



Monitoring Performance on Virtual Services Platform 9000

Release 4.0
NN46250-701
Issue 06.01
December 2014

© 2014 Avaya Inc.

All Rights Reserved.

Notice

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

Documentation disclaimer

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

Link disclaimer

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

Warranty

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

Licenses

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

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

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

License types

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

Heritage Nortel Software

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

Copyright

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

Third Party Components

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

Preventing Toll Fraud

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

result in substantial additional charges for your telecommunications services.

Avaya Toll Fraud intervention

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

Downloading Documentation

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

Contact Avaya Support

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

Trademarks

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

Avaya is a registered trademark of Avaya Inc.

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

Contents

Chapter 1: Introduction	11
Purpose.....	11
Related resources.....	11
Documentation.....	11
Training.....	11
Viewing Avaya Mentor videos.....	12
Support.....	12
Searching a documentation collection.....	13
Chapter 2: New in this release	14
Features.....	14
Other changes.....	18
Chapter 3: Performance and fault management fundamentals	19
Switch management tools.....	19
Dynamic network applications.....	20
Digital diagnostic monitoring.....	20
Layer 2 redundancy and High Availability clarification.....	21
Local alarms.....	21
Connectivity Fault Management.....	21
Chapter 4: Performance management chassis and port	23
Chassis performance management using ACLI.....	23
Viewing system performance.....	23
Configuring CPU-HA.....	23
Chassis performance management using EDM.....	24
Viewing system performance.....	24
Configuring CPU-HA for Layer 3 redundancy.....	24
Port performance management using ACLI.....	24
Viewing DDI module information.....	24
Viewing DDI temperature information.....	28
Viewing DDI voltage information.....	28
Port performance management using EDM.....	29
Configuring rate limits.....	30
Enabling learning limits on a port.....	30
Viewing DDI information.....	31
Chapter 5: IPFIX	35
Internet Protocol Flow Information eXport.....	35
IPFIX configuration using ACLI.....	37
Enabling IPFIX globally.....	38
Configuring IPFIX metering using filters.....	39
Configuring IPFIX on a port.....	40

Configuring IPFIX slot parameters.....	41
Configuring collector parameters.....	43
Viewing flow information.....	44
Flushing IPFIX flow information.....	46
Viewing global IPFIX information.....	47
Viewing collector information.....	47
Viewing exporter information.....	48
Viewing IPFIX information for an interface.....	49
IPFIX configuration using EDM.....	50
Enabling IPFIX globally.....	50
Configuring collector parameters.....	51
Configuring slot parameters.....	52
Configuring IPFIX on a port.....	52
Configuring IPFIX metering using filters.....	53
Chapter 6: Key Health Indicators.....	55
Key Health Indicators using ACLI.....	55
Displaying KHI information.....	55
Clearing KHI information.....	70
Displaying KHI performance information.....	70
Displaying KHI control processor information.....	77
Key Health Indicators using EDM.....	78
Clearing KHI statistics.....	78
Viewing KHI forwarding information for first generation modules.....	79
Viewing KHI forwarding information for second generation modules.....	80
Viewing protocol drop counters.....	81
Displaying KHI port information.....	83
Viewing COP statistics.....	83
Displaying second generation QE statistics.....	84
Displaying second generation IFP drop information.....	86
Displaying second generation RSP drop information.....	87
Displaying second generation RSP error information.....	89
Displaying second generation RSP forwarding information.....	90
Displaying second generation RSP statistics.....	95
Displaying second generation sierra drop statistics.....	97
Displaying second generation sierra forwarding statistics.....	98
Displaying second generation sierra extended forwarding statistics.....	98
Displaying second generation sierra state information.....	100
Displaying second generation zag status.....	102
Displaying second generation zagros error information.....	106
Displaying second generation zag forward statistics.....	107
Displaying second generation zag drop statistics.....	110
Displaying second generation zag drop extended statistics.....	112
Chapter 7: Link state change control.....	114

Link state change control using ACLI.....	114
Controlling link state changes.....	114
Displaying link state changes.....	115
Link state change control using EDM.....	116
Controlling link state changes.....	116
Chapter 8: Log and trap information.....	117
Log and trap fundamentals.....	117
Simple Network Management Protocol.....	117
Overview of traps and logs.....	118
Log message format.....	119
Log files.....	121
Log file transfer.....	122
Log configuration using ACLI.....	123
Configuring a UNIX system log and syslog host.....	124
Configuring logging.....	127
Configuring the remote host address for log transfer.....	128
Configuring system logging to external storage.....	129
Configuring system message control.....	130
Extending system message control.....	131
Viewing logs.....	132
Configuring ACLI logging.....	135
Log configuration using EDM.....	137
Configuring the system log.....	137
Configuring the system log table.....	138
SNMP trap configuration using ACLI.....	139
Configuring an SNMP host.....	140
Configuring an SNMP notify filter table.....	141
Configuring SNMP interfaces.....	142
Enabling SNMP trap logging.....	145
SNMP trap configuration using EDM.....	147
Configuring an SNMP host target address.....	147
Configuring target table parameters.....	149
Configuring an SNMP notify table.....	150
Configuring SNMP notify filter profiles.....	150
Configuring SNMP notify filter profile table parameters.....	151
Enabling SNMP trap logging.....	152
Viewing the trap sender table.....	152
Chapter 9: Remote monitoring.....	154
Remote monitoring fundamentals.....	154
RMON alarm variables.....	157
RMON configuration using ACLI.....	176
Configuring RMON alarms and events.....	176
Viewing RMON settings.....	178

