before this entry ages out. The value 0 indicates that the entry is not subject to aging.
Protocol	Indicates the outgoing mechanism through which the switch learns this route. For IP Multicast over Fabric Connect, this value is spb-access or spb-network. Spb-access indicates the datastream learned was from the UNI ports. Spb-network indicates that the datastream learned was from the SPBM cloud.

Showing the SPBM Multicast Database

Determine the database used by the SPBM multicast module.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **ISIS** > **SPBM**.
2. Click the **IpMcastRoutes** tab.

IpMcastRoutes Field Descriptions

Use the information in the following table to use the **IpMcastRoutes** tab.

Name	Description
VsnIsid	Specifies the VSN I-SID.
Group	Specifies the group IP address for the IP multicast route.
Source	Specifies the IP address where the IP multicast route originated from.
SourceBeb	Specifies the Source Backbone Edge Bridge (BEB) for the IP multicast route.
VlanId	Specifies the VLAN ID.
VrfName	Specifies the VRF name.
DataIsid	Specifies the VRF ID for the multicast route.
Type	Specifies the type for the IP multicast route.
Bvlan	Specifies the Backbone VLAN (B-VLAN).
NniInterfaces	Specifies the Network-to-Network Interface ports.

Troubleshoot MACsec

Use the information in this section to troubleshoot problems with the MACsec feature.

The switch also supports viewing MACsec performance statistics. For more information on the supported statistics and procedures to view them, see Monitoring Performance for VOSS.

Viewing the MACsec Connectivity Association Details

Perform this procedure to view the MACsec connectivity association (CA) details.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. View the MACsec CA details:

   ```
   show macsec connectivity-association [WORD<5-15>]
   ```

   Note:

   This command displays the MACsec CA details, including the MD5 hashed value of the CA key.

Example

View the MACsec connectivity association details:

Note:

Slot and port information can differ depending on hardware platform. For more information about specific hardware, see your hardware documentation.

```
Switch:1>show macsec connectivity-association
================================================================================================
                     MACSEC Connectivity Associations Info

================================================================================================
  Connectivity                Connectivity                AN_Mode /      Port
Association Name          Association Key Hash            TxKeyParity     Members
------------------------------------------------------------------------------------------------
caname50               ba6b005bef79e7b95f3e08181e2501ce    2AN / NA       1/49
caname51               5b41f44ecaa54f3873e781557b39230b    4AN / odd
caname15               053f26fb96b011191f2da28849f08677    4AN / Even     1/50

Switch:1#show macsec statistics 1/50 secure-channel inbound

================================================================================================
                  MACSEC Port Inbound Secure Channel Statistics

================================================================================================
           UnusedSA     NoUsingSA       Late         NotValid      Invalid
PortId     Packets      Packets        Packets       Packets       Packets
------------------------------------------------------------------------------------------------
1/47       0               0              0             0             0

           Delayed     Unchecked        Ok         Octets          Octets
PortId     Packets      Packets        Pkts      Validated       Decrypted
------------------------------------------------------------------------------------------------
1/47       0               0              1796          0             169282
```

```
Switch:1#show macsec statistics 1/50 secure-channel outbound

================================================================================
                MACSEC Port Outbound Secure Channel Statistics

================================================================================
          Protected      Encrypted      Octets          Octets
PortId     Packets        Packets      Protected       Encrypted
--------------------------------------------------------------------------------
1/47       0               2628            0               277182
```

Note:

For MACsec Key Agreement (MKA) MACsec , the **show macsec statistics [{slot/port[/sub-port][-slot/port[/sub-port]][,...]}] secure-channel [inbound | outbound]** command displays only the statistics for the current Secure Association (SA), instead of displaying the statistics for all SAs created under that Secure Channel.

Viewing MACsec Status

Perform this procedure to view MACsec status.

About this task

This command displays the status for the following:

- MACsec status
- MACsec encryption status
- MACsec encryption cipher suite, if supported on your hardware platform
- The associated Connectivity Association (CA) name

Note:

If you do not specify a port number, the information on all MACsec capable interfaces is displayed.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. View the MACsec status:

   ```
   show macsec status {slot/port[/sub-port] [-slot/port[/sub-port]]
   [,...]}
   ```

3. Display all MACsec related information:

   ```
   show macsec
   ```

Example

Note:

Slot and port information can differ depending on hardware platform. For more information about specific hardware, see your hardware documentation.

The switch does not support replay protect.

Note:

Configuration of a MACsec cipher suite is not supported on all hardware platforms. If you do not see it in your switch output, it is not supported. For more information on the hardware restrictions, see your hardware documentation.

View MACsec status:

```
Switch:1#show macsec status

================================================================================================================
                                MACSEC Port Status
================================================================================================================
         MACSEC     Encryption  Replay      Replay       Encryption       Cipher     CA       MKA-Profile
PortId   Status     Status      Protect     Protect W'dow   Offset         Suite      Name     Name
----------------------------------------------------------------------------------------------------------------
1/1      enabled    disabled    disabled      --              none         AES-128    SMLTCONN mkapro1
1/2      disabled   disabled    disabled      --              none         AES-128    Nil      --
1/3      disabled   disabled    disabled      --              none         AES-128    Nil      --
1/4      disabled   disabled    disabled      --              none         AES-128    Nil      --
1/5      disabled   disabled    disabled      --              none         AES-128    Nil      --
1/6      disabled   disabled    disabled      --              none         AES-128    Nil      --
1/7      disabled   disabled    disabled      --              none         AES-128    Nil      --
1/8      disabled   disabled    disabled      --              none         AES-128    Nil      --
1/9      disabled   disabled    disabled      --              none         AES-128    Nil      --
1/10     disabled   disabled    disabled      --              none         AES-128    Nil      --
1/11     disabled   disabled    disabled      --              none         AES-128    Nil      --
1/12     disabled   disabled    disabled      --              none         AES-128    Nil      --
1/13     disabled   disabled    disabled      --              none         AES-128    Nil      --

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

View MACsec status on port `1/1`:

```
Switch:1>show macsec status 1/1

================================================================================================================
                                MACSEC Port Status
================================================================================================================
         MACSEC     Encryption  Replay      Replay       Encryption       Cipher     CA       MKA-Profile
PortId   Status     Status      Protect     Protect W'dow   Offset         Suite      Name     Name
----------------------------------------------------------------------------------------------------------------
1/1      enabled    disabled    disabled      --              none         AES-128    SMLTCONN mkaprof1
```

Display all MACsec information:

```
Switch:1#show macsec

================================================================================================
                    MACSEC Connectivity Associations Info
================================================================================================
  Connectivity               Connectivity                AN_Mode /      Port
Association Name          Association Key Hash           TxKeyParity     Members
------------------------------------------------------------------------------------------------
caname33              d4433e901bae92d0cc472706f66cfc18      4AN / odd

All 1 out of 1 Total Num of Macsec connectivity associates displayed


================================================================================================
                                MACSEC Port Status
================================================================================================
```

```
         MACSEC     Encryption  Replay      Replay       Encryption       Cipher      CA
PortId   Status     Status     Protect     Protect W'dow   Offset         Suite       Name
-------------------------------------------------------------------------------------------
1/1      disabled   disabled   disabled      --             none          AES-128     Nil
1/2      disabled   disabled   disabled      --             none          AES-128     Nil
1/3      enabled    enabled    disabled      --        ipv4Offset(30)     AES-256    caname33
1/4      enabled    enabled    disabled      --        ipv4Offset(30)     AES-128    caname34
1/5      disabled   disabled   disabled      --             none          AES-128     Nil
1/6      disabled   disabled   disabled      --             none          AES-128     Nil

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

Troubleshoot MACsec Using EDM

Use the information in this section to troubleshoot problems with the MACsec feature using Enterprise Device Manager (EDM) interface.

View MACsec Connectivity Association Details

Perform this procedure to view the MACsec connectivity association (CA) details.

Procedure

1. In the navigation pane, expand **Configuration** > **Edit**.
2. Select **Chassis**.
3. Select the **MACSec** tab.

MACSec Field Descriptions

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

Name	Description
AssociationName	Specifies a name for each connectivity association configured on the device. **Tip:** Configure the Connectivity Association key name (CKN) in multiples of 4 characters to avoid MKA interoperability issues between VOSS switches and EXOS switches. For example, Macsecma (8 chararcters) or Macsecmka123 (12 characters) are valid, but Macsec (6 characters) is not valid.
AssociationKey	Specifies a pre-shared, connectivity association key associated with each connectivity association configured on the device.

Table continues…

Name	Description
AssociationPortMembers	Specifies the set of ports for which this connectivity association is associated.
AssociationTxKeyParity	Specifies Tx key parity using the following values: • **None** — key parity is not specified **Note:** The none value only applies to platforms that support 2AN mode. If you do not specify a key parity value, the system defaults to 2AN mode. • **Even** — generates even-numbered keys • **Odd** — generates odd-numbered keys

Troubleshooting Fabric Attach

The following sections help you troubleshoot problems with Fabric Attach (FA) using either the Command Line Interface (CLI) or the Enterprise Device Manager (EDM).

Troubleshooting workflow

Troubleshoot FA in the following sequence:

- **Verify FA configuration:**

  As a first step, for proper operation, verify that FA is enabled properly at both the global and interface levels. Use the procedures in this section to verify FA configuration.

- **Verify LLDP port-level transmission and reception:**

  LLDP operates at the interface level. Enabling FA at the port level automatically enables LLDP transmission and reception at the port level. Similarly, enabling FA at the MLT level automatically enables LLDP transmission and reception for all ports in that MLT. Use the procedures in this section to verify LLDP interface (port or MLT) statistics.

- **Verify FA discovery, I-SID-to-VLAN mapping assignments and Switched UNI I-SID creation:**

  After you verify LLDP transmission, verify that FA element discovery completed successfully. After a successful FA discovery, FA clients can send I-SID-to-VLAN mapping assignments to the FA Server on an FA-enabled port or MLT. The FA server accepts or rejects these mapping assignments. A prerequisite to successful mapping assignments is that IS-IS and SPBM are properly configured on the FA server. Successful FA mappings result in the creation of ELAN I-SIDs with end-points of type Switched UNI on the FA Server switch.

Troubleshooting Fabric Attach using the CLI

Verify configuration of Fabric Attach

Viewing Fabric Attach Configuration

To operate properly, Fabric Attach (FA) must be configured properly at both the global and interface level on the switch. Use this procedure to verify FA configuration.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. To verify that FA is enabled globally, enter one of the following commands:

   - `show fa`
   - `show fa agent`

3. To view all FA interfaces (ports and MLTs), enter:

   ```
   show fa interface
   ```

4. To view FA interface configuration on ports, use one of the following commands:

   - To view FA configuration on all ports, enter:

     ```
     show fa interface port
     ```

   - To view FA configuration on a specific port, enter:

     ```
     show fa interface port [{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}]
     ```

5. To view FA interface configuration on MLTs, use one of the following commands:

   - To view FA configuration on all MLTs, enter:

     ```
     show fa interface mlt
     ```

   - To view FA configuration on a specific MLT, enter:

     ```
     show fa interface mlt [<1-512>]
     ```

   - To view FA interface configuration based on the authentication status, enter:

     ```
     show fa interface [enabled-auth] [disabled-auth]
     ```

Example

Verify that FA is configured globally.

```
Switch:1#show fa

================================================================
                        Fabric Attach Configuration
================================================================
                          FA Service : enabled
```

```
                      FA Element Type : server
                  FA Assignment Timeout : 240
                   FA Discovery Timeout : 240
                      FA Provision Mode : spbm
```

Verify FA configuration at the interface (port or MLT) level, on all interfaces.

In the following example output, note that:

- FA is enabled on interfaces `2/10`, `4/11` and `Mlt2`.
- Both FA and message authentication are disabled on port `4/6`.
- Both FA and message authentication are enabled on port `4/11`.

```
Switch:1#show fa interface

============================================================
                            Fabric Attach Interfaces
============================================================
INTERFACE    SERVER   MGMT     MGMT     MSG AUTH MSG AUTH
             STATUS   ISID     CVID     STATUS   KEY
------------------------------------------------------------
Port2/10     enabled  0        0        disabled ****
Port4/6      disabled 0        0        disabled ****
Port4/11     enabled  0        0        enabled  ****
Mlt2         enabled  0        0        disabled ****

------------------------------------------------------------
 4 out of 4 Total Num of fabric attach interfaces displayed
------------------------------------------------------------
```

Verify FA configuration on a specific port, for example, on port `2/10`.

```
Switch:1#show fa interface port 2/10

============================================================
                            Fabric Attach Interfaces
============================================================
INTERFACE    SERVER   MGMT     MGMT     MSG AUTH MSG AUTH
             STATUS   ISID     CVID     STATUS   KEY
------------------------------------------------------------
Port2/10     enabled  0        0        disabled ****

------------------------------------------------------------
 1 out of 1 Total Num of fabric attach interfaces displayed
------------------------------------------------------------
```

Verify FA configuration on an MLT, for example, on `Mlt2`.

```
Switch:1#show fa interface mlt 2

============================================================
                            Fabric Attach Interfaces
============================================================
INTERFACE    SERVER   MGMT     MGMT     MSG AUTH MSG AUTH
             STATUS   ISID     CVID     STATUS   KEY
------------------------------------------------------------
Mlt2         enabled  0        0        disabled ****

------------------------------------------------------------
 1 out of 1 Total Num of fabric attach interfaces displayed
------------------------------------------------------------
```

View the FA interfaces that have authentication enabled:

```
Switch:1#show fa interface enabled-auth

================================================================================
                             Fabric Attach Interfaces
================================================================================
INTERFACE    SERVER   MGMT     MGMT     MSG AUTH MSG AUTH
             STATUS   ISID     CVID     STATUS   KEY
--------------------------------------------------------------------------------
Port4/11     enabled  0        0        enabled  ****
--------------------------------------------------------------------------------
1 out of 1 Total Num of fabric attach interfaces displayed
--------------------------------------------------------------------------------
```

View the FA interfaces that have authentication disabled:

```
Switch:1#show fa interface disabled-auth

================================================================================
                             Fabric Attach Interfaces
================================================================================
INTERFACE    SERVER   MGMT     MGMT     MSG AUTH MSG AUTH
             STATUS   ISID     CVID     STATUS   KEY
--------------------------------------------------------------------------------
Port2/10     enabled  0        0        disabled ****
Port4/6      disabled 0        0        disabled ****
Mlt2         enabled  0        0        disabled ****

--------------------------------------------------------------------------------
3 out of 3 Total Num of fabric attach interfaces displayed
--------------------------------------------------------------------------------
```

Verify LLDP port-level transmission and reception

Viewing Port-based LLDP Statistics

Use this procedure to verify port-based LLDP statistics.

About this task

LLDP operates at the interface level. Enabling FA on a port automatically enables LLDP transmission and reception on the port. It also enables traffic tagging and disables spanning tree on that port.

Enabling FA on an MLT enables LLDP transmission and reception on *all* ports in the MLT.

Note:

When FA is enabled on ports in an MLT or LACP MLT, tagging is enabled and spanning tree is disabled on those ports.

When a port is removed from an MLT, LLDP transmission on that port stops and spanning tree is enabled. Any I-SID-to-VLAN mappings on that port are removed, if not already learned on any other port in the MLT. This also causes the Switched UNI I-SID to be deleted from the MLT. If however, the mappings are learned on another port on the MLT, then the Switched UNI I-SID continues to exist for that MLT.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. To verify successful LLDP transmission on a port, enter:

   ```
   show lldp tx-stats port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}
   ```

3. To verify that a port receives LLDP PDUs successfully, enter:

   ```
   show lldp rx-stats port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}
   ```

4. **(Optional)** To clear LLDP statistics on a port, or ports, enter:

   ```
   clear lldp stats {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}
   ```

Example

Verify LLDP transmission statistics on a port:

```
Switch:1>en
Switch:1#show lldp tx-stats port 1/2
=============================================================
                                LLDP Tx-Stats
=============================================================

PORT NUM                 FRAMES
-------------------------------------------------------------
1/2                      100
```

Verify that the port is receiving LLDP PDUs:

```
Switch:1#show lldp rx-stats port 1/2

================================================================================
                                LLDP Rx-Stats
================================================================================

Port      Frames      Frames   Frames   TLVs        TLVs           AgeOuts
Num       Discarded   Errors   Total    Discarded   Unsupported
                                        (Non FA)    (Non FA)
--------------------------------------------------------------------------------
1/2       0           0        46       0           0              0
```

Variable Definitions

Use the data in the following table to use the **`show lldp tx-stats`** and the **`show lldp rx-stats`** commands.

Variable	Value
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Verify FA discovery and I-SID-to-VLAN mapping assignments

Displaying Learned LLDP Neighbors

Use this procedure to verify details of the LLDP neighbors learned.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Verify details of LLDP neighbors learned:

   ```
   show lldp neighbor
   ```

3. Verify details of LLDP neighbors learned on a specific port:

   ```
   show lldp neighbor port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}
   ```

Example

The following example shows how two switches—an FA Server and an FA Proxy discover each other as LLDP neighbors. Switch A, which is the FA Server is an VSP 7200 Series switch (model 7254XSQ) and switch B which is the proxy device is an ERS 4826GTS switch.

The following examples shows neighbor discovery on non-channelized and channelized ports (if your platform supports channelization).

On the non-channelized port `1/1` on the FA Server, verify neighbor discovery of the proxy switch.

```
SwitchA:1>en
SwitchA:1#show lldp neighbor

================================================================================================
                                        LLDP Neighbor
================================================================================================

Port: 1/1       Index    : 1
                Time: 1 day(s), 04:03:52
                ChassisId: MAC Address        70:30:18:5a:05:00
                PortId   : MAC Address        70:30:18:5a:05:07
                SysName  :
                SysCap   : Br / Br
                PortDescr: Port 7
                SysDescr : Ethernet Routing Switch 4826GTS HW:10 FW:5.8.0.1 SW:v6.9.2.027

------------------------------------------------------------------------------------------------
Total Neighbors : 1
------------------------------------------------------------------------------------------------
Capabilities Legend: (Supported/Enabled)
B= Bridge,    D= DOCSIS,    O= Other,     R= Repeater,
S= Station,   T= Telephone, W= WLAN,      r= Router
Switch:1(config)#
```

On the proxy switch, verify discovery of the FA Server switch.

```
SwitchB:1>en
SwitchB:1#show lldp neighbor
--------------------------------------------------------------------------------
                            LLDP neighbor
--------------------------------------------------------------------------------
```

```
Port: 7     Index: 71
        Time: 12 days, 21:40:30
        ChassisId: MAC address        a4:25:1b:52:70:00
        PortId:    MAC address        a4:25:1b:52:70:04
        SysName:   BEB1-7254XSQ
        SysCap:    rB / rB            (Supported/Enabled)
        PortDesc:  Virtual Services Platform 7254XSQ - Gbic1000BaseT Port
1/1
        SysDescr:  VSP-7254XSQ (6.0.0.0_GA)
--------------------------------------------------------------------------------
Sys capability: O-Other; R-Repeater; B-Bridge; W-WLAN accesspoint; r-Router;
T-Telephone; D-DOCSIS cable device; S-Station only.
Total neighbors: 1
```

On the channelized port `1/1/1` on the FA Server switch, verify discovery of the proxy switch.

```
SwitchA:1>en
SwitchA:1#show lldp neighbor

================================================================================================
                                        LLDP Neighbor
================================================================================================

Port: 1/1/1     Index    : 1
                Time: 1 day(s), 04:03:52
                ChassisId: MAC Address        70:30:18:5a:05:00
                PortId   : MAC Address        70:30:18:5a:05:07
                SysName  :
                SysCap   : Br / Br
                PortDescr: Port 7
                SysDescr : FA Proxy 4826GTS HW:10 FW:5.8.0.1  SW:v5.9.2.027

------------------------------------------------------------------------------------------------
Total Neighbors : 1
------------------------------------------------------------------------------------------------
Capabilities Legend: (Supported/Enabled)
B= Bridge,    D= DOCSIS,    O= Other,     R= Repeater,
S= Station,   T= Telephone, W= WLAN,      r= Router
Switch:1(config)#
```

Verify neighbor discovery on the proxy switch.

```
SwitchB:1>en
SwitchB:1#show lldp neighbor
--------------------------------------------------------------------------------
                            LLDP neighbor
--------------------------------------------------------------------------------
Port: 7     Index: 71
        Time: 12 days, 21:40:30
        ChassisId: MAC address        a4:25:1b:52:70:00
        PortId:    MAC address        a4:25:1b:52:70:04
        SysName:   BEB1-7254XSQ
        SysCap:    rB / rB            (Supported/Enabled)
        PortDesc:  Virtual Services Platform 7254XSQ - 40GbCR4-Channel Port
1/1/1
        SysDescr:  VSP-7254XSQ (6.0.0.0_GA)
--------------------------------------------------------------------------------
Sys capability: O-Other; R-Repeater; B-Bridge; W-WLAN accesspoint; r-Router;
T-Telephone; D-DOCSIS cable device; S-Station only.
Total neighbors: 1
```

Variable Definitions

Use the data in the following table to use the **`show lldp neighbor`** command.

Variable	Value
port *{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}*	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port. Displays LLDP neighbor information on the specified port.

Viewing Fabric Attach Discovered Elements

Use this procedure to view Fabric Attach discovered elements.

About this task

When FA is enabled on an FA Server switch, LLDP PDUs are exchanged between the FA Server and FA Clients or FA Proxies. Standard LLDPs allow neighbors to be learned. With the help of organizational-specific element discovery TLVs, the client or proxy recognizes that it has attached to the FA Server. Only after the discovery handshake is complete, an FA Client or FA Proxy can transmit I-SID-to-VLAN assignments to join the SPB Fabric network through the FA Server.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Display FA discovered elements:

   ```
   show fa elements
   ```

3. Display FA discovered elements on a specific port:

   ```
   show fa elements [{slot/port[/sub-port] [-slot/port[/sub-port]]
   [,...]}]
   ```

Example

The following example displays the sample output for the `show fa elements` command.

```
Switch:1#show fa elements

================================================================================
                        Fabric Attach Discovery Elements
================================================================================
                         MGMT                                          ELEM ASGN
PORT   TYPE              VLAN STATE  SYSTEM ID                         AUTH AUTH
--------------------------------------------------------------------------------
1/5    proxy             710  T / S  50:61:84:ee:8c:00:20:00:00:01    AP   AP
1/6    proxy             710  T / S  50:61:84:ee:8c:00:20:00:00:01    AP   AP

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

```
                      Fabric Attach Authentication Detail
================================================================================
       ELEM OPER                      ASGN OPER
PORT   AUTH STATUS                    AUTH STATUS
--------------------------------------------------------------------------------
1/5    successAuth                    successAuth
1/6    successAuth                    successAuth

State Legend: (Tagging/AutoConfig)
T= Tagged,    U= Untagged,    D= Disabled,    S= Spbm,    V= Vlan,    I= Invalid

Auth Legend:
AP= Authentication Pass,  AF= Authentication Fail,
NA= Not Authenticated,  N= None



--------------------------------------------------------------------------------

2 out of 2 Total Num of fabric attach discovery elements displayed
```

Viewing Fabric Attach Statistics

If FA discovery fails, use this procedure to display FA statistics to determine if FA discovery TLVs were processed. You can also view the FA assignment statistics to determine the number of FA assignments that were accepted or rejected by the FA Server.

You can view the statistics at either the global level or at the port (interface) level.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. View global level FA statistics:

   ```
   show fa statistics [summary]
   ```

3. View FA statistics at the slot/port level:

   ```
   show fa statistics [{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}]
   ```

   Note:

   If a slot is removed from the switch chassis, the FA statistics are not displayed on the slot ports. When the slot is inserted back again, the statistics counters are reset.

4. **(Optional)** Clear FA statistics:

   ```
   clear fa statistics [summary] [{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}]
   ```

Examples

Viewing FA discovery and assignment statistics:

```
Switch:1>en
Switch:1#show fa statistics

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

```
                            Fabric Attach STATISTICS
================================================================================
Port      DiscElem   DiscElem   DiscElem   DiscAuth
          Received   Expired    Deleted    Failed
--------------------------------------------------------------------------------
1/1       3057       0          1          0
1/2       2000       0          1          0

================================================================================
                       Fabric Attach ASSIGNMENTS STATISTICS
================================================================================
Port      Asgn       Asgn       Asgn       Asgn       Asgn       AsgnAuth
          Received   Accepted   Rejected   Expired    Deleted    Failed
--------------------------------------------------------------------------------
1/1       3149       3          1          3          0          0
1/2       1500       0          1          2          0          0
```

View a summary of the FA discovery and assignment statistics:

```
Switch:1#show fa statistics summary

================================================================================
                            Fabric Attach STATISTICS SUMMARY
================================================================================
Port      DiscElem   DiscElem   DiscElem   DiscAuth
          Received   Expired    Deleted    Failed
--------------------------------------------------------------------------------
1/1       3057       0          1          0
1/2       2000       0          1          0

================================================================================
                       Fabric Attach ASSIGNMENTS STATISTICS SUMMARY
================================================================================
Port      Asgn       Asgn       Asgn       Asgn       Asgn       AsgnAuth
          Received   Accepted   Rejected   Expired    Deleted    Failed
--------------------------------------------------------------------------------
1/1       3149       3          1          3          0          0
1/2       1500       0          1          2          0          0
```

Viewing FA statistics on a specific port (port `1/1`):

```
Switch:1>en
Switch:1#show fa statistics 1/1

================================================================================
                            Fabric Attach STATISTICS
================================================================================
Port      DiscElem   DiscElem   DiscElem   DiscAuth
          Received   Expired    Deleted    Failed
--------------------------------------------------------------------------------
1/1       3057       0          1          0

================================================================================
                       Fabric Attach ASSIGNMENTS STATISTICS
================================================================================
Port      Asgn       Asgn       Asgn       Asgn       Asgn       AsgnAuth
          Received   Accepted   Rejected   Expired    Deleted    Failed
--------------------------------------------------------------------------------
1/1       3149       3          1          3          0          0
```

Optionally, clear FA statistics and verify that the statistics are cleared.

```
Switch:1#clear fa statistics
Switch:1#show fa statistics

================================================================================
                            Fabric Attach STATISTICS
================================================================================
Port      DiscElem   DiscElem   DiscElem   DiscAuth
          Received   Expired    Deleted    Failed
--------------------------------------------------------------------------------
1/1       0          0          0          0
1/2       0          0          0          0

================================================================================
                       Fabric Attach ASSIGNMENTS STATISTICS
================================================================================
Port      Asgn       Asgn       Asgn       Asgn       Asgn       AsgnAuth
          Received   Accepted   Rejected   Expired    Deleted    Failed
--------------------------------------------------------------------------------
1/1       0          0          0          0          0          0
1/2       0          0          0          0          0          0
```

Variable Definitions

Use the data in the following table to use the **`show fa statistics`** command.

Variable	Value
summary	Displays global level fabric attach statistics
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port. Displays Fabric Attach statistics on ports.

Use the data in the following table to use the **`clear fa statistics`** command.

Variable	Value
summary	Clears global level fabric attach statistics
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port. Clears Fabric Attach statistics on ports.

Viewing Fabric Attach I-SID-to-VLAN Assignments

Use this procedure to display the I-SID-to-VLAN assignments advertised by an FA Client or an FA Proxy, to be supported on the FA Server. These assignments can be accepted or rejected by the FA Server. An assignment that is successfully accepted by the FA Server results in the creation of a Switched UNI I-SID on the interface.

Before you begin

Verify that IS-IS and SPBM are properly configured on the FA Server switch.

- Verify SPBM configuration using the command **`show running-config module spbm`**.
- Verify IS-IS configuration using one of the following commands:
  - **`show isis`**
  - **`show isis interface`**
  - **`show isis adjacency`**
  - **`show isis lsdb`**

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Display FA I-SID-to-VLAN assignments:

   ```
   show fa assignment
   ```

3. Display FA I-SID-to-VLAN assignments on specific ports:

   ```
   show fa assignment [{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}]
   ```

Example

The following example displays a sample output for the `show fa assignment` command.

Note:

The state of I-SID-to-VLAN assignments on a client or proxy device is `pending` until it is changed by the FA Server to `active` or `reject`.

```
Switch:>en
Switch:1#show fa assignment
================================================================================
                          Fabric Attach Assignment Map
================================================================================
Interface  I-SID      Vlan       State      Origin
--------------------------------------------------------------------------------
1/1        2          2          active     proxy
1/2        3          3          active     proxy
1/2        4          4          active     proxy
1/3        5          5          reject     proxy
--------------------------------------------------------------------------------

 4 out of 4 Total Num of fabric attach assignment mappings displayed
```

```
--------------------------------------------------------------------------
```

Variable Definitions

Use the data in the following table to use the **`show fa assignment`** command.

Variable	Value
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Displaying Switched UNI (ELAN) I-SID Information

Use this procedure to display information on FA-created Switched UNI (ELAN) I-SIDs.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Display all Switched UNI (ELAN) I-SIDs:

   ```
   show i-sid elan
   ```

3. Display ELAN I-SID information on an MLT:

   ```
   show mlt i-sid [<1-512>]
   ```

   Note:

   Viewing ELAN I-SID information on an MLT is useful to understand the origin of the I-SID when multiple client or proxy devices connecting to the FA Server using SMLT MLT advertise the *same* I-SID-to-VLAN mappings. In the event of a link failure on an MLT, the origin of the I-SID helps determine on which MLT, and thereby from which proxy or client device, the mappings were successfully learnt.

4. Display ELAN I-SID information on ports:

   ```
   show interfaces gigabitEthernet i-sid [{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}]
   ```

Example

Display information on all Switched UNI (ELAN) I-SIDs.

The following sample output displays, for example, the I-SID information on one of the peer switches of the FA Server, in a dual-homed SMLT configuration.

```
Switch:1>en
Switch:1#show i-sid elan

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

```
                                        Isid Info
==========================================================================================
ISID        ISID                    PORT               MLT                ORIGIN
ID          TYPE       VLANID       INTERFACES         INTERFACES
------------------------------------------------------------------------------------------
2002        ELAN       N/A          c2002:1/10         -                  DISC_LOCAL
4000        ELAN       N/A          -                  c4000:1            DISC_BOTH
4001        ELAN       N/A          -                  c4001:1            DISC_LOCAL
4030        ELAN       N/A          -                  c4030:1            DISC_BOTH
4051        ELAN       N/A          -                  c4051:1            DISC_BOTH
10200       ELAN       N/A          -                  c200:1             DISC_REMOTE

c: customer vid    u: untagged-traffic

All 6 out of 6 Total Num of Elan i-sids displayed
```

Note:

The `I-SID TYPE` field displays once for each I-SID. The `I-SID TYPE` of an I-SID that is either learned through FA mapping assignments or configured as an FA management I-SID, is always ELAN. If a platform VLAN has the same I-SID value as that of the I-SID in an FA mapping assignment or in an FA management I-SID configuration, then the platform VLAN is associated with the I-SID endpoint and appears in the `VLANID` column.

Note:

- The `ORIGIN` field displays once for each I-SID. It indicates the origin of the I-SID and *not* the origin of the I-SID endpoint. To view the origin of the I-SID endpoints, execute either the `show mlt i-sid` or the `show interfaces gigabitEthernet i-sid` command.
  - The origin of I-SID `4000` displays as `DISC_BOTH`because it is discovered on both vIST peers.
  - The origin of I-SID `4001` displays as `DISC_LOCAL` because it is first discovered on the local FA Server switch.
  - The origin of I-SID `10200` displays as `DISC_REMOTE` because it is first discovered on the peer switch and then synchronized with the local switch.
- If the origin of an I-SID is DISC_LOCAL, DISC_REMOTE, DISC_BOTH or MANAGEMENT, it changes to CONFIG, after you manually configure an endpoint on the I-SID.

Display MLT I-SID information for MLT `1`.

In this sample output, the `ORIGIN` field indicates the origin of the I-SID endpoint.

```
Switch:1#show mlt i-sid 1

==============================================================================
                                 MLT Isid Info
==============================================================================
              ISID                    ISID
MLTID IFINDEX ID       VLANID C-VID  TYPE     ORIGIN           BPDU
------------------------------------------------------------------------------
1     6144    4000     N/A    4000   ELAN     DISC_BOTH
1     6144    4001     N/A    4001   ELAN     DISC_LOCAL
1     6144    4030     N/A    4030   ELAN     DISC_BOTH
1     6144    4051     N/A    4051   ELAN     DISC_BOTH
1     6144    10200    N/A    200    ELAN     DISC_REMOTE
```

```
--------------------------------------------------------------------------------

5 out of 6 Total Num of i-sid endpoints displayed
```

Display I-SID information on the port `1/10`:

In this sample output, the `ORIGIN` field indicates the origin of the I-SID endpoint.

Note:

The origin of I-SID `2002` displays as `DISC_LOCAL` because it is first discovered on the local FA Server switch. The FA Server switch synchronzies I-SID information only to the peer switch in the case of learning on a port interface.

```
Switch:1#show interfaces gigabitEthernet i-sid 1/10

================================================================================
                                PORT Isid Info
================================================================================
                ISID                    ISID
PORTNUM IFINDEX ID       VLANID C-VID   TYPE     ORIGIN           BPDU
--------------------------------------------------------------------------------
1/10    201     2002     N/A    601     ELAN     DISC_LOCAL
--------------------------------------------------------------------------------