RMON configuration using EDM.....	179
Enabling RMON globally.....	179
Enabling RMON history.....	180
Disabling RMON history.....	182
Viewing RMON history statistics.....	182
Viewing the RMON log.....	184
Creating an alarm.....	185
Creating a port history alarm.....	187
Viewing RMON alarms.....	188
Deleting an alarm.....	188
Creating a default RMON event.....	188
Creating a nondefault RMON event.....	189
Viewing RMON events.....	190
Deleting an event.....	191
Chapter 10: Service Level Agreement Monitor.....	192
SLA Mon server and agent.....	192
QoS tests.....	193
Limitations.....	194
SLA Mon configuration using ACLI.....	194
Configuring the SLA Mon agent.....	194
SLA Mon configuration using EDM.....	197
Configuring the SLA Mon agent.....	197
Chapter 11: Statistics.....	200
SPBM packet drop statistics.....	200
Viewing statistics using ACLI.....	201
Viewing TCP statistics.....	201
Viewing port routing statistics.....	202
Displaying bridging statistics for specific ports.....	203
Displaying DHCP-relay statistics for specific ports.....	204
Displaying DHCP-relay statistics for all interfaces.....	205
Viewing IPv6 DHCP Relay statistics.....	207
Displaying LACP statistics for specific ports.....	208
Displaying RMON statistics for specific ports.....	209
Displaying detailed statistics for ports.....	210
Displaying IS-IS statistics and counters.....	212
Displaying SPBM packet drop statistics by port.....	214
Clearing SPBM packet drop statistics.....	217
Displaying policing statistics.....	217
Clearing ACL statistics.....	218
Viewing ACE statistics.....	219
Viewing MSTP statistics.....	220
Viewing RSTP statistics.....	221
Viewing RSTP port statistics.....	221

Viewing MLT statistics.....	222
Showing OSPF error statistics on a port.....	224
Viewing OSPF interface statistics.....	225
Viewing OSPF range statistics.....	226
Viewing basic OSPF statistics for a port.....	228
Showing extended OSPF statistics.....	229
Viewing IPv6 OSPF statistics.....	230
Showing the EAPoL status of the device.....	231
Showing EAPoL authenticator statistics.....	231
Showing EAPoL session statistics.....	232
Showing RADIUS server statistics.....	233
Viewing RMON statistics.....	235
Viewing PCAP statistics.....	236
Viewing IPFIX statistics.....	237
Clearing IPFIX statistics.....	238
Clearing IPv6 statistics.....	239
Viewing multicast routing process statistics.....	240
Viewing IPv6 VRRP statistics.....	241
Viewing ICMP statistics.....	244
Viewing IPv6 statistics on an interface.....	245
Viewing statistics using EDM.....	245
Enabling RMON statistics.....	245
Disabling RMON statistics.....	246
Viewing RMON statistics.....	246
Graphing chassis statistics.....	249
Graphing port statistics.....	250
Viewing chassis system statistics.....	251
Viewing chassis SNMP statistics.....	251
Viewing chassis IP statistics.....	253
Viewing chassis ICMP In statistics.....	254
Viewing chassis ICMP Out statistics.....	255
Viewing ICMP statistics.....	256
Viewing chassis TCP statistics.....	258
Viewing chassis UDP statistics.....	259
Configuring Switch Fabric statistics capture.....	260
Viewing Switch Fabric statistics.....	261
Viewing port interface statistics.....	264
Viewing port Ethernet errors statistics.....	266
Viewing port bridging statistics.....	268
Viewing port spanning tree statistics.....	269
Viewing port routing statistics.....	270
Viewing IPv6 statistics for an interface.....	270
Viewing DHCP statistics for an interface.....	273

Viewing IPv6 DHCP Relay statistics for a port.....	273
Graphing DHCP statistics for a port.....	274
Viewing DHCP statistics for a port.....	274
Graphing DHCP statistics for a VLAN.....	275
Displaying DHCP-relay statistics for Option 82.....	275
Viewing LACP port statistics.....	278
Viewing port policer statistics.....	279
Displaying file statistics.....	279
Viewing QoS policy statistics.....	280
Graphing QoS policy statistics.....	281
Viewing statistics for a specific QoS policy.....	282
Viewing ACE port statistics.....	282
Viewing ACL statistics.....	283
Clearing ACL statistics.....	284
Viewing VLAN and Spanning Tree CIST statistics.....	284
Viewing VLAN and Spanning Tree MSTI statistics.....	285
Viewing VRRP interface stats.....	286
Viewing VRRP statistics.....	287
Viewing IPv6 VRRP statistics for an interface.....	287
Viewing IPv6 VRRP statistics.....	289
Viewing SMLT statistics.....	290
Viewing RSTP status statistics.....	291
Viewing MLT interface statistics.....	292
Viewing MLT Ethernet error statistics.....	293
Viewing RIP statistics.....	295
Viewing OSPF chassis statistics.....	295
Viewing IPv6 OSPF statistics.....	296
Graphing OSPF statistics for a VLAN.....	297
Graphing OSPF statistics for a port.....	298
Viewing BGP global stats.....	299
Viewing BGP peer general statistics.....	303
Viewing BGP peer advanced statistics.....	304
Viewing BGP peer receive statistics.....	305
Viewing BGP peer transmit statistics.....	307
Viewing statistics for a VRF.....	309
Viewing EAPoL Authenticator statistics.....	310
Viewing EAPoL diagnostic statistics.....	311
Viewing EAPoL session statistics.....	313
Showing the Authenticator session statistics.....	314
Showing RADIUS server statistics.....	315
Showing SNMP statistics.....	316
Viewing PCAP stats.....	317
Enabling multicast routing process statistics.....	318

Contents

Viewing multicast routing process statistics.....	318
Viewing IPFIX