1 out of 6 Total Num of i-sid endpoints displayed
```

Variable Definitions

Use the data in the following table to use the **`show i-sid`** command.

Variable	Value
elan	Displays all ELAN I-SIDs.

Use the data in the following table to use the **`show mlt i-sid`** command.

Variable	Value
<1–512>	The valid range for MLT ID.

Use the data in the following table to use the **`show interfaces gigabitEthernet i-sid`** command.

Variable	Value
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Troubleshooting Fabric Attach using the EDM

Verify configuration of Fabric Attach

Configuring Fabric Attach Globally

Use this procedure to configure FA globally or view existing FA global configuration.

Procedure

1. In the navigation pane, expand the **Configuration** > **Edit** folders.
2. Click **Fabric Attach**.
3. Click the **Globals** tab.
4. To enable or disable the Fabric Attach service, click **enabled** or **disabled** in the **Service** field.

   Caution:

   Disabling FA flushes all FA element discovery and mappings.
5. View the element type in the **ElementType** field.

   Note:

   The only supported element type is **faServer** (FA Server).
6. To specify the assignment time-out, enter a time-out value in seconds in the **AsgnTimeout** field.
7. View the provision mode in the **ProvisionMode** field.

   Note:

   The supported provision mode is **spbm**.
8. To specify the discovery time-out, enter a time-out value in seconds in the **DiscTimeout** field.
9. To clear the FA statistics, select the **Clear FA Statistics** checkbox.
10. To clear the error counters, select the check boxes **ClearErrorCounters** and/or **ClearGlobalErrorCounters**.
11. Click **Apply**.

Fabric Attach Globals Field Descriptions

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

Name	Description
Service	Enables or disables Fabric Attach service globally.

Table continues…

Name	Description
	The default is enable.
ElementType	Specifies the Fabric Attach element type. The supported element type is Fabric Attach Server.
AsgnTimeout	Specifies the Fabric Attach assignment time-out in seconds. The range is 45 to 480 seconds. The default is 240 seconds.
ProvisionMode	Specifies the Fabric Attach provision mode. The supported provision mode is SPB.
DiscTimeout	Specifies the Fabric Attach discovery time-out in seconds. The range is 45 to 480 seconds. The default is 240 seconds.
Clear FA Statistics	Clears Fabric Attach statistics.
ClearGlobalErrorCounters	Clears Fabric Attach global error counters. Disabled by default.

Configuring Fabric Attach Interface-level Settings

Use this procedure to configure FA interface-level settings or view existing interface-level settings.

You can enable Fabric Attach on a port, static MLT or an LACP MLT. Enabling FA on a port not only enables tagging but also disables spanning tree on that port. Enabling FA on an MLT enables FA on all ports of the MLT. When FA is enabled on ports in an MLT or LACP MLT, tagging is enabled and spanning tree is disabled on all those ports.

Before you begin

Ensure that FA is enabled globally on the switch.

About this task

Enabling FA on a port or MLT is necessary for element discovery. On the FA Server, FA is enabled globally by default. However, you must explicitly enable FA on a desired port or MLT interface, following which the FA Server can begin transmitting LLDP PDUs that contain the element discovery TLVs. This information is received by FA Client and FA Proxy devices which in turn also transmit their FA capabilities and settings. After the element handshake completes, the FA Server receives I-SID-to-VLAN assignment mappings from the connected client or proxy devices, on that port or MLT.

Procedure

1. In the navigation pane, expand the **Configuration** > **Edit** folders.
2. Click **Fabric Attach**.
3. Click the **Ports** tab.

   The FA interface-level settings are displayed.
4. To modify existing settings, double-click on the fields on this window. After making the required changes, click **Apply** to save your changes.
5. To configure FA on a new port or MLT interface:
   a. Click **Insert**.

The **Insert Ports** dialog box appears.

b. To configure FA on a port, enter a port number in the format slot/port[/sub-port], or click **Port** to select from a list of available ports.

c. To configure FA on an MLT, enter an MLT ID or click **Mlt** to select from a list of configured MLTs.

Note:

FA is successfully enabled on the MLT, only if all ports of the MLT have FA successfully enabled. Enabling FA enables LLDP on all ports. Tagging is enabled and spanning tree is disabled.

d. Click **Insert** to save your changes.

6. To remove (delete) FA on a port or MLT:

a. In the content pane, select a port or MLT from the list.

b. Click **Delete**.

Caution:

Removing FA on an interface flushes all FA element discovery and I-SID-to-VLAN mappings associated with that interface.

Ports Field Descriptions

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

Name	Description
IfIndex	Specifies the interface (port or MLT) on which Fabric Attach is configured.
State	Specifies the current state of the Fabric Attach port. It is either enabled or disabled. This field indicates whether LLDP PDUs (that include FA TLVs) are generated on the port (enabled) or not (disabled).
MsgAuthStatus	Specifies the Fabric Attach message authentication status on the port. It is either enabled or disabled.
MsgAuthKey	Specifies the Fabric Attach message authentication key for the associated port. The maximum length of this key is 32 characters.
MgmtIsid	Specifies the Fabric Attach management I-SID for the associated port. The range is 0 to 16777215. A zero value indicates that the management I-SID is not specified for the interface.
MgmtCvid	Specifies the Fabric Attach management customer VLAN ID (C-VID) for the interface. A zero value indicates that no C-VID is specified for the interface. Using the maximum configuration value for your switch indicates the port is untagged. Platform support determines the C-VID range.

Verify port-level transmission and reception

Viewing LLDP Reception Statistics

Use this procedure to view the LLDP reception statistics. You can also view these statistics graphically.

About this task

LLDP operates at the port interface level. Enabling FA on a port automatically enables LLDP transmission and reception on that port. It also enables traffic tagging and disables spanning tree on that port.

Enabling FA on an MLT enables LLDP transmission and reception on *all* ports in that MLT.

Note:

When a port is removed from an MLT, LLDP transmission on that port stops and spanning tree is enabled. Any I-SID-to-VLAN mappings on that port are removed, if not already learned on any other port in the MLT. This also causes the Switched UNI I-SID to be deleted from the MLT. If however, the mappings are learned on another port on the MLT, then the Switched UNI I-SID continues to exist for that MLT.

For ports in an LACP MLT, when FA is enabled, tagging is enabled on all ports in the LACP MLT. The consistency check for FA is based on key membership. If all ports with the same key do not support FA, FA is not successfully enabled on those ports.

Note:

If a slot is removed from the switch chassis, the statistics are not displayed on the slot ports. When the slot is inserted back again, the statistics counters are reset.

Procedure

1. In the navigation pane, expand the **Configuration** > **Edit** > **Diagnostics** > **802_1ab** folders.
2. Click **LLDP**.
3. Click the **RX Stats** tab.
4. To view the reception statistics graphically for a port:
   a. Select a row and click **Graph**.

      The **RX Stats-Graph,<port-number>** tab displays.

      You can view a graphical representation of the following data:

      - **FramesDiscardedTotal** — Total number of LLDP received frames that were discarded.
      - **FramesErrors** — Total number of erroneous LLDP frames received.
      - **FramesTotal** — Total number of frames received.
      - **TLVsDiscardedTotal** — Total number of received TLVs that were discarded.
      - **TLVsUnrecognizedTotal** — Total number of unrecognized TLVs received.

   b. Select one of the above parameters and click the appropriate icon on the top left-hand-side corner of the menu bar to draw a line chart, area chart, bar chart or a pie chart.

c. Click **Clear Counters** to clear the existing counters, and fix a reference point in time to restart the counters.

d. Click **Export**, to export the statistical data to a file.

e. To fix a poll interval, select an appropriate value from the **Poll Interval** drop-down list.

RX Stats Field Descriptions

Use the data in the following table to use the RX Stats tab.

Name	Description
PortNum	Specifies the port number.
FramesDiscardedTotal	Specifies the number of LLDP frames received on the port, but discarded, for any reason. This counter provides an indication of possible LLDP header formatting problems in the sending system, or LLDP PDU validation problems in the receiving system.
FramesErrors	Specifies the number of invalid LLDP frames received on the port.
FramesTotal	Specifies the total number of LLDP frames received on the port.
TLVsDiscardedTotal	Specifies the number of LLDP TLVs discarded on the port, for any reason.
TLVsUnrecognizedTotal	Specifies the number of LLDP TLVs on the port, that are unrecognized on that port. An unrecognized TLV is referred to as the TLV whose type value is in the range of reserved TLV types (000 1001–111 1110). An unrecognized TLV could be, for example, a basic management TLV from a later LLDP version.
AgeoutsTotal	Specifies the number of LLDP age-outs that occur on a specific port. An age-out is the number of times the complete set of information advertised by a particular MSAP is deleted, because the information timeliness interval has expired.

Graphing LLDP Reception Statistics

Use this procedure to graphically view the LLDP reception statistics.

About this task

LLDP operates at the port interface level. Enabling FA on a port automatically enables LLDP transmission and reception on that port. It also enables traffic tagging and disables spanning tree on that port.

Enabling FA on an MLT enables LLDP transmission and reception on *all* ports in that MLT.

Note:

When a port is removed from an MLT, LLDP transmission on that port stops and spanning tree is enabled. Any I-SID-to-VLAN mappings on that port are removed, if not already learned on any other port in the MLT. This also causes the Switched UNI I-SID to be deleted from the MLT. However, if the mappings are learned on another port on the MLT, the Switched UNI I-SID continues to exist for that MLT.

For ports in an LACP MLT, when FA is enabled, tagging is enabled on all ports in the LACP MLT. The consistency check for FA is based on key membership. If all ports with the same key do not support FA, FA is not successfully enabled on those ports.

Note:

If a slot is removed from the switch chassis, the statistics are not displayed on the slot ports. When the slot is inserted back again, the statistics counters are reset.

Procedure

1. In the navigation pane, expand the **Configuration** > **Edit** > **Diagnostics** > **802_1ab.LLDP** folders.
2. Click the **RX Stats** tab.
3. To view the reception statistics graphically for a port:
   a. Select a row, and click **Graph**.

      The system displays the **RX Stats-Graph,<port-number>** tab.

   b. Select one of the parameters, and click the appropriate icon in the upper-left corner of the menu bar to draw a line chart, area chart, bar chart or a pie chart.
4. To clear the existing counters, and fix a reference point in time to restart the counters, click **Clear Counters**.
5. To export the statistical data to a file, click **Export**.
6. To configure a poll interval, select an appropriate value from the **Poll Interval** drop-down list.

RX Stats Field Descriptions

Use the data in the following table to use the RX Stats tab.

Name	Description
FramesDiscardedTotal	Specifies the number of LLDP frames received on the port, but discarded, for any reason. This counter provides an indication of possible LLDP header formatting problems in the sending system, or LLDP PDU validation problems in the receiving system.
FramesErrors	Specifies the number of invalid LLDP frames received on the port.
FramesTotal	Specifies the total number of LLDP frames received on the port.
TLVsDiscardedTotal	Specifies the number of LLDP TLVs discarded on the port, for any reason.
TLVsUnrecognizedTotal	Specifies the number of LLDP TLVs on the port, that are unrecognized on that port.

Table continues…

Name	Description
	An unrecognized TLV is referred to as the TLV whose type value is in the range of reserved TLV types (000 1001–111 1110). An unrecognized TLV could be, for example, a basic management TLV from a later LLDP version.

Viewing LLDP Transmission Statistics

Use this procedure to view the LLDP transmission statistics. You can also view the statistics graphically.

About this task

LLDP operates at the port interface level. Enabling FA on a port automatically enables LLDP transmission and reception on that port. It also enables traffic tagging and disables spanning tree on that port.

Enabling FA on an MLT enables LLDP transmission and reception on *all* ports in that MLT.

Note:

When a port is removed from an MLT, LLDP transmission on that port stops and spanning tree is enabled. Any I-SID-to-VLAN mappings on that port are removed, if not already learned on any other port in the MLT. This also causes the Switched UNI I-SID to be deleted from the MLT. If however, the mappings are learned on another port on the MLT, then the Switched UNI I-SID continues to exist for that MLT.

For ports in an LACP MLT, when FA is enabled, tagging is enabled on all ports in the LACP MLT. The consistency check for FA is based on key membership. If all ports with the same key do not support FA, FA is not successfully enabled on those ports.

Note:

If a slot is removed from the switch chassis, the statistics are not displayed on the slot ports. When the slot is inserted back again, the statistics counters are reset.

Procedure

1. In the navigation pane, expand the **Configuration** > **Edit** > **Diagnostics** > **802_1ab** folders.
2. Click **LLDP**.
3. Click the **TX Stats** tab.

   The transmission statistics are displayed.
4. To view the transmission statistics graphically for a port:
   a. In the content pane (on the right-hand-side), select a row and click the **Graph** button.

      The **TX Stats-Graph,<port-number>** tab displays.

      You can view a graphical representation of the LLDP frames transmitted (**FramesTotal**), for the following parameters:

      - AbsoluteValue

- Cumulative
- Average/sec
- Minimum/sec
- Maximum/sec
- LastVal/sec

b. To view the graph, select one of the above parameters and click the appropriate icon on the top left-hand-side of the menu bar to draw a line chart, area chart, bar chart or a pie chart.

c. Click **Clear Counters** to clear the existing counters, and fix a reference point in time to restart the counters.

d. Click **Export**, to export the statistical data to a file.

e. To fix a poll interval, select an appropriate value from the **Poll Interval** drop-down list.

TX Stats Field Descriptions

Use the data in the following table to use the TX Stats tab.

Name	Description
PortNum	Specifies the port number.
FramesTotal	Specifies the total number of LLDP frames transmitted.

Graphing LLDP Transmission Statistics

Use this procedure to view the LLDP transmission (TX) statistics. You can also view the statistics graphically.

About this task

LLDP operates at the port interface level. Enabling FA on a port automatically enables LLDP transmission and reception on that port. It also enables traffic tagging and disables spanning tree on that port.

Enabling FA on an MLT enables LLDP transmission and reception on all ports in that MLT.

Note:

When a port is removed from an MLT, LLDP transmission on that port stops and spanning tree is enabled. Any I-SID-to-VLAN mappings on that port are removed, if not already learned on any other port in the MLT. This also causes the Switched UNI I-SID to be deleted from the MLT. However, if the mappings are learned on another port on the MLT,the Switched UNI I-SID continues to exist for that MLT.

For ports in an LACP MLT, when FA is enabled, tagging is enabled on all ports in the LACP MLT. The consistency check for FA is based on key membership. If all ports with the same key do not support FA, FA is not successfully enabled on those ports.

Note:

If a slot is removed from the switch chassis, the statistics are not displayed on the slot ports. When the slot is inserted back again, the statistics counters are reset.

Procedure

1. In the navigation pane, expand the **Configuration** > **Edit** > **Diagnostics** > **802_1ab.LLDP** folders.
2. Click the **TX Stats** tab.

   The system displays the transmission statistics.
3. To view the transmission statistics graphically for a port:
   a. Select a row, and click **Graph**.

      The system displays the **TX Stats-Graph,<port-number>** tab.
   b. To view the graph, select one of the parameters, and click the appropriate icon on the upper-left corner of the menu bar to draw a line chart, area chart, bar chart or a pie chart.
4. To clear the existing counters, and fix a reference point in time to restart the counters, click **Clear Counters**.
5. To export the statistical data to a file, click **Export**.
6. To fix a poll interval, select an appropriate value from the **Poll Interval** drop-down list.

TX Stats Field Descriptions

Use the data in the following table to use the TX Stats tab.

Name	Description
FramesTotal	Specifies the total number of LLDP frames transmitted.

Viewing Global FA Statistics Graphically

Use this procedure to view the global FA statistics graphically.

Procedure

1. In the navigation pane, expand the **Configuration** > **Graph** folders.
2. Click **Chassis**.
3. Click the **Fabric Attach** tab.
4. To view a graphical representation of the statistics, select a row and click the appropriate icon on the top left-hand-side of the menu bar to draw a line chart, area chart, bar chart or a pie chart.
5. Click **Clear Counters** to clear the existing counters, and fix a reference point in time to restart the counters.

6. Click **Export**, to export the statistical data to a file.
7. To fix a poll interval, select an appropriate value from the **Poll Interval** drop-down list.

Fabric Attach field descriptions

Use the data in the following table to use the Fabric Attach tab.

Name	Description
DiscElemReceived	Specifies the number of discovery elements received globally.
AsgnReceived	Specifies the number of remote I-SID-to-VLAN assignments received globally.
AsgnAccepted	Specifies the number of remote I-SID-to-VLAN assignments accepted globally.
AsgnRejected	Specifies the number of remote I-SID-to-VLAN assignments rejected globally.
AsgnExpired	Specifies the number of remote I-SID-to-VLAN assignments that expired globally.
AuthFailed	Specifies the number of authentications that failed globally.
DiscAuthFailed	Specifies the number of discovery authentications that failed globally.
DiscElemExpired	Specifies the number of discovery elements that expired globally.
DiscElemDeleted	Specifies the number of discovery elements that were deleted globally.
AsgnDeleted	Specifies the number of remote assignments that were deleted globally.

Viewing FA Port Statistics Graphically

Use this procedure to view the FA port statistics graphically.

Before you begin

Ensure that a switch port is selected in the **Device Physical View** tab.

Procedure

1. In the navigation pane, expand the **Graph** > **Port** folders.
2. Click the **Fabric Attach** tab.

   The FA port statistics are displayed.
3. To view a graphical representation of the port statistics, select a row and click the appropriate icon on the top left-hand-side of the menu bar to draw a line chart, area chart, bar chart or a pie chart.
4. Click **Clear Counters** to clear the existing counters, and fix a reference point in time to restart the counters.

5. Click **Export**, to export the statistical data to a file.
6. To fix a poll interval, select an appropriate value from the **Poll Interval** drop-down list.

Fabric Attach Field Descriptions

Use the data in the following table to use the Fabric Attach tab.

Name	Description
DiscElemReceived	Specifies the number of discovery elements received on a given port.
AsgnReceived	Specifies the number of remote I-SID-to-VLAN assignments received on a given port.
AsgnAccepted	Specifies the number of remote I-SID-to-VLAN assignments accepted on a given port.
AsgnRejected	Specifies the number of remote I-SID-to-VLAN assignments rejected on a given port.
AsgnExpired	Specifies the number of remote I-SID-to-VLAN assignments that expired on a given port.
AuthFailed	Indicates the number of received TLVs for which authentication was attempted and failed on the identified port.
DiscElemExpired	Specifies the number of discovery elements that expired on a given port.
DiscElemDeleted	Specifies the number of discovery elements that were deleted on a given port.
AsgnDeleted	Specifies the number of remote assignments that were deleted on a given port.
AsgnAuthFailed	Specifies the number of remote assignment authentications that failed on a given port.

Verify FA discovery and I-SID-to-VLAN mapping assignments

Viewing LLDP Neighbor Information

Use this procedure to view the LLDP neighbor information.

Procedure

1. In the navigation pane, expand the **Configuration** > **Edit** > **Diagnostics** > **802_1ab** folders.
2. Click **LLDP**.
3. Click the **Neighbor** tab.

Neighbor Field Descriptions

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

Name	Description
TimeMark	Indicates the time filter. For more information about TimeFilter, see the TimeFilter textual convention in IETF RFC 2021.

Table continues…

Name	Description
LocalPortNum	Identifies the port on which the remote system information is received.
Index	Indicates a particular connection instance that is unique to the remote system.
ProtocolType	Indicates whether the entry protocol is CDP or LLDP.
SysName	Indicates the name of the remote system.
IpAddress	Indicates the neighbor's IP address.
PortIdSubType	Indicates the type of encoding used to identify the remote port.
PortId	Indicates the remote port ID.
PortDesc	Indicates the remote port description.
ChassisIdSubtype	Indicates the type of encoding used to identify the remote system chassis. • chassisComponent • interfaceAlias • portComponent • macAddress • networkAddress • interfaceName • local
ChassisId	Indicates the chassis ID of the remote system.
SysCapSupported	Identifies the system capabilities supported on the remote system.
SysCapEnabled	Identifies the system capabilities enabled on the remote system.
SysDesc	Indicates the description of the remote system.

Configuring Fabric Attach I-SID-to-VLAN Assignments

Use this procedure to view or configure FA I-SID-to-VLAN assignment information.

Procedure

1. In the navigation pane, expand the **Configuration** > **Edit** folders.
2. Click **Fabric Attach**.
3. Click the **Assignment** tab.
4. If you make configuration changes, click **Apply** to save changes.

Assignments Field Descriptions

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

Name	Description
IfIndex	Specifies the interface identifier of the I-SID-to-VLAN assignment.
Isid	Specifies the I-SID value of the I-SID-to-VLAN assignment.

Table continues…

Name	Description
Vlan	Specifies the VLAN ID component of the I-SID-to-VLAN assignment.
State	Specifies the current state of the I-SID-to-VLAN assignment. It can be one of the following values: • Other • Pending • Active • Rejected
Origin	Specifies the origin information of the I-SID-to-VLAN assignment.

Fabric Attach troubleshooting example

Troubleshooting FA Server Rejection of I-SID-to-VLAN Assignments Using Trace

Consider an FA solution where the FA Server receives I-SID-to-VLAN assignment requests from a proxy device and some of these assignment requests are rejected by the FA Server. Use this procedure to help you troubleshoot the cause of the rejection.

Note:

When the FA Server rejects an I-SID-to-VLAN assignment request, the error message in the log file lists a generic reason for the failure, such as `rejected due to application error (status 9)`. To troubleshoot further, you must use trace.

This procedure also demonstrates how you can configure trace for enhanced troubleshooting.

Procedure

Begin troubleshooting on the FA Server

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Verify that router IS-IS is enabled. This is required for proper FA operation.

   ```
   show isis
   ```

   Note:

   I-SID-to-VLAN assignments are *always* rejected if router IS-IS is disabled.

3. Verify that FA is enabled on the interface on which I-SID-to-VLAN assignments are expected.

   ```
   show fa interface [disabled-auth] [enabled-auth] [mlt <1-512>] [port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}]
   ```

4. Verify the discovery and authentication status of the proxy device, on the interface.

```
show fa elements [{slot/port[/sub-port] [-slot/port[/sub-port]]
[,...]}]
```

5. Determine the I-SID-to-VLAN assignments received on the interface and which ones are rejected.

   ```
   show fa assignment [{slot/port[/sub-port] [-slot/port[/sub-port]]
   [,...]}]]
   ```

6. View the log file to determine the cause of the assignment rejection.

   ```
   show log file module fa
   ```

   Note:

   When the FA Server rejects an I-SID-to-VLAN assignment request, only a generic reason for the rejection is logged.

Enhanced troubleshooting using trace

7. Configure trace:

   a. Enable keyword search in the trace output:

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

   b. Set the trace level for FA:

      ```
      trace level <Module_ID>
      ```

      Note:

      - *<Module_ID>* specifies the module for the trace. Different hardware platforms support different ID ranges because of feature support differences. To see which module IDs are available on the switch, use the **`show trace modid-list`** command or CLI command completion Help.
      - FA uses the trace level `221`.

   c. Turn on trace:

      ```
      trace screen [enable]|[disable]
      ```

Example:

The following example *simulates* a configuration error on the FA Server as a result of which the FA Server rejects I-SID-to-VLAN assignments from the proxy device. When the FA Server rejects an I-SID-to-VLAN assignment, the error message listed in the log file is a generic reason for the rejection, as demonstrated in this example. To troubleshoot further, set up trace.

On the FA Server, assume that the interface MLT 1 consists of ports `1/5` and `1/6`. Assume that a proxy device sends I-SID-to-VLAN assignment mapping requests with I-SID 9005 and CVID 400, on this interface.

Simulate a configuration error on the FA Server:

Configure a management I-SID with a C-VID value that is different from that of the C-VID in the I-SID-to-VLAN assignment request from the proxy. So, for example, configure a management I-SID

with C-VID 999, which is different from the C-VID advertised by the proxy, which is 400. This causes rejection of I-SID-to-VLAN assignment requests on the interface.

```
Switch:1>en
Switch:1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface mlt 1
Switch:1(config-mlt)#no fa enable
Switch:1(config-mlt)#fa management i-sid 9005 c-vid 999
Switch:1(config-mlt)#fa enable
Switch:1(config-mlt)#exit
Switch:1(config)#exit
```

At this stage, the FA Server rejects I-SID-to-VLAN assignments as shown below.

```
Switch:1#show fa assignment

================================================================================
                          Fabric Attach Assignment Map
================================================================================
Interface  I-SID      Vlan       State      Origin
--------------------------------------------------------------------------------
1/5        312        710        active     proxy
1/5        9005       400        reject     proxy
1/6        312        710        active     proxy
1/6        9005       400        reject     proxy
--------------------------------------------------------------------------------
4 out of 4 Total Num of fabric attach assignment mappings displayed
--------------------------------------------------------------------------------
```

Begin troubleshooting on the FA Server:

Verify that IS-IS is enabled.

```
Switch:1>en
Switch:1#show isis

================================================================================
                              ISIS General Info
================================================================================
                            AdminState : enabled
                            RouterType : Level 1
                             System ID : 8404.bcb1.0043
                  Max LSP Gen Interval : 900
                                Metric : wide
                   Overload-on-startup : 20
                              Overload : false
                           Csnp Interval : 10
                           PSNP Interval : 2
                     Rxmt LSP Interval : 5
                              spf-delay : 100
                            Router Name : FAServer
                     ip source-address : 43.43.43.43
                   ipv6 source-address : 1:43:0:0:0:0:0:43
               ip tunnel source-address : 12.43.43.43
                             Tunnel vrf : 12
                          ip tunnel mtu :
                     Num of Interfaces : 4
                  Num of Area Addresses : 1
                         inband-mgmt-ip :
                               backbone : disabled
               Dynamically Learned Area : 00.0000.0000
                             FAN Member : No
Switch:1#
```

Verify that FA is enabled on the interface `MLT 1`, on which the I-SID-to-VLAN assignments are expected. View the `SERVER STATUS` field.

```
Switch:1#show fa interface mlt 1

================================================================
                          Fabric Attach Interfaces
================================================================
INTERFACE    SERVER   MGMT     MGMT     MSG AUTH MSG AUTH
             STATUS   ISID     CVID     STATUS   KEY
----------------------------------------------------------------
Mlt1         enabled  0        0        enabled  ****

----------------------------------------------------------------
1 out of 1 Total Num of fabric attach interfaces displayed
----------------------------------------------------------------
```

Verify the discovery and authentication status of the proxy device on the interface. Note that the proxy is successfully discovered and authenticated on ports `1/5` and `1/6` of `MLT 1`.

```
Switch:1#show fa elements

================================================================================
                       Fabric Attach Discovery Elements
================================================================================
                        MGMT                                          ELEM ASGN
PORT   TYPE             VLAN STATE  SYSTEM ID                         AUTH AUTH
--------------------------------------------------------------------------------
1/5    proxy            1    T / S  10:cd:ae:09:40:00:20:00:00:01    AP   AP
1/6    proxy            1    T / S  10:cd:ae:09:40:00:20:00:00:01    AP   AP

================================================================================
                     Fabric Attach Authentication Detail
================================================================================
       ELEM OPER                    ASGN OPER
PORT   AUTH STATUS                  AUTH STATUS
--------------------------------------------------------------------------------
1/5    successAuth                  successAuth
1/6    successAuth                  successAuth


State Legend: (Tagging/AutoConfig)
T= Tagged,    U= Untagged,    D= Disabled,    S= Spbm,    V= Vlan,    I= Invalid


Auth Legend:
AP= Authentication Pass,  AF= Authentication Fail,
NA= Not Authenticated,  N= None



--------------------------------------------------------------------------------

2 out of 2 Total Num of fabric attach discovery elements displayed

--------------------------------------------------------------------------------
```

View the log file to determine the cause of the rejection. The log file displays the generic error `rejected due to application error (status 9)` as follows:

```
Switch:1#show log file module fa
...
CP1  [12/04/15 00:45:51.185:UTC] 0x00374583 00000000 GlobalRouter FA INFO Fabric Attach
```

```
Element Discovered on interface 1/5 Element type proxy (3) Id 50:61:84:ee:8c:
00:20:00:00:01 CP1  [12/04/15 00:45:51.187:UTC] 0x0037458f 00000000 GlobalRouter FA INFO
Fabric Attach Assignment rejected: interface 1/5 i-sid 9005 cvid 400 rejected due to
application error (status 9)
...
...
```

To troubleshoot further, use trace.

```
Switch:1#trace grep fa
Switch:1#trace level 221 3
Switch:1#trace screen enable
Screen tracing is on
```

View the trace output. The trace output displays that the error was caused because the FA interface (`MLT 1`) was configured with a different C-VID for I-SID 9005.

```
Switch:1#0:07:57.801252 1  fa.c                  :858 [lcy-ve][12898-13062]cbcp-
main.x:faUpdateSwitchedUni      :FA: faUpdateSwitchedUni port 196 isid 9005 cvid 400
0:07:57.801283 1  fa_swuni.c            :2900[lcy-ve][12898-13062]cbcp-
main.x:faUpdateSwitchedUniCheck :FA: Call faUpdateSwitchedUniCheckSmlt for mlt 1
0:07:57.801644 1  fa_swuni.c            :2421[lcy-ve][12898-13062]cbcp-
main.x:faSwitchedUniCheckEndpointParms:FA: Failed
rcIsidElanEndPointTblConsistencyCheckCommon for Ifindex 6144 Isid 9005 Cvid 400 error
Switched UNI/Fabric Attach MLT cannot be configured for different c-vid for same I-SID
0:07:57.802074 1  fa.c                  :858 [lcy-ve][12898-13062]cbcp-
main.x:faUpdateSwitchedUni      :FA: faUpdateSwitchedUni port 197 isid 9005 cvid 400
0:07:57.802086 1  fa_swuni.c            :2900[lcy-ve][12898-13062]cbcp-
main.x:faUpdateSwitchedUniCheck :FA: Call faUpdateSwitchedUniCheckSmlt for mlt 1
0:07:57.802276 1  fa_swuni.c            :2421[lcy-ve][12898-13062]cbcp-
main.x:faSwitchedUniCheckEndpointParms:FA: Failed
rcIsidElanEndPointTblConsistencyCheckCommon for Ifindex 6144 Isid 9005 Cvid 400 error
Switched UNI/Fabric Attach MLT cannot be configured for different c-vid for same I-SID”
```

Troubleshooting FAN Transit

Use the following section to troubleshoot Fabric Area Network (FAN) Transit information on a switch.

Viewing FAN Transit information - Detailed

About this task

Use this procedure to verify detailed FAN Transit information of a switch acting as a transit node in a FAN. Transit nodes are not members of the FAN and only forward FAN traffic.

This procedure also includes verification on the FAN member nodes.

Procedure

Verification on FAN member nodes:

1. Verify that the node is a FAN member:

   ```
   show isis
   ```

2. Verify that the node signals FAN membership on TLV 147:

```
show isis lsdb local tlv 147 detail
```

3. Verify that the node creates FAN multicast FIB entries for itself and the other member nodes:

   ```
   show isis spbm multicast-fib
   ```

4. Verify that the transit node is a part of the FAN tree:

   ```
   l2 tracetree-fan
   ```

5. Verify that the transit node is in the SPB path between the FAN member nodes:

   ```
   show isis spbm unicast-tree <2-4059>
   ```

Verification on the transit node:

6. Verify that the transit node is not a member of the FAN and does not signal FAN membership on TLV 147:

   ```
   show isis
   ```

   ```
   show isis lsdb local tlv 147 detail
   ```

7. Verify that the transit node creates FAN multicast FIB entries for the FAN member nodes:

   ```
   show isis spbm multicast-fib
   ```

Example

Use the following example to verify FAN Transit information on a transit node between two end-point SPB nodes, BEB-A and BEB-B.

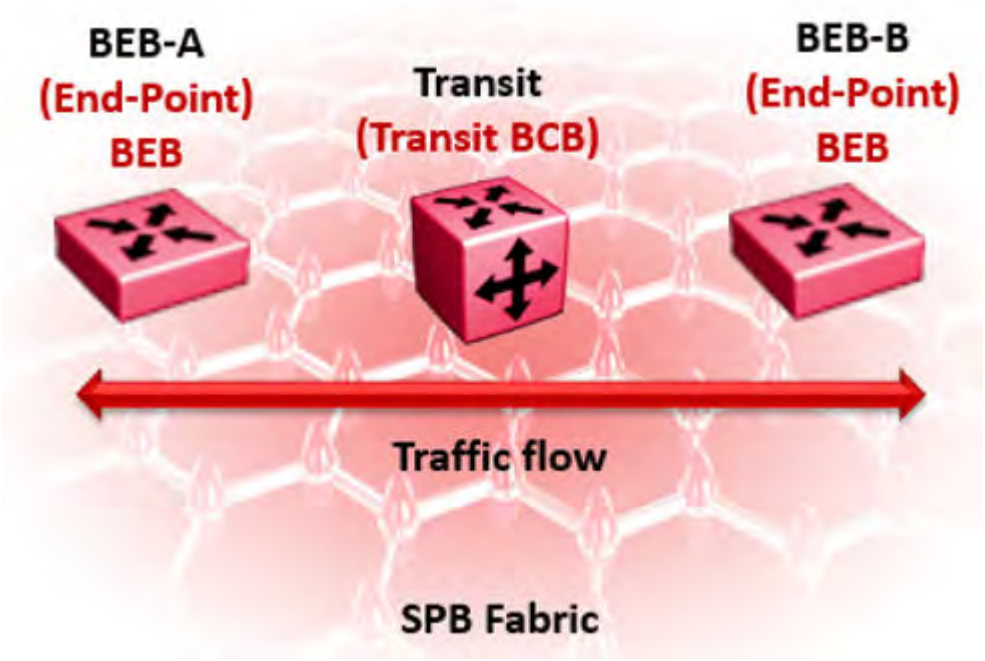

Verification on BEB-A:

Note:

You can execute the following commands on either BEB-A or BEB-B. The following example displays the sample outputs for BEB-A.

Verify that BEB-A is a member of the FAN. The `FAN Member` attribute displays as `Yes`.

```
BEB-A:1#show isis
================================================================================
                              ISIS General Info
================================================================================
                            AdminState : enabled
                            RouterType : Level 1
                             System ID : b0ad.aa41.9c84
                 Max LSP Gen Interval : 900
                               Metric : wide
                Overload-on-startup : 20
                            Overload : false
                          Csnp Interval : 10
                          PSNP Interval : 2
                    Rxmt LSP Interval : 5
                             spf-delay : 100
                           Router Name : BEB-A
                    ip source-address :
                  ipv6 source-address :
             ip tunnel source-address :
                            Tunnel vrf :
                         ip tunnel mtu :
                    Num of Interfaces : 1
               Num of Area Addresses : 1
                          inband-mgmt-ip :
                                backbone : disabled
             Dynamically Learned Area : 00.0000.0000
                            FAN Member : Yes
```

Verify that BEB-A signals FAN membership. FAN membership is signaled when TLV 147 displays the FAN multicast address.

```
BEB-A:1#show isis lsdb local tlv 147 detail

================================================================================
                              ISIS LSDB (DETAIL)
================================================================================
--------------------------------------------------------------------------------

Level-1 LspID: b0ad.aa41.9c84.00-00     SeqNum: 0x00000010      Lifetime:   412
        Chksum: 0xac95  PDU Length: 122
        Host_name: BEB-A
        Attributes: IS-Type 1
TLV:147 Chassis MAC: b0:ad:aa:41:9c:00

TLV:147 FAN Mcast Addr: b1:ad:aa:41:9c:84
```

Verify that BEB-A creates FAN multicast FIB entries for itself and other FAN nodes on ISID 16777001:

```
BEB-A:1#show isis spbm multicast-fib
================================================================================================================
                                        SPBM MULTICAST FIB ENTRY INFO
================================================================================================================
MCAST DA           ISID      BVLAN SYSID           HOST-NAME           OUTGOING-INTERFACES    INCOMING
                                                                                              INTERFACE
----------------------------------------------------------------------------------------------------------------
65:6a:52:ce:04:84  16777001  4058  646a.52ce.0484  BEB-B                                       3/20
65:6a:52:ce:04:84  16777001  4059  646a.52ce.0484  BEB-B                                        /20
b1:ad:aa:41:9c:84  16777001  4058  b0ad.aa41.9c84  BEB-A                3/20                   cpp
```

```
-------------------------------------------------------------------------------------------------------
Total number of SPBM MULTICAST FIB entries 3
```

Verify that the transit node is in the SPB path between the FAN end-points BEB-A and BEB-B. `4058` and `4059` are the SPB B-VIDs.

Note:

You can view the SPB B-VIDs by executing the command **`show isis spbm`**.

```
BEB-A:1#show isis spbm unicast-tree 4058
Node:646a.52ce.0484 (BEB-B) -> Node:f873.a202.53df (TRANSIT) -> ROOT
Node:f873.a202.53df (TRANSIT) -> ROOT
```

```
BEB-A:1#show isis spbm unicast-tree 4059
Node:646a.52ce.0484 (BEB-B) -> Node:f873.a202.53df (TRANSIT) -> ROOT
Node:f873.a202.53df (TRANSIT) -> ROOT
```

Verify that the transit node is a part of the FAN tree:

```
BEB-A:1#l2 tracetree-fan

Please wait for l2tracetree to complete or press any key to abort

l2tracetree to b1:ad:aa:41:9c:84, vlan 4058 i-sid 16777001 nickname 0.11.03 hops 64
1   BEB-A           b0:ad:aa:41:9c:84 -> TRANSIT          f8:73:a2:02:53:df
2   TRANSIT         f8:73:a2:02:53:df -> BEB-B            64:6a:52:ce:04:84
```

Verification on the Transit node:

Verify that the transit node is not a member of the FAN. The `FAN Member` parameter does not display.

```
TRANSIT:1#show isis
===============================================================================
                              ISIS General Info
===============================================================================
                           AdminState : enabled
                           RouterType : Level 1
                            System ID : f873.a202.53df
                  Max LSP Gen Interval : 900
                               Metric : wide
                 Overload-on-startup : 20
                             Overload : false
                          Csnp Interval : 10
                          PSNP Interval : 2
                       Rxmt LSP Interval : 5
                              spf-delay : 100
                            Router Name : TRANSIT
                       ip source-address :
              ip tunnel source-address :
                           Tunnel vrf :
                        ip tunnel mtu :
                    Num of Interfaces : 2
                Num of Area Addresses : 1
                       inband-mgmt-ip :
                             backbone : disabled
```

Verify that the transit node does not signal FAN membership on TLV 147. The `FAN Mcast Addr` parameter is not displayed.

```
TRANSIT:1#show isis lsdb local tlv 147 detail

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

```
                                ISIS LSDB (DETAIL)
================================================================================================
------------------------------------------------------------------------------------------------

Level-1 LspID: f873.a202.53df.00-00     SeqNum: 0x00000013      Lifetime:  1149
        Chksum: 0x424d  PDU Length: 135
        Host_name: TRANSIT
        Attributes:     IS-Type 1
TLV:147 Chassis MAC: f8:73:a2:02:50:00
```

Verify the transit node creates the FAN Multicast FIB entries for the FAN end-points on ISID 16777001.

```
TRANSIT:1#show isis spbm multicast-fib
================================================================================================================
                                            SPBM MULTICAST FIB ENTRY INFO
================================================================================================================
MCAST DA           ISID      BVLAN SYSID           HOST-NAME            OUTGOING-INTERFACES    INCOMING
                                                                                               INTERFACE
----------------------------------------------------------------------------------------------------------------
65:6a:52:ce:04:84  16777001  4058  646a.52ce.0484  BEB-B                 5/20                    5/9
65:6a:52:ce:04:84  16777001  4059  646a.52ce.0484  BEB-B                 5/20                    5/9
b1:ad:aa:41:9c:84  16777001  4058  b0ad.aa41.9c84  BEB-A                 5/9                     5/20
b1:ad:aa:41:9c:84  16777001  4059  b0ad.aa41.9c84  BEB-A                 5/9                     5/20

----------------------------------------------------------------------------------------------------------------
Total number of SPBM MULTICAST FIB entries 4
----------------------------------------------------------------------------------------------------------------
```

Chapter 10: Upper Layer Troubleshooting

Use the information in this chapter to troubleshoot Layer 4 to 7 applications.

This chapter includes the following sections:

- Troubleshooting SNMP
- Troubleshooting DHCP
- Troubleshooting BGP configuration

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).

If you enable enhanced secure mode, the switch does not support the default SNMPv1 and default SNMPv2 community strings, and default SNMPv3 user name. The individual in the administrator access level role can configure a non-default value for the community strings, and the switch can continue to support SNMPv1 and SNMPv2. The individual in the administrator access level role can also configure a non-default value for the SNMPv3 user name and the switch can continue to support SNMPv3. If you disable enhanced secure mode, the SNMPv1 and SNMPv2 support for community strings remains the same, and the default SNMPv3 user name remains the same. Enhanced secure mode is disabled by default.

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.

SNMP Trap not Received

Perform the following procedure to troubleshoot issues in which an SNMP trap is not received.

Procedure

1. Enter Global Configuration mode:

   ```
   enable

   configure terminal
   ```

2. Show the BPDU Guard status for the port:

   ```
   show spanning-tree bpduguard {slot/port[/sub-port][-slot/port[/sub-port]][,...]}
   ```

3. Configure the correct SNMP target information:

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

Example

In the following example, BPDU guard is enabled on port 1/8, BPDU packets are received, port 1/8 is disabled, and the TimerCount is incrementing, but no SNMP trap is ever received.

```
Switch:1>enable
Switch:1#show spanning-tree bpduguard 1/8

================================================================================
                                  Bpdu Guard
================================================================================
Port           PORT        PORT                TIMER BPDUGUARD
NUM      MLTID ADMIN_STATE OPER_STATE  TIMEOUT  COUNT ADMIN_STATE
--------------------------------------------------------------------------------
1/8            Down        Down        120      0     Disabled
```

Variable Definitions

Use the data in the following table to use the **`show spanning-tree`** command.

Variable	Value
{slot/port[/sub-port][-slot/port[/sub-port]][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

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

Variable	Value
filter *WORD<1-32>*	Specifies a filter profile name.
host *WORD<1-256>*	Specifies the IPv4 or IPv6 host address
inform [timeout *<1-2147483647>*]	Specifies the notify type. The optional timeout parameter configures the timeout value, which specifies the time to wait for a reply before resending the inform message. Time is specified in centiseconds
noAuthNoPriv|authNoPriv|authPriv *WORD<1-32>*	Specifies the security level.
port *<1-65535>*	Specifies the port number that will be set as the destination port at the UDP level in the trap packet.
retries *<0-255>*	Specifies the number of packets to be sent if no reply is received.
{slot/port[/sub-port][-slot/port[/sub-port]][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

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[/sub-port]
[-slot/port[/sub-port]][,...]}]
```

Note:

Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-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

```
Switch:1# ping 192.0.2.31
```

```
Switch:1:#show spanning-tree mstp port role
================================================================================
                          CIST Port Roles and States
================================================================================
Port-Index  Port-Role    Port-State   PortSTPStatus  PortOperStatus
--------------------------------------------------------------------------------
1/1         Disabled     Forwarding   Disabled       Disabled
1/2         Disabled     Forwarding   Disabled       Disabled
1/3         Disabled     Discarding   Enabled        Disabled
1/4         Disabled     Discarding   Enabled        Disabled
1/5         Disabled     Forwarding   Disabled       Disabled
1/6         Disabled     Forwarding   Disabled       Disabled
```

```
1/7         Disabled     Forwarding   Disabled       Disabled
1/8         Disabled     Forwarding   Disabled       Disabled
1/9         Disabled     Discarding   Enabled        Disabled
1/10        Disabled     Discarding   Enabled        Disabled
1/11        Disabled     Discarding   Enabled        Disabled
1/12        Designated   Forwarding   Enabled        Enabled
1/13        Disabled     Forwarding   Disabled       Disabled
1/14        Disabled     Forwarding   Disabled       Disabled

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

```
Switch:1:# trace level 9 3
```

```
Switch:1:# 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.

Variable	Value
{slot/port[/sub-port][-slot/port[/sub-port]][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
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 switch only participates in forwarding the requests and replies to and from the client and the DHCP server. The switch always acts as the relay agent, on which you configure the forward path to the server.

To troubleshoot DHCP Relay issues on the switch, use the following procedure.

Procedure

1. Verify that the DHCP server is reachable using ping. If ping is working and the DHCP server is reachable, DHCP should work.
2. Verify that the relay agents and the forward path configured are reachable. Ping the server and the gateway to the server.
3. Check that the relay agent configurations are correct. Also verify that DHCP is enabled on the switch:

   ```
   show ipv6 dhcp-relay interface {gigabitEthernet {slot/port[/sub-port][-slot/port[/sub-port]][,...]}|vlan <1-4059>
   ```

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 <1-4059>
   ```

6. In a scenario with VRRP and SMLT, configure VRRP IP as the DHCP relay agent.
7. When using the VRRP VRID as the relay agent, make sure the VRRP configurations are proper.
8. To verify that relay forward and relay receive are working, enable trace for DHCP with IPv6, and grep trace for relay:

   ```
   trace level 66 3
   ```

   ```
   trace grep relay
   ```

   ```
   trace screen enable
   ```

9. Display the count of DHCP Relay requests and replies to verify the system received requests and replies:

   ```
   show ipv6 dhcp-relay counters
   ```

IPv6 DHCP Relay Server Side Troubleshooting

Use the following procedure to troubleshoot IPv6 DHCP Relay on the server side.

Procedure

1. Enable the services on the server side, and then create an IP pool.

   The IP pool must contain the range of addresses that you want to assign to the clients.

   Configure the IP pool with the same network subnet as that of the relay agent.
2. When the configuration is complete, initiate a DHCP request from a client.
3. Check the log file available on the server to verify the reason for packet drop.
4. Capture the packets on the server side using Ethereal.
5. From the server side, use ping to verify that the relay agent address is reachable.

   Ensure that a route to the relay is configured.
6. For more configuration aspects, see the Microsoft webpage for troubleshooting and configuration issues.

   Note:

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

   Example 1: if you receive the message `0x00108796 (relayMsgSend): cannot find route entry for destination` on the console, you must ping the server. If the server is not reachable, the system cannot forward the packet. This is not a DHCP issue.

   Example 2: if you receive the message `0x00108705` this indicates a problem at the transmission level. Check the server reachability and ensure that MAC learning is correct before you pursue DHCP issues.

IPv6 DHCP Relay Client Side Troubleshooting

You can collect a client console dump, which can be used to analyze why the received packet cannot be processed and the allocated address cannot be used by the client.

In addition, restarting the client can also fix the issue in some cases.

Make sure the client supports IPv6 requests.

Connect the server directly to the client. If the IP is assigned, then the problem is with the relay.

Enabling Trace Messages for IPv6 DHCP Relay

Use this procedure to enable trace for IPv6 DHCP Relay and enable IPv6 forwarding trace.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. To troubleshoot IPv6 DHCP Relay, you can enable rcip6 trace messages using the following command:

   ```
   trace level 66 3
   ```

3. You can also enable IPv6 forwarding trace using the following command:

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

Example

Enable rcip6 trace messages and enable IPv6 forwarding trace:

```
Switch:1>enable
Switch:1#trace level 66 3
Switch:1#trace ipv6 forwarding
```

Downgrading or Upgrading from Releases that Support Different Key Sizes

Use this procedure if you need to downgrade or upgrade from a release that supports different key sizes.

Different releases can support different DSA host key, RSA host key, and DSA user key sizes. If you need to upgrade or downgrade to an earlier release that does not support the same key size, you must delete all of the keys from the .ssh directory and generate new keys for SSH. If you do not do this, key sizes that are no longer supported will no longer function.

You only need to perform this procedure if you have previously generated DSA host, RSA host, or DSA user keys with a release that supports different key sizes.

Procedure

1. Use the following command to disable SSH:

   ```
   no ssh
   ```

2. From the config terminal go to the .ssh directory using the command:

   ```
   cd /intflash/.ssh
   ```

3. After you upgrade or downgrade, delete the following keys from the .ssh directory.

```
ssh_dss.key
ssh_rsa.key
moc_sshc_dsa_file
moc_sshc_rsa_file
id_dsa_rwa
id_dsa_rwa.pub
id_rsa_rwa
id_rsa_rwa.pub
moc_sshc_dsa_file_fed
moc_sshc_rsa_file_fed
known_hosts
ssh_ecdsa.key
dsa_key_<access level like rwa/rw/ro/admin/security/privilege/operator/auditor>,
example: dsa_key_rwa
rsa_key_<access level like rwa/rw/ro/admin/security/privilege/operator/auditor>,
example: rsa_key_rwa
```

4. Generate a new DSA host key:

   ```
   ssh dsa-host-key [<1024-1024>]
   ```

5. Generate a new SSH DSA user key:

   ```
   ssh dsa-user-key WORD<1-15> [size <1024-1024>]
   ```

6. Generate a new RSA host key:

   ```
   ssh rsa-host-key [<1024-2048>]
   ```

Troubleshooting TACACS+

The switch supports the Terminal Access Controller Access Control System plus (TACACS+) client. TACACS+ is a remote authentication protocol that provides centralized validation of users who attempt to gain access to a router or network access server (NAS). The TACACS+ feature is disabled by default.

The switch implementation of TACACS+ does not support:

- Earlier versions of TACACS
- Point-to-Point Protocol (PPP) authentication and accounting
- IPv6 addresses

TACACS+ is part of the Base Software License. For more information about licensing, see Administering VOSS.

See the following sections to troubleshoot TACACS+.

Unable to Log On Using Telnet or Rlogin

If you cannot log on using Telnet or rlogin, perform the following steps.

Procedure

1. Check whether the TACACS+ server is available or unreachable.
2. On the TACACS+ server, check whether you configured the privilege level correctly. On successful authorization, the TACACS+ server returns an access level to the switch for the current user, which determines the user access privileges. The switch supports access levels 1 to 6 and access level 15.

   The following table maps user accounts to TACACS+ privilege level.

Switch access level	TACACS+ privilege level	Description
NONE	0	If the TACACS+ server returns an access level of 0, the user is denied access. You cannot log into the device if you have an access level of 0.
READ ONLY	1	Permits you to view only configuration and status information.
LAYER 1 READ WRITE	2	Permits you to view most of the switch configuration and status information and change physical port settings.
LAYER 2 READ WRITE	3	Permits you to view and change configuration and status information for Layer 2 (bridging and switching) functions.
LAYER 3 READ WRITE	4	Permits you to view and change configuration and status information for Layer 2 and Layer 3 (routing) functions.
READ WRITE	5	Permits you to view and change configuration and status information across the switch. This level does not allow you to change security and password settings.
READ WRITE ALL	6	Permits you to have all the rights of read-write access and the ability to change security settings, including command line interface (CLI) and web-based management user names and passwords, and the SNMP community strings.
NONE	7 to 14	If the TACACS+ server returns an access level of 7 to 14, the

Table continues…

Switch access level	TACACS+ privilege level	Description
		user is denied access. You cannot log into the device if you have an access level of 7 to 14.
READ WRITE ALL	15	Permits you to have all the rights of read-write access and the ability to change security settings, including command line interface (CLI) and Web-based management user names and passwords, and the SNMP community strings. **Note:** Access level 15 is internally mapped to access level 6, which ensures consistency with other vendor implementations. The switch does not differentiate between an access level of 6 and an access level of 15.

After you enable TACACS+ authorization, the current privilege-level to command mapping on the switch is no longer relevant because the TACACS+ server has complete responsibility for command authorization. TACACS+ authorization provides access to the system based on username, not based on privilege level.

Note:

If you want to switch to a privilege level 'X' using **`tacacs switch level <1-15>`** command, you must create a user "$enabX$" on the TACACS+ server. X is the privilege level that you want to change.

3. On the TACACS+ server, check whether you configured the password and user name correctly.
4. On the TACACS+ server, check whether you configured the switch IP address in the trust list.
5. Check whether you configured the encryption key, connection mode (single connection or per-session connection), and TCP port number the same on the TACACS+ server and switch.
6. If you can log on to the switch, check whether the TACACS+ server configured on the platform has the correct IP address:

   ```
   show tacacs
   ```

7. Use the output from the preceding step to verify whether the key field configured on the platform is the same as that on the TACACS+ server.
8. Also use the output from the **show tacacs** command to verify whether you configured the single connection option on the platform, and whether the TACACS+ server supports the single connection.

Example

Check whether the TACACS+ server configured on the platform has the correct IP address:

```
Switch:1>enable
Switch:1(config)#show tacacs

Global Status:

   global enable : false

   authentication enabled for : cli

   accounting enabled for : none

   authorization : disabled

   User privilege levels set for command authorization : None

Server:

                       create :

Prio    Status  Key     Port  IP address  Timeout Single Source SourceEnabled
Primary NotConn ******   3    192.0.2.254      30   true 5.5.5.5  true
Backup  NotConn ******  47    198.51.100.1      10  false 0.0.0.0 false
```

Job Aid

The following table describes the fields in the output for the **show tacacs** command.

Name	Description
Global Status	
global enable	Displays if the TACACS+ feature is enabled globally.
authentication enabled for	Displays which application is authenticated by TACACS+. The possibilities are CLI, web, or all.
accounting enabled for	Displays if accounting is enabled. You can only enable accounting for CLI. By default, accounting is not enabled.
authorization	Displays if authorization is enabled.
User privilege levels set for command authorization	Displays the privilege levels set for command authorization. When you configure command authorization for a particular level, all commands that you execute are sent to the TACACS+ server for authorization. The device can only execute the commands the TACACS+ server authorizes.

Table continues…

Name	Description
	The user privilege levels are: • 0: denied access • 1: read only (ro) access • 2: Layer 1 read and write (l1) access • 3: Layer 2 read and write (l2) access • 4: Layer 3 read and write (l3) access • 5: read and write (rw) access • 6: read and write all (rwa) access • 7-14: denied access • 15: read and write all (rwa) access
`Server`	
`Prio`	Displays the priority of the TACACS+ server. The switch attempts to use the primary server first, and the secondary server second.
`Status`	Displays the connection status between the server and the switch – connected or not connected.
`Key`	Displays as ****** instead of the actual key. The key is secret and is not visible.
`Port`	Displays the TCP port used to establish the connection to the server. The default port is 49.
`IP address`	Displays the IP address for the primary and secondary TACACS+ servers.
`Timeout`	Displays the period of time, in seconds, the switch waits for a response from the TACACS+ daemon before it times out and declares an error. The default is 10 seconds.
`Single`	Displays if a single open connection is maintained between the switch and TACACS+ daemon, or if the switch opens and closes the TCP connection to the TACACS+ daemon each time they communicate. The default is false, which means the device does not maintain the single open connection.
`Source`	Displays the fixed source IP address, if you configure one, for all outgoing TACACS+ packets.
`SourceEnabled`	Displays if the fixed source IP address is enabled for all outgoing TACACS+ packets.

Unable to Log On Using SSH

If you cannot log on using Secure Shell (SSH), perform the following steps.

Procedure

1. Verify that the network, the switch, and the TACACS+ server is reachable.
2. Verify whether you configured the SSH client correctly.
3. Verify whether you enabled and configured the SSH function correctly on the switch:

   ```
   show ssh global
   ```

Example

Verify whether you enabled and configured SSH function correctly on the switch:

```
Switch:1>enable
Switch:1#show ssh global

Total Active Sessions  : 0
        version            : v2only
        port               : 22
        max-sessions       : 4
        timeout            : 60
        action rsa-keygen  : rsa-keysize 2048
        action dsa-keygen  : dsa-keysize 2048
        rsa-auth           : true
        dsa-auth           : true
        pass-auth          : false
        enable             : true
```

Job Aid

The following table describes the fields in the output for the **`show ssh global`** command.

Parameter	Description
Total active sessions	Specifies the number of active SSH sessions underway.
version	Specifies if SSH is version 1 or version 2. The default is v2. Extreme Networks recommends that you configure the version to v2 only.
port	Specifies the SSH connection port. The default is 22. You cannot configure the following TCP ports as SSH connection ports: 0 to 1024 (except port 22), 1100, 4095, 5000, 5111, 6000, or 999.
max-sessions	Specifies the maximum number of SSH sessions allowed. The default is 4.
timeout	Specifies the SSH connection authentication timeout in seconds. The default is 60 seconds.
action rsa-keygen	Specifies the SSH RSA key size.

Table continues…

Parameter	Description
action dsa-keygen	Specifies the SSH DSA key size.
rsa-auth	Specifies if RSA authentication is enabled or disabled. The default is enabled.
dsa-auth	Specifies if DSA authentication is enabled or disabled. The default is enabled.
pass-auth	Specifies if password authentication is enabled or disabled. The default is enabled.
enable	Specifies if SSH secure mode is enabled. False is disabled. Secure is enabled.

Unable to Log On by any Means (Telnet, Rlogin, or SSH)

If you cannot log on by any means, perform the following steps.

Procedure

1. Check whether the TACACS+ server runs properly and try to restart the TACACS+ server.
2. Check whether you enabled both TACACS+ and RADIUS on the switch.

   ```
   show radius
   ```

   ```
   show tacacs
   ```

   If TACACS+ fails, RADIUS can take over the authentication, authorization, and accounting (AAA) process.
3. Check whether you configured the TACACS+ server to unencrypted mode, as the switch always sends encrypted TACACS+ messages.
4. Check whether you configured the switch properly. In particular, check the IP address and key.

   ```
   show tacacs
   ```
5. Check whether you configured the encryption key, connection mode (single connection or per-session connection), and TCP port number the same on the TACACS+ server and switch.
6. If the server connects directly, check whether the administrative and operation status of the port is up:

   ```
   show interface gigabitethernet {slot/port[/sub-port][-slot/port[/sub-port]][,...]}
   ```
7. If the server is connected in a network, check whether the switch has a route configured to the server network:

   ```
   show ip route
   ```

Example

Check whether you enabled both TACACS+ and RADIUS on the switch:

```
Switch:1>enable
Switch:1(config)#show tacacs

Global Status:

   global enable : false

   authentication enabled for : cli

   accounting enabled for : none

   authorization : disabled

   User privilege levels set for command authorization : None

Server:
                     create :

Prio   Status  Key     Port  IP address  Timeout Single Source SourceEnabled
Primary NotConn ******   3    192.0.2.254      30   true 5.5.5.5  true
Backup  NotConn ******  47    198.51.100.1      10  false 0.0.0.0 false

Switch:1(config)#show radius
             acct-attribute-value : 193
                     acct-enable : false
        acct-include-cli-commands : false
        access-priority-attribute : 192
             auth-info-attr-value : 91
         command-access-attribute : 194
           cli-commands-attribute : 195
                    cli-cmd-count : 40
               cli-profile-enable : false
                           enable : false
                 igap-passwd-attr : standard
           igap-timeout-log-fsize : 512
                        maxserver : 10
            mcast-addr-attr-value : 90
                    sourceip-flag : false
```

Check whether the administrative and operation status of the port is up:

```
Switch:1#show interface gigabitethernet 1/2

================================================================================
                                Port Interface
================================================================================
PORT                       LINK  PORT           PHYSICAL          STATUS
NUM   INDEX DESCRIPTION    TRAP  LOCK     MTU   ADDRESS           ADMIN  OPERATE
--------------------------------------------------------------------------------
1/2   257   1000BaseTX     true  false    1950  00:24:7f:a1:70:61 up     up


================================================================================
                                  Port Name
================================================================================
PORT                                              OPERATE  OPERATE  OPERATE

NUM   NAME                          DESCRIPTION   STATUS   DUPLEX    SPEED    VL
AN
--------------------------------------------------------------------------------
```

```
1/2                                  1000BaseTX    up       full     1000     Ta
gged


================================================================================
                                  Port Config
================================================================================
PORT                DIFF-SERV   QOS   MLT   VENDOR

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

Check whether the switch has a route configured to the server network:

```
Switch:1(config)#show ip route


================================================================================================
                                    IP Route - GlobalRouter
================================================================================================
                                                    NH
INTER
DST             MASK            NEXT                VRF/ISID       COST FACE  PROT AGE
TYPE PRF
------------------------------------------------------------------------------------------------
198.51.100.1        255.255.255.255 192.0.2.65      GlobalRouter     1   100   OSPF 0
IB   125
198.51.100.5         255.255.255.255 192.0.2.5       -               1   0     LOC  0
DB   0
198.51.100.13        255.255.255.255                      GlobalRouter     10  1000  ISIS
0   IBS  7
198.51.100.200       255.255.255.255                      GlobalRouter     10  1000  ISIS
0   IBS  7
4 out of 4 Total Num of Route Entries, 4 Total Num of Dest Networks displayed.
------------------------------------------------------------------------------------------------
TYPE Legend:
I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best Route, E=Ecmp Rout
e,
U=Unresolved Route, N=Not in HW, F=Replaced by FTN, V=IPVPN Route, S=SPBM Route
PROTOCOL Legend:
v=Inter-VRF route redistributed
```

Job Aid

The following table describes the fields in the output for the **`show radius`** command.

Parameter	Description
acct-attribute-value	Specifies the accounting attribute value.
acct-enable	Specifies if the accounting attribute is enabled.
acct-include-cli-commands	Specifies if the accounting attribute includes CLI commands. The default is false.
access-priority-attribute	Specifies the value of the access priority attribute. The default is 192.
auth-info-attr-value	Specifies the value of the authentication information attribute. The default is 91.
command-access-attribute	Specifies the value of the command access attribute. The default is 194.

Table continues…

Parameter	Description
cli-commands-attribute	Specifies the value of the CLI commands attribute. The default is 195.
cli-cmd-count	Specifies how many CLI commands before the system sends a RADIUS accounting interim request. The default is 40.
cli-profile-enable	Specifies if RADIUS CLI profiling is enabled. CLI profiling grants or denies access to users being authenticated by way of the RADIUS server. You can add a set of CLI commands to the configuration on the RADIUS server, and you can specify the command-access mode for these commands. The default is false.
enable	Specifies if RADIUS authentication is globally enabled on the switch.
igap-passwd-attr	Specifies the IGMP for user Authentication Protocol (IGAP) password attribute.
igap-timeout-log-fsize	Specifies the IGMP for user Authentication Protocol (IGAP) timeout log file size.
maxserver	Specifies the maximum number of servers allowed for the device. The default is 10.
mcast-addr-attr-value	Specifies the value of the multicast address attribute. The default is 90.
sourceip-flag	Specifies if the switch can use a configured source IP address. If the outgoing interface on the switch fails, a different source IP address is used, which requires that you make configuration changes to define the new RADIUS client on the RADIUS server. To simplify RADIUS server configuration, you can configure the switch to use a circuitless IP (CLIP) address as the source IP and NAS IP address when transmitting RADIUS packets. By default, the switch uses the IP address of the outgoing interface as the source IP, and the NAS IP address for RADIUS packets that it transmits.

Administrator Unable to Obtain Accounting Information from the TACACS+ Server

If the administrator is unable to obtain accounting information from the TACACS+ server, perform the following steps.

Procedure

1. Check whether you enabled accounting on the switch:

```
show tacacs
```

2. Check whether you enabled accounting on the TACACS+ server.

Example

Check whether accounting is enabled on the switch:

```
Switch:1>enable
Switch:1(config)#show tacacs

Global Status:

   global enable : false

   authentication enabled for : cli

   accounting enabled for : none

   authorization : disabled

   User privilege levels set for command authorization : None

Server:

                       create :

Prio    Status  Key     Port  IP address  Timeout Single Source SourceEnabled
Primary NotConn ******   3    192.0.2.254      30   true 5.5.5.5  true
Backup  NotConn ******  47    198.51.100.1      10  false 0.0.0.0 false
```

Trap Server Cannot Receive Trap Packets from the Switch

If the trap server cannot receive trap packets from the switch, perform the following steps.

Procedure

1. Check whether you configured the trap server correctly on the switch:

   ```
   show snmp-server host
   ```

2. Check whether a firewall exists between the switch and the trap server.

Example

Check whether you configured the trap server correctly on the switch:

```
Switch:1>enable
Switch:1#show snmp-server host

================================================================================
                              Notify Configuration
================================================================================
Notify Name                     Tag                                 Type
--------------------------------------------------------------------------------
Inform                          informTag                           inform
Trap                            trapTag                             trap


================================================================================
                         Notify Profile Configuration
```

```
================================================================================
Params Name                     Profile Name
--------------------------------------------------------------------------------
AuthNoPriv-md5                  profile2
AuthPriv-md5                    profile3
NoAuthNoPriv-md5                profile1


================================================================================
                         Target Address Configuration
================================================================================

Target Name                     TDomain   TAddress                 TMask

--------------------------------------------------------------------------------
4c20cc369925edbd1fe3cf8e2584c498 ipv4      47.17.142.155:162

55fca382ffba169e986783bbbdedc334 ipv4      47.17.143.57:162



================================================================================
                         Target Address Configuration
================================================================================
Target Name                     Timeout Retry TagList
Params                          MMS
--------------------------------------------------------------------------------
4c20cc369925edbd1fe3cf8e2584c498 1500    3     trapTag
4c20cc369925edbd1fe3cf8e2584c498 484
55fca382ffba169e986783bbbdedc334 1500    3     trapTag
55fca382ffba169e986783bbbdedc334 484

================================================================================
                         Target Params Configuration
================================================================================
Target Name                     MP Model  Security Name                     Sec
Level
--------------------------------------------------------------------------------
4c20cc369925edbd1fe3cf8e2584c498 snmpv1    readview                          noAu
thNoPriv
55fca382ffba169e986783bbbdedc334 snmpv2c   secret                            noAu
thNoPriv
TparamV1                        snmpv1    readview                          noAu
thNoPriv
TparamV2                        snmpv2c   readview                          noAu
thNoPriv
```

Troubleshooting TACACS+ Problems

Use the **`trace level`** command to check traps and log files to see any TACACS+ failure. If TACACS+ experiences failure conditions, the TACACS+ module sends SNMP traps to notify the user. The TACACS+ module also logs the failure information into the system log file.

About this task

Caution:

Using the trace tool inappropriately can cause primary CPU lockup conditions, loss of access to the device, loss of protocols, and service degradation. If you use trace level 3 (verbose) or trace level 4 (very verbose), do not use the screen to view commands due to the volume of information the system generates and the effect on the system.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Configure the trace level for the TACACS+ module:

   ```
   trace level 109 <1-4>
   ```

   The TACACS+ module ID is 109.

3. Stop trace:

   ```
   trace shutdown
   ```

4. View the trace results on screen:

   ```
   trace screen enable
   ```

5. View trace saved to a file:

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

6. Save the trace file for retrieval:

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

   If you do not specify a file name, the file name is systrace.txt.

Variable Definitions

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

Variable	Value
level [<*Module_ID*>] *[<1–4>]*	Starts the trace by specifying the module ID and level. Module ID 23 represents the IGMP module <*Module_ID*> specifies the module for the trace. Different hardware platforms support different ID ranges because of feature support differences. To see which module IDs are available on the switch, use the **`show trace modid-list`** command or CLI command completion Help. <0-4> specifies the trace level: • 0 — Disabled

Table continues…

Variable	Value
	• 1 — Very terse • 2 — Terse • 3 — Verbose • 4 — Very verbose
shutdown	Stops the trace operation.
screen {disable|enable}	Enables or disables the display of trace output to the screen. **Important:** Extreme Networks recommends you to avoid using the screen to view commands if you use trace level 3 (verbose) or trace level 4 (very verbose) due to the volume of information generated and the effect on the system.

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

Variable	Value
file [tail]	Displays the trace results saved to a file.
level	Displays the current trace level for all modules.
modid-list	Specifies the module ID list.

Using BGP Debugging Commands

Use global and peer debug commands to display specific debug messages for the global and peer Border Gateway Protocol (BGP) configuration, including the BGP neighbors.

You can use these commands to troubleshoot the BGP configuration.

Procedure

1. Enter BGP Router Configuration mode:

   ```
   enable
   configure terminal
   router bgp
   ```

2. Show specific debug messages for the global BGP configuration:

   ```
   global-debug mask WORD<1-100>
   ```

3. Display specific debug messages for the global BGP neighbors:

   ```
   neighbor-debug-all mask WORD<1-100>
   ```

4. Display specific debug messages for BGP peers or peer groups:

   ```
   neighbor <nbr_ipaddr|peer-group-name> neighbor-debug-mask
   WORD<1-100>
   ```

5. Display debug messages on the console:

   ```
   debug-screen <on|off>
   ```

Example

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

Display the global debug messages for error and packet:

```
Switch:1(router-bgp)#global-debug mask error,packet
```

End (disable) the display of global debug messages:

```
Switch:1(router-bgp)#global-debug mask none
```

Display specific debug messages for the global BGP neighbors:

```
Switch:1(router-bgp)#neighbor-debug-all mask packet,event
```

Display specific debug messages for BGP peers or peer groups:

```
Switch:1(router-bgp)#neighbor 192.0.2.10 neighbor-debug-mask event,trace
```

Display debug messages on the console:

```
Switch:1(router-bgp)#debug-screen on
```

Variable Definitions

Use the data in the following table to use the **`global-debug mask`** and **`neighbor-debug-all mask`** commands.

Variable	Value
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.

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

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

Table continues…

Variable	Value
	[<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 26: Mask categories and messages

Mask category	Message
none	None disables the display of all debug messages.
all	All configures the device to show all categories of debug messages.
error	Error configures the device to show error debug messages.
packet	Packet configures the device to show packet debug messages.
event	Event configures the device to show event debug messages.
warning	Warning configures the device to show warning debug messages.
init	Init configures the device to show initialization debug messages.
filter	Filter configures the device to show filter-related debug messages.
update	Update configures the device to show update-related debug messages.

Chapter 11: Traps Reference

The switch generates alarms, traps, and logs. This section provides information about traps.

Proprietary Traps

The following tables describe proprietary traps for the switch. All of the following traps have a status of current.

Table 27: 1.3.6.1.4.1.45.4.8.0.xx series

OID	Notification type	Objects	Description
1.3.6.1.4.1.45.4.8.0.1	slaMonitorAgentExceptionDetected	slaMonitorAgentExceptionDetected	The SLA Monitor (SLA Mon) agent process has terminated unexpectedly. You must reenable SLA Monitor to restart the SLA Monitor agent.

Table 28: 1.3.6.1.4.1.45.5.17.0.xx series

OID	Notification type	Objects	Description
1.3.6.1.4.1.45.5.17.0.1	bsDhcpSnoopingBindingTableFull	bsDhcpSnoopingNotificationClientMACAddr	A bsDhcpSnoopingBindingTableFull notification is generated when you try adding a new DHCP binding entry, and the binding table is full. The value of bsDhcpSnoopingNotificationClientMACAddr gives the MAC address that cannot be added to the binding table. This notification also indicates that additional untrusted DHCP packets will not be

Table continues…

OID	Notification type	Objects	Description
			added to the binding table, and will be dropped.
1.3.6.1.4.1.45.5.17.0.2	bsDhcpSnoopingTrap	• bsDhcpSnoopingNotificationSourcePort • bsDhcpSnoopingNotificationMsgTypebsDhcpSnoopingNotificationSourceMACAddr • bsDhcpSnoopingNotificationClientMACAddr • bsDhcpSnoopingIfTrusted	A bsDhcpSnoopingTrap notification is generated when a DHCP packet is dropped.
1.3.6.1.4.1.45.5.17.0.4	bsDhcpSnoopingStaticEntryMACConflict	bsDhcpSnoopingNotificationSourceMACAddr bsDhcpSnoopingIfIndex	A bsDhcpSnoopingStaticEntryMACConflict notification is generated when a DHCP packet is dropped, because a static entry with the same MAC address was found in the binding table.

Table 29: 1.3.6.1.4.1.45.5.18.0.xx series

OID	Notification type	Objects	Description
1.3.6.1.4.1.45.5.18.0.1	bsaiArpPacketDroppedOnUntrustedPort	• bsArpInspectionIfTrusted • bsArpInspectionNotificationSourceMACAddr	A bsaiArpPacketDroppedOnUntrustedPort trap signifies that an ARP packet is dropped on an untrusted port, due to an invalid IP or MAC address binding. The port is identified by the instance of bsArpInspectionIfTrusted generated by the trap.

Table 30: 1.3.6.1.4.1.45.5.20.0.xx series

OID	Notification type	Objects	Description
1.3.6.1.4.1.45.5.20.0.2	bsSourceGuardCannotEnablePort	bsSourceGuardConfigMode	A bsSourceGuardCannotEnablePort trap signifies

Table continues…

OID	Notification type	Objects	Description
			that there are insufficient resources to enable IP Source Guard checking on a port because of internal state changes within the system such as system initialization. The port is identified by the instance of bsSourceGuardConfigMode generated by the trap.
1.3.6.1.4.1.45.5.20.0.1	bsSourceGuardReachedMaxIpEntries	bsSourceGuardConfigMode	A bsSourceGuardReachedMaxIpEntries trap signifies that the maximum number of IP address entries on a port has been reached. The port is identified by the instance of bssourceGuardConfigMode generated by the trap.

Table 31: 1.3.6.1.4.1.45.5.34.0.xx series

OID	Notification type	Objects	Description
1.3.6.1.4.1.45.5.34.0.1	bsnesGloballyEnabled	—	A bsnesGloballyEnabled trap signifies that the Energy Saver feature was enabled globally.
1.3.6.1.4.1.45.5.34.0.2	bsnesGloballyDisabled	—	A bsnesGloballyDisabled trap signifies that the Energy Saver feature was disabled globally.
1.3.6.1.4.1.45.5.34.0.3	bsnesManuallyActivated	—	A bsnesManuallyActivated trap signifies that the Energy Saver was activated manually.
1.3.6.1.4.1.45.5.34.0.4	bsnesManuallyDeactivated	—	A bsnesManuallyDeactivated trap signifies that the Energy Saver feature was deactivated manually.

Table continues…

OID	Notification type	Objects	Description
1.3.6.1.4.1.45.5.34.0.5	bsnesScheduleNotApplied	—	A bsnesScheduleNotApplied trap signifies that a schedule was not applied because SNTP is not synchronized.
1.3.6.1.4.1.45.5.34.0.6	bsnesScheduleApplied	—	A bsnesScheduleApplied trap signifies that SNTP is synchronized and the schedule is being applied.
1.3.6.1.4.1.45.5.34.0.7	bsnesActivated	—	A bsnesActivated trap signifies that the Energy Saver feature was activated by the schedule.
1.3.6.1.4.1.45.5.34.0.8	bsnesDeactivated	—	A bsnesDeactivated trap signifies that the Energy Saver feature was deactivated by the schedule.

Table 32: 1.3.6.1.4.1.45.5.43.0.xx series

ODID	Notification type	Objects	Description
1.3.6.1.4.1.45.5.43.0.1	bsLstInterfaceStatusChanged	• ifIndex • bsLstInterfaceStatus • bsLstGroupIndex	A bsLstInterfaceStatusChanged trap signifies that a physical or logical interface changed status in a Link-state tracking (LST) tracking group.
1.3.6.1.4.1.45.5.43.0.2	bsLstInterfaceOperStateChanged	• ifIndex • bsLstInterfaceStatus • bsLstGroupIndex	A bsLstInterfaceOperStateChanged trap signifies that the operational status of an LST group changed due to an interface status change. For example, when the last interface in an LST group is down, the operational status of theLST group changes to down.

Table 33: 1.3.6.1.4.1.2272.1.4.10.1.1.xxx series

OID	Notification type	Objects	Description
1.3.6.1.4.1.2272.1.4.10.1.1.114	rcPortShutdownReason	• rcPortVLacpTotalFlapCount • rcPortVLacpFirstFlapTimeStamp • rcPortVLacpLastFlapTimeStamp	An rcPortShutdownReason trap signifies that the specific VLACP link port status is down as VLACP flaps are detected on that port.

Table 34: 1.3.6.1.4.1.2272.1.21.0.xx series

OID	Notification type	Objects	Description
1.3.6.1.4.1.2272.1.21.0.3	rcnErrorNotification	rcErrorLevel rcErrorCode rcErrorText	An rcnErrorNotification trap signifies that the SNMPv2 entity, acting in an agent role, has detected an error condition.
1.3.6.1.4.1.2272.1.21.0.4	rcnStpNewRoot	rcStgId	An rcnStpNewRoot trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the Spanning Tree Protocol 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	An rcnStpTopologyChange trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the Spanning Tree Protocol has gone due to a topology change event.
1.3.6.1.4.1.2272.1.21.0.6	rcnChasPowerSupplyDown	• rcChasPowerSupplyId • rcChasPowerSupplyOperStatus	An rcnChasPowerSupplyDown trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the rcChasPowerSupplyOperStatus object for one of its power supply units is about to transition into the down state.
1.3.6.1.4.1.2272.1.21.0.7	rcnChasFanDown	• rcChasFanId • rcChasFanOperStatus	An rcnChasFanDown trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the rcChasFanOperStatus object for one of its power supply units is about to transition into the down state.
1.3.6.1.4.1.2272.1.21.0.8	rcnLinkOscillation	rcPortIndex	An rcnLinkOscillation trap signifies that the SNMPv2 entity, acting in an agent role, has detected an excessive number of link state transitions on the specified port.

Table continues…

OID	Notification type	Objects	Description
1.3.6.1.4.1.2272.1.21.0.9	rcnMacViolation	• rcErrorText • rcPortIndex	An rcnMacViolation trap signifies that the SNMPv2 entity, acting in an agent role, has received a PDU with an invalid source MAC address.
1.3.6.1.4.1.2272.1.21.0.13	rcn2kTemperature	rc2kChassisTemperature	An rcn2kTemperature trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the chassis is overheating.
1.3.6.1.4.1.2272.1.21.0.14	rcnChasPowerSupplyUp	• rcChasPowerSupplyId • rcChasPowerSupplyOperStatus	An rcnChasPowerSupplyUp trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the rcChasPowerSupplyOperStatus object for one of its power supply units is about to transition into the up state.
1.3.6.1.4.1.2272.1.21.0.16	rcnStpTCN	• rcStgId • rcPortIndex • rcStgBridgeAddress	An rcnStpTCN trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the Spanning Tree Protocol has gone due to a topology change event.
1.3.6.1.4.1.2272.1.21.0.17	rcnSmltIstLink Up	—	An rcnSmltIstLinkUp trap signifies that the split MLT link is from down to up.
1.3.6.1.4.1.2272.1.21.0.18	rcnSmltIstLink Down	rcnSmltIstLinkDown	An rcnSmltIstLinkDown 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	An rcnSmltLinkUp trap signifies that the split SMLT link is up.
1.3.6.1.4.1.2272.1.21.0.20	rcnSmltLinkDown	rcMltSmltId	An rcnSmltLinkDown trap signifies that the split SMLT link is down.
1.3.6.1.4.1.2272.1.21.0.21	rcnChasFanUp	• rcChasFanId • rcChasFanOperStatus	An rcnChasFanUp trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the rcChasFanOperStatus object for one of its power supply units is about to transition into the up state.
1.3.6.1.4.1.2272.1.21.0.22	rcnPasswordChange	• rcCliPasswordChange • rcCliPassChangeResult	An rcnPasswordChange trap signifies that the SNMPv2 entity, acting in an agent role, has detected that one of the CLI passwords is changed.
1.3.6.1.4.1.2272.1.21.0.23	rcnEmError	• rc2kCardIndex • rcChasEmModeError	An rcnEmError trap signifies that the SNMPv2 entity, acting in an agent role, has detected an Em error.
1.3.6.1.4.1.2272.1.21.0.26	rcnSmartCpldTimerFired	rc2kCardIndex	An rcnSmartCpldTimerFired trap signifies that the CP ID timer fired.

Table continues…

OID	Notification type	Objects	Description
1.3.6.1.4.1.2272.1.21.0.27	rcnCardCpldNotUpDate	rc2kCardIndex	An rcnCardCpldNotUpDate trap signifies that the CP ID is not up-to-date.
1.3.6.1.4.1.2272.1.21.0.28	rcnIgapLogFileFull	—	An rcnIgapLogFileFull trap signifies that the IGAP accounting time-out Log File has reached the maximum.
1.3.6.1.4.1.2272.1.21.0.30	rcnSshServerEnabled	rcSshGlobalPort	An rcnSshServerEnabled trap signifies that the SSH server is enabled.
1.3.6.1.4.1.2272.1.21.0.31	rcnSshServerDisabled	rcSshGlobalPort	An rcnSshServerDisabled trap signifies that the SSH server is disabled.
1.3.6.1.4.1.2272.1.21.0.37	rcnSaveConfigAction	rcSysActionL1	An rcnSaveConfigAction trap indicates that the switch run time or boot configuration is being saved.
1.3.6.1.4.1.2272.1.21.0.38 **Note:** This trap does not apply to all platforms.	rcnLoopDetectOnPort	• rcVlanId • rcPortIndex	An rcnLoopDetectOnPort trap indicates that a loop has been detected on a port. The VLAN on that port will be disabled.
1.3.6.1.4.1.2272.1.21.0.41	rcnAggLinkUp	rcMltId	An rcnAggLinkUp trap is generated when the operational state of the aggregator changes from down to up.
1.3.6.1.4.1.2272.1.21.0.42	rcnAggLinkDown	rcMltId	An rcnAggLinkDown trap is generated when the operational state of the aggregator changes from up to down.
1.3.6.1.4.1.2272.1.21.0.59	rcnFdbProtectViolation	• rcVlanId • rcPortIndex	The rcnFdbProtectViolation trap signifies that the port 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	An rcnLogMsgControl trap signifies whether the number of times of repetition of the particular Log message has exceeded the particular frequency or count.
1.3.6.1.4.1.2272.1.21.0.61	rcnSaveConfigFile	• rcSysActionL1 • rcSysConfigFileName	An rcnSaveConfig trap signifies that either the runtime config or the boot config has been saved on the switch.
1.3.6.1.4.1.2272.1.21.0.62	rcnDNSRequestResponse	• rcSysDnsServerListIpAddr • rcSysDnsRequestType	An rcnDnsRequestResponse trap signifies that the switch sent a query to the DNS server or that it received a successful response from the DNS Server.

Table continues…

OID	Notification type	Objects	Description
1.3.6.1.4.1.2272.1.21.0.63	rcnDuplicateIpAddress	• ipNetToMediaNetAddress • ipNetToMediaPhysAddress	An rcnDuplicateIpAddress 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	An rcnLoopDetectPortDown 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	An rcnLoopDetectMacDiscard 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	An rcnAutoRecoverPort trap signifies that autorecovery has reenabled a port that was previously disabled by link flap.
1.3.6.1.4.1.2272.1.21.0.69	rcnAutoRecoverLoopDetectedPort	rcVlanNewLoopDetectedAction	An rcnAutoRecoverPort trap signifies that autorecovery has cleared the loop detect action that was taken on a port.
1.3.6.1.4.1.2272.1.21.0.74	rcnTacacsAuthFailure	rcTacacsGlobalLastUserName	An rcnTacacsAuthFailure trap signifies that TACACS+ authentication failed for a user.
1.3.6.1.4.1.2272.1.21.0.75	rcnTacacsNoServers	—	An rcnTacacsNoServers trap signifies that you are unable to use any TACACS+ servers for authentication.
1.3.6.1.4.1.2272.1.21.0.76	rcnTacacsRxUnsupportedFrame	• rcTacacsGlobalLastAddressType • rcTacacsGlobalLastAddress	An rcnTacacsRxUnsupportedFrame trap signifies that an unsupported frame was received from the TACACS+ server.
1.3.6.1.4.1.2272.1.21.0.77	rcnTacacsExceededMaxLogins	—	An rcnTacacsExceededMaxLogins trap signifies that there was an attempt to exceed the maximum number of allowed TACACS+ logins.
1.3.6.1.4.1.2272.1.21.0.78	rcnTacacsClientFailure	—	An rcnTacacsClientFailure trap signifies that the TACACS+ Client application is down.
1.3.6.1.4.1.2272.1.21.0.80	rcnVlacpPortDown	rcPortIndex	An rcnVlacpPortDown trap signifies that VLACP is down on the port specified.
1.3.6.1.4.1.2272.1.21.0.81	rcnVlacpPortUp	rcPortIndex	An rcnVlacpPortUp trap signifies that VLACP is up on the port specified.

Table continues…

OID	Notification type	Objects	Description
1.3.6.1.4.1.2272.1.21.0.83	rcnEapMacIntrusion	• rcSysIpAddr • rcRadiusPaePortNumber • rcRadiusEapLastAuthMac • rcRadiusEapLastRejMac	An rcnEapMacIntrusion trap signifies that an EAP MAC intrusion has occurred on this port.
1.3.6.1.4.1.2272.1.21.0.110	rcnMaxRouteWarnClear	rcVrfName	An 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	An rcnMaxRouteWarnSet trap signifies that the routing table for the virtual router is reaching its maximum size. Take action to prevent this.
1.3.6.1.4.1.2272.1.21.0.112	rcnMaxRouteDropClear	rcVrfName	An rcnMaxRouteDropClear trap signifies that the routing table for the virtual router 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	An rcnMaxRouteDropSet trap signifies that the routing table for the virtual router 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	An rcnMstpNewCistRoot 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	An rcnMstpNewMstiRoot 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	An rcnMstpNewCistRegionalRoot 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.

Table continues…

OID	Notification type	Objects	Description
1.3.6.1.4.1.2272.1.21.0.120	rcnRstpNewRoot	rcStgBridgeAddress	An rcnRstpNewRoot 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	An 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.
1.3.6.1.4.1.2272.1.21.0.143	rcn2kGbicRemoved	• rc2kCardIndex • rcPortIndex • rcPortUserLabel1 • rcPortUserLabel2 • rc2kChassisUserLabel1	An rcGbicRemoved trap signifies that the SNMPv2 entity, acting in an agent role has detected that an XFP or SFP is removed from the specified slot or port.
1.3.6.1.4.1.2272.1.21.0.144	rcn2kGbicInserted	• rc2kCardIndex • rcPortIndex • rcPortUserLabel1 • rcPortUserLabel2 • rc2kChassisUserLabel1	An rcGbicInserted trap signifies that the SNMPv2 entity, acting in an agent role has detected that the an XFP or SFP is inserted in the specified slot or port.
1.3.6.1.4.1.2272.1.21.0.167	rcnChasPowerSupplyNoRedundancy	—	An rcnChasPowerSupplyNoRedundancy trap signifies that the chassis is running on a power supply without redundancy.
1.3.6.1.4.1.2272.1.21.0.168	rcnChasPowerSupplyRedundancy	—	An rcnChasPowerSupplyRedundancy trap signifies that the chassis is running on a power supply with redundancy.
1.3.6.1.4.1.2272.1.21.0.171	rcnLicenseTrialPeriodExpired	—	An rcnLicenseTrialPeriodExpired trap signifies that the Trial Period License has expired.
1.3.6.1.4.1.2272.1.21.0.172	rcnLicenseTrialPeriodExpiry	rcSysLicenseTrialDaysLeft	An rcnLicenseTrialPeriodExpiry trap signifies the time remaining, in days, before the Trial Period License 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.

Table continues…

OID	Notification type	Objects	Description
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 • rcIpResourceUsageGlobalIngressThreshold	This notification is generated when the number of mroute ingress records exceeds the ingress threshold.
1.3.6.1.4.1.2272.1.21.0.176	rcnMrouteEgressThresholdExceeded	• rcIpResourceUsageGlobalIngressRecInUse • rcIpResourceUsageGlobalIngressThreshold	This notification is generated when the number of mroute egress records exceeds the egress threshold.
1.3.6.1.4.1.2272.1.21.0.185	rcnChasPowerSupplyRunningLow	—	An rcnChasPowerSupplyRunningLow trap signifies that the chassis is running on low power supply.
1.3.6.1.4.1.2272.1.21.0.192	rcnIsisPlsbMetricMismatchTrap	• rcIsisLocalLspId • rcIsisLocalL1Metric • rcIsisNgbLspId • rcIsisNgbL1Metric • rcIsisPlsbTrapType • rcIsisTrapIndicator • rcIsisLocalHostName • rcIsisNgbHostName	An rcnIsisPlsbMetricMismatchTrap signifies that a Link State Packet (LSP) with a different value of Level 1 metric is received.
1.3.6.1.4.1.2272.1.21.0.193	rcnIsisPlsbDuplicateSysidTrap	• rcIsisLocalSysId • rcIsisLocalInterface • rcIsisPlsbTrapType • rcIsisTrapIndicator	An rcnIsisPlsbduplicateSysidTrap signifies that a Hello packet with a duplicate system ID is received.
1.3.6.1.4.1.2272.1.21.0.194	rcnIsisPlsbLsdbUpdateTrap	rcIsisPlsbTrapType	An rcnIsisPlsbLsdbUpdateTrap signifies that link state database (LSDB) information has changed.
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.
1.3.6.1.4.1.2272.1.21.0.278	rcnIsisPlsbBvidMismatchTrap	• rcIsisLocalSysId • rcIsisLocalPrimaryBvid	An rcnIsisPlsbBvidMismatchTrap signifies when a backbone VLAN ID (BVID) Type-Length-Value (TLV) from a

Table continues…

OID	Notification type	Objects	Description
		• rcIsisLocalPrimaryTieBrkAlg • rcIsisLocalSecondaryBvid • rcLocalSecondaryTieBrkAlg • rcIsisNgbSysId • rcIsisNgbPrimaryBvid • rcIsisNgbPrimaryTieBrkAlg • rcIsisNgbSecondaryBvid • rcIsisNgbSecondaryTieBrkAlg • rcIsisLocalBvidCounter • rcIsisNgbBvidCounter • rcIsisPlsbTrapType • rcIsisTrapIndicator • rcIsisNgbHostName	neighbor node does not match the local configuration.
1.3.6.1.4.1.2272.1.21.0.279	rcnIsisPlsbSmltVirtBmacMismatchTrap	• rcIsisLocalVirtualBmac • rcIsisPeerVirtualBmac • rcIsisPlsbTrapType • rcIsisTrapIndicator	An rcnIsisPlsbSmltVirtBmacMismatchTrap signifies that the virtual Backbone MAC (BMAC) configured in the switch is different from the virtual BMAC configured on the interswitch trunking (IST) peer.
1.3.6.1.4.1.2272.1.21.0.280	rcnIsisPlsbSmltPeerBmacMismatchTrap	• rcIsisSysId • rcIsisSmltPeerSysId • rcIsisPlsbTrapType • rcIsisTrapIndicator	An rcnIsisPlsbSmltPeerBmacMismatchTrap signifies one of the following situations: • The Split MultiLink Trunking (SMLT) peer Backbone MAC (BMAC) configured in the interswitch trunking (IST) peer is different from the Intermediate-System-to- Intermediate-System (IS-IS) System ID of the local switch. • The SMLT peer BMAC configured on the local switch is different from the IS-IS System ID of the IST peer.

Table continues…

OID	Notification type	Objects	Description
1.3.6.1.4.1.2272.1.21.0.281	rcnIsisPlsbAdjStateTrap	• rcIsisNgbSysId • rcIsisLocalInterface • rcIsisPlsbTrapType • rcIsisAdjState • rcIsisNgbHostName	An rcnIsisPlsbAdjStateTrap signifies when IS-IS adjacency state changes.
1.3.6.1.4.1.2272.1.21.0.282	rcnIsisPlsbDuplicateNNameTrap	• rcIsisNgbNickname • rcIsisPlsbTrapType • rcIsisTrapIndicator • rcIsisNgbSysId • rcIsisDuplicateNnameCounter • rcIsisNgbHostName	An rcnIsisPlsbDuplicateNNameTrap signifies that a Link State Packet (LSP) with a duplicate nickname is received. The trap should be generated by all the switches in the network.
1.3.6.1.4.1.2272.1.21.0.283	rcnIsisPlsbSmltSplitBebMismatchTrap	• rcIsisLocalSmltSplitBeb • rcIsisPeerSmltSplitBeb • rcIsisPlsbTrapType • rcIsisTrapIndicator	An rcnIsisPlsbSmltSplitBebMismatchTrap signifies that the SMLT Split Backbone Edge Bridge (BEB) configured on the local switch and on the IST peer are the same. One IST switch must be configured as the primary Split BEB and the other IST peer must be configured as the secondary Split BEB.
1.3.6.1.4.1.2272.1.21.0.284	rcnIsisPlsbMultiLinkAdjTrap	• rcIsisNgbSysId • rcIsisLocalInterface • rcIsisPrevInterface • rcIsisPlsbTrapType • rcIsisNgbHostName • rcIsisTrapIndicator	An rcnIsisPlsbMultiLinkAdjTrap signifies when the Intermediate-System-to-Intermediate-System (IS-IS) protocol forms more than one adjacency with the same IS-IS.
1.3.6.1.4.1.2272.1.21.0.285	rcnaSshSessionLogout	rcSshGlobalHostIpAddr	An rcnaSshSessionLogout trap signifies a Secure Shell (SSH) session logout.
1.3.6.1.4.1.2272.1.21.0.286	rcnaSshUnauthorizedAccess	rcSshGlobalHostIpAddr	An rcnaSshUnauthorizedAccess trap signifies that 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.

Table continues…

OID	Notification type	Objects	Description
1.3.6.1.4.1.2272.1.21.0.288	rcnaSshSessionLogin	rcSshGlobalHostIpAddr	An rcnaSshSessionLogin trap signifies that there is a Secure Shell (SSH) session login.
1.3.6.1.4.1.2272.1.21.0.305	RcIsisPlsbSmltVirtBmacMisconfigTrap	• rcIsisSmltVirtBmacMisconfigNodeSysId • rcIsisPlsbTrapType • rcIsisSmltVirtBmacMisconfigNodeHostName • rcIsisTrapIndicator	An SPBM ISIS trap signifies that SMLT virtual BMAC has been used by nodes other than the SMLT nodes as system-id or MAC.
1.3.6.1.4.1.2272.1.21.0.306 **Note:** This trap does not apply to all platforms.	rcnPortChannelizedStateChangeTrap	• rcPortIndex • rcChannelizedPortAdminMode	An rcnPortChannelizedStateChangeTrap notification signifies that a port channelized state has changed by administratively enabling or disabling.
1.3.6.1.4.1.2272.1.21.0.335	rcnSystemUsbInternalAccessErrorTrap	—	An rcnSystemUsbInternalAccessErrorTrap notification signifies that the system has lost internal access to the USB. This trap only applies to platforms that require the USB as part of the operating system.
1.3.6.1.4.1.2272.1.21.0.341	rcnDvrVistPeerDomainMismatchErrorTrap	rcIsisPeerVirtualBmac	An rcnDvrVistPeerDomainMismatchErrorTrap notification is generated when the VIST link comes up between two DvR peers that are in different DvR domains.
1.3.6.1.4.1.2272.1.21.0.342	rcnDvrVistPeerDomainMismatchErrorClearTrap	rcIsisPeerVirtualBmac	An rcnDvrVistPeerDomainMismatchErrorTrap notification is generated when the error condition of having a VIST link up between two DvR peers from different DvR domains is cleared.
1.3.6.1.4.1.2272.1.21.0.351	rcnIsisPlsbIsisEnabledWithZeroNicknameTrap	• rcIsisLocalSysId • rcIsisPlsbTrapType • rcIsisTrapIndicator	An rcnIsisPlsbIsisEnabledWithZeroNicknameTrap notification is generated when the IS-IS is enabled with a zero nickname.
1.3.6.1.4.1.2272.1.21.0.352	rcnRestConfServerOperationStatusTrap	rcnRestConfServerOperationStatus	An rcnRestConfServerOperationStatusTrap notification is generated when the RESTCONF server is enabled or disabled to indicate the operational status of the RESTCONF server.

Table 35: 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 Virtual Router Redundancy Protocol (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 Virtual Router Redundancy Protocol (VRRP), for instance, a transition from master to backup when shutdown is received.

Standard Traps

The following table describes standard traps that the switch can generate.

Table 36: Standard traps

OID	Notification type	Objects	Description
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
1.3.6.1.2.1.46.1.3.0.3	vrrpTrapStateTransition	ifIndex vrrpTrapStateTransitionType vrrpTrapStateTransitionCause vrrpOperVrId	A vrrpTrapStateTransition trap signifies a state transition has occurred on a particular Virtual Router Redundancy Protocol (VRRP) interface. Implementation of this trap is optional.

Table continues…

OID	Notification type	Objects	Description
		vrrpOperIpAddr ipAdEntAddr	vrrpOperIpAddr contains the IP address of the VRRP interface while ipAdEntAddr contains the IP address assigned to the physical interface.
1.3.6.1.2.1.68.0.1	vrrpTrapNewMaster	vrrpOperMasterIpAddr	The newMaster trap indicates that the sending agent has transitioned to Master state.
1.3.6.1.2.1.68.0.2	vrrpTrapAuthFailure	vrrpTrapPacketSrc vrrpTrapAuthErrorType	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	pingProbeFailed	pingCtlTargetAddressType pingCtlTargetAddress pingResultsOperStatus pingResultsIpTargetAddressType pingResultsIpTargetAddress pingResultsMinRtt pingResultsMaxRtt pingResultsAverageRtt pingResultsProbeResponse pingResultsSentProbes pingResultsRttSumOfSquares pingResultsLastGoodProbe	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	pingCtlTargetAddressType pingCtlTargetAddress pingResultsOperStatus pingResultsIpTargetAddressType pingResultsIpTargetAddress pingResultsMinRtt pingResultsMaxRtt pingResultsAverageRtt pingResultsProbeResponses pingResultsSentProbes pingResultsRttSumOfSquares pingResultsLastGoodProbe	This trap is generated after a ping test fails when the corresponding pingCtlTrapGeneration object is configured to testFailure(1). In this instance pingCtlTrapTestFailureFilter specifies the number of probes in a test required to fail to consider the test as failed.

Table continues…

OID	Notification type	Objects	Description
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.
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.

Table continues…

OID	Notification type	Objects	Description
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	—	—

Glossary

attenuation	The decrease in signal strength in an optical fiber caused by absorption and scattering.
Backbone Core Bridge (BCB)	Backbone Core Bridges (BCBs) form the core of the SPBM network. The BCBs are SPBM nodes that do not terminate the VSN services. BCBs forward encapsulated VSN traffic based on the Backbone MAC Destination Address (B-MAC-DA). A BCB can access information to send that traffic to any Backbone Edge Bridges (BEBs) in the SPBM backbone.
Backbone Edge Bridge (BEB)	Backbone Edge Bridges (BEBs) are SPBM nodes where Virtual Services Networks (VSNs) terminate. BEBs handle the boundary between the core MAC-in-MAC Shortest Bath Bridging MAC (SPBM) domain and the edge customer 802.1Q domain. A BEB node performs 802.1ah MAC-in-MAC encapsulation and decapsulation for the Virtual Services Network (VSN).
Backbone MAC (B-MAC)	Provider Backbone Bridging (PBB) MAC-in-MAC encapsulation encapsulates customer MAC addresses in Backbone MAC (B-MAC) addresses. MAC-in-MAC encapsulation defines a BMAC-DA and BMAC-SA to identify the backbone source and destination addresses. The originating node creates a MAC header that SPBM uses for delivery from end to end. As the MAC header stays the same across the network, no need exists to swap a label or perform a route lookup at each node, allowing the frame to follow the most efficient forwarding path end to end. In Shortest Path Bridging MAC (SPBM), each node has a System ID, which is used in the topology announcement. This same System ID also serves as the switch Backbone MAC address (B-MAC), which is used as the source and destination MAC address in the SPBM network.
Backbone VLAN identifier (B-VID)	The Backbone VLAN identifier (B-VID) indicates the Shortest Path Bridging MAC (SPBM) B-VLAN associated with the SPBM instance.
Complete Sequence Number Packets (CSNP)	Complete Sequence Number Packets (CSNP) contain the most recent sequence numbers of all Link State Packets (LSPs) in the database. When all routers update their LSP database, synchronization is complete.
Connectivity Fault Management (CFM)	Connectivity Fault Management is a mechanism to debug connectivity issues and to isolate faults within the Shortest Path Bridging MAC (SPBM) network. CFM operates at Layer 2 and provides the equivalent of ping and traceroute. IEEE 802.1ag Connectivity Fault Management (CFM) divides or

	separates a network into administrative domains called Maintenance Domains (MD).
Customer MAC (C-MAC)	For customer MAC (C-MAC) addresses, which is customer traffic, to forward across the service provider back, SPBM uses IEEE 802.1ah Provider Backbone Bridging MAC-in-MAC encapsulation. The system encapsulates C-MAC addresses within a backbone MAC (B-MAC) address pair made up of a BMAC destination address (BMAC-DA) and a BMAC source address (BMAC-SA).
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.
designated router (DR)	A single router elected as the designated router for the network. In a broadcast or nonbroadcast multiple access (NBMA) network running the Open Shortest Path First (OSPF) protocol, a DR ensures all network routers synchronize with each other and advertises the network to the rest of the Autonomous System (AS). In a multicast network running Protocol Independent Multicast (PIM), the DR acts as a representative router for directly connected hosts. The DR sends control messages to the rendezvous point (RP) router, sends register messages to the RP on behalf of directly connected sources, and maintains RP router status information for the group.
Dynamic Random Access Memory (DRAM)	A read-write random-access memory, in which the digital information is represented by charges stored on the capacitors and must be repeatedly replenished to retain the information.
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.
Enterprise Device Manager (EDM)	A web-based embedded management system to support single-element management. EDM provides complete configuration management functionality for the supported devices and is supplied to the customer as embedded software in the device.
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 certain regular expression (RE).

I/O module

An I/O module is a module that provides network connectivity for various media (sometimes called Layer 0) and protocol types. I/O modules are also called Ethernet modules.

Intermediate System to Intermediate System (IS-IS)

Intermediate System to Intermediate System(IS-IS) is a link-state, interior gateway protocol. ISO terminology refers to routers as Intermediate Systems (IS), hence the name Intermediate System to Intermediate System (IS-IS). IS-IS operation is similar to Open Shortest Path First (OSPF).

In Shortest Path Bridging MAC (SPBM) networks, IS-IS discovers network topology and builds shortest path trees between network nodes that IS-IS uses for forwarding unicast traffic and determining the forwarding table for multicast traffic. SPBM employs IS-IS as the interior gateway protocol and implements additional Type-Length-Values (TLVs) to support additional functionality.

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)

IGMP is a host membership protocol used to arbitrate membership in multicast services. IP multicast routers use IGMP to learn the existence of host group members on their directly attached subnets.

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.

IS-IS Hello packets

Intermediate System to Intermediate System (IS-IS) uses Hello packets to initialize and maintain adjacencies between neighboring routers. IS-IS Hello packets contain the IP address of the interface over which the Hello transmits. These packets are broadcast to discover the identities of neighboring IS-IS systems and to determine whether the neighbor is a Level 1 router.

Layer 1

Layer 1 is the Physical Layer of the Open System Interconnection (OSI) model. Layer 1 interacts with the MAC sublayer of Layer 2, and performs character encoding, transmission, reception, and character decoding.

Layer 2	Layer 2 is the Data Link Layer of the OSI model. Examples of Layer 2 protocols are Ethernet and Frame Relay.
Layer 2 Virtual Services Network	The Layer 2 Virtual Services Network (L2 VSN) feature provides IP connectivity over SPBM for VLANs. Backbone Edge Bridges (BEBs) handle Layer 2 virtualization. At the BEBs you map the end-user VLAN to a Service Instance Identifier (I-SID). BEBs that have the same I-SID configured can participate in the same Layer 2 Virtual Services Network (VSN).
Layer 3	Layer 3 is the Network Layer of the OSI model. An example of a Layer 3 protocol is Internet Protocol (IP).
Layer 3 Virtual Services Network	The Layer 3 Virtual Services Network (L3 VSN) feature provides IP connectivity over SPBM for VRFs. Backbone Edge Bridges (BEBs) handle Layer 3 virtualized. At the BEBs through local provisioning, you map the end-user IP enabled VLAN or VLANs to a Virtualized Routing and Forwarding (VRF) instance. Then you map the VRF to a Service Instance Identifier (I-SID). VRFs that have the same I-SID configured can participate in the same Layer 3 Virtual Service Network (VSN).
Layer 4	The Transport Layer of the OSI model. An example of a Layer 4 protocol is Transmission Control Protocol (TCP).
light emitting diode (LED)	A semiconductor diode that emits light when a current passes through it.
Link State Protocol Data Unit (LSPDUs)	Link State Protocol Data Unit is similar to a Link State Advertisement in Open Shortest Path First (OSPF). Intermediate System to Intermediate System (IS-IS) runs on all nodes of Shortest Path Bridging-MAC (SPBM). Since IS-IS is the basis of SPBM, the device must first form the IS-IS adjacency by first sending out hellos and then Link State Protocol Data Units. After the hellos are confirmed both nodes sends Link State Protocol Data Units (LSPDUs) that contain connectivity information for the SPBM node. These nodes also send copies of all other LSPDUs they have in their databases. This establishes a network of connectivity providing the necessary information for each node to find the best and proper path to all destinations in the network.
link trace message	The link trace message (LTM) is often compared to traceroute. A MEP transmits the LTM packet. This packet specifies the target MAC address of an MP, which is the SPBM system id or the virtual SMLT MAC. MPs on the path to the target address respond with an LTR. LTM contains: • Time to live (TTL) • Transaction Identifier • Originator MAC address

- Target MAC address

link-state database (LSDB)

A database built by each OSPF router to store LSA information. The router uses the LSDB to calculate the shortest path to each destination in the autonomous system (AS), with itself at the root of each path.

Local Area Network (LAN)

A data communications system that lies within a limited spatial area, uses a specific user group and topology, and can connect to a public switched telecommunications network (but is not one).

Loopback Messages (LBM)

A Loopback Message (LBM) is a unicast message triggered by the operator issuing an operational command. LBM can be addressed to either a Maintenance End Point (MEP) or Maintenance Intermediate Point (MIP), but only a MEP can initiate an LBM. The destination MP can be addressed by its MAC address. The receiving MP responds with a Loopback Response (LBR). LBM can contain an arbitrary amount of data that can be used to diagnose faults as well as performance measurements. The receiving MP copies the data to the LBR. The system achieves fault verification through the use of Loopback Messages (LBM).

Loopback Response (LBR)

Loopback Response (LBR) is the response from a Maintenance Point (MP).

MAC-in-MAC encapsulation

MAC-in-MAC encapsulation defines a BMAC-DA and BMAC-SA to identify the backbone source and destination addresses. The originating node creates a MAC header that the device uses for delivery from end to end. As the MAC header stays the same across the network, there is no need to swap a label or do a route lookup at each node, allowing the frame to follow the most efficient forwarding path end to end.

Maintenance Associations (MA)

Maintenance Associations (MA) are administrative associations in a network that is divided by the 802.1ag Connectivity Fault Management (CFM) feature. CFM groups MAs within Maintenance Domains. Each MA is defined by a set of Maintenance Points (MP). An MP is a demarcation point on an interface that participates in CFM within an MD. Connectivity Fault Management is a mechanism to debug connectivity issues and to isolate faults within the Shortest Path Bridging MAC (SPBM) Network.

Maintenance Domains (MD)

Maintenance Domains (MD) are administrative domains that divides a network by the 802.1ag Connectivity Fault Management (CFM) feature. Each MD is further subdivided into logical groupings called Maintenance Associations (MA). Connectivity Fault Management is a mechanism to debug connectivity issues and to isolate faults within the Shortest Path Bridging MAC (SPBM) Network.

Maintenance Points (MP)

Maintenance Points (MP) are a demarcation point on an interface that participates in Connectivity Fault Management (CFM) within a Maintenance Domain (MD). There are two types of MP: Maintenance End Point (MEP) and Maintenance Intermediate Point (MIP). Connectivity Fault Management

	is a mechanism to debug connectivity issues and to isolate faults within the Shortest Path Bridging MAC (SPBM) Network.
management information base (MIB)	The MIB defines system operations and parameters used for the Simple Network Management Protocol (SNMP).
mask	A bit string that the device uses along with an IP address to indicate the number of leading bits in the address that correspond with the network part.
maximum transmission unit (MTU)	The largest number of bytes in a packet—the maximum transmission unit of the port.
Media 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 system mirrors the traffic.
mirroring port	The port to which the system mirrors all traffic, also referred to as the destination port.
mirroring VLAN	The virtual Local Area Network (VLAN) to which the system mirrors the traffic.
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 switch directs the data to an MGID. The switch maintains a table that maps MGIDs to their member ports. Both virtual LAN (VLAN) and IP multicast (IPMC) use MGIDs.
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.
Partial Sequence Number Packets (PSNP)	Partial Sequence Number Packets (PSNP) are requests for missing Link State Packets (LSPs). When a receiving router detects a missing LSP, it sends a PSNP to the router that sent the Complete Sequence Number Packets (CSNP).
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 specific layer and that consists of protocol-control information of the specific layer and possibly user data of that layer.
Provider Backbone Bridge (PBB)	To forward customer traffic across the service-provider backbone, SPBM uses IEEE 802.1ah Provider Backbone Bridging (PBB) MAC-in-MAC encapsulation, which hides the customer MAC (C-MAC) addresses in a backbone MAC (B-MAC) address pair. MAC-in-MAC encapsulation defines a BMAC-DA and BMAC-SA to identify the backbone source and destination addresses.
QSFP+	A hot pluggable, quad small form-factor pluggable plus (QSFP+) transceiver, which is used in 40 Gbps and 4x10 Gbps Ethernet applications. 4x10 Gbps requires channelization support.
QSFP28	A hot pluggable, quad small form-factor pluggable 28 (QSFP28) transceiver, which is used in 100 Gbps and 4x25 Gbps Ethernet applications. 4x25 Gbps requires channelization support. It is similar in physical appearance to QSFP+ transceivers.
quality of service (QoS)	QoS features reserve resources in a congested network, allowing you to configure a higher priority for certain devices. For example, you can configure a higher priority for IP deskphones, which need a fixed bit rate and split the remaining bandwidth between data connections if calls in the network are more important than the file transfers.
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 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 Network Monitoring (RMON)	Creates and displays alarms for user-defined events, gathers cumulative statistics for Ethernet interfaces, and tracks statistical history for Ethernet interfaces.
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. You often use RIP as a very simple interior gateway protocol (IGP) within small networks.
Secure Shell (SSH)	SSH uses encryption to provide security for remote logons and data transfer over the Internet.
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.
Service Instance Identifier (I-SID)	The SPBM B-MAC header includes a Service Instance Identifier (I-SID) with a length of 24 bits. SPBM uses this I-SID to identify and transmit any virtualized traffic in an encapsulated SPBM frame. SPBM uses I-SIDs to virtualize VLANs (Layer 2 Virtual Services Network [VSN]) or VRFs (Layer 3 Virtual Services Network [VSN]) across the MAC-in-MAC backbone. With Layer 2 VSNs, you associate the I-SID with a customer VLAN, which is then virtualized across the backbone. With Layer 3 VSNs, you associate the I-SID with a customer VRF, which is also virtualized across the backbone.
SFP	A hot pluggable, small form-factor pluggable (SFP) transceiver, which is used in Ethernet applications up to 1 Gbps.
SFP+	A hot pluggable, small form-factor pluggable plus (SFP+) transceiver, which is used in Ethernet applications up to 10 Gbps. It is similar in physical appearance to SFP transceivers.
Shortest Path Bridging (SPB)	Shortest Path Bridging is a control Link State Protocol that provides a loop-free Ethernet topology. There are two versions of Shortest Path Bridge: Shortest Path Bridging VLAN and Shortest Path Bridging MAC. Shortest Path Bridging VLAN uses the Q-in-Q frame format and encapsulates the source bridge ID into the VLAN header. Shortest Path Bridging MAC uses the 802.1 ah MAC-in-MAC frame format and encapsulates the source bridge identifier into the B-MAC header.
Shortest Path Bridging MAC (SPBM)	Shortest Path Bridging MAC (SPBM) uses the Intermediate-System-to-Intermediate-System (IS-IS) link-state routing protocol to provide a loop-free Ethernet topology that creates a shortest-path topology from every node to every other node in the network based on node MAC addresses. SPBM uses the 802.1ah MAC-in-MAC frame format and encapsulates the

	source bridge identifier into the B-MAC header. SPBM eliminates the need for multiple overlay protocols in the core of the network by reducing the core to a single Ethernet-based link-state protocol, which can provide virtualization services, both layer 2 and layer 3, using a pure Ethernet technology base.
Simple Loop Prevention Protocol (SLPP)	Simple Hello Protocol that prevents loops in a Layer 2 network (VLAN).
Simple Network Management Protocol (SNMP)	SNMP administratively monitors network performance through agents and management stations.
spanning tree	A simple, fully-connected active topology formed from the arbitrary physical topology of connected bridged Local Area Network components by relaying frames through selected bridge ports. The protocol parameters and states that are used and exchanged to facilitate the calculation of the active topology and to control the bridge relay function.
Spanning Tree Group (STG)	A collection of ports in one spanning-tree instance.
Split MultiLink Trunking (SMLT)	An extension to IEEE 802.1AX (link aggregation), provides nodal and link failure protection and flexible bandwidth scaling to improve on the level of Layer 2 resiliency.
time-to-live (TTL)	The field in a packet used to determine the valid duration for the packet. The TTL determines the packet lifetime. The system discards a packet with a TTL of zero.
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 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.
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.

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 Local Area Network (VLAN)	A Virtual Local Area Network is a group of hosts that communicate as if they are attached to the same broadcast domain regardless of their physical location. VLANs are layer 2 constructs.
virtual router	An abstract object managed by the Virtual Router Redundancy Protocol (VRRP) that acts as a default router for hosts on a shared LAN.
virtual router forwarding (VRF)	Provides traffic isolation between customers operating over the same node. Each virtual router emulates the behavior of a dedicated hardware router by providing separate routing functionality, and the network treats each VRF as a separate physical router.
Virtual Router Redundancy Protocol (VRRP)	A protocol used in static routing configurations, typically at the edge of the network. This protocol operates on multiple routers on an IP subnet and elects a primary gateway router. When the primary router fails, a backup router is quickly available to take its place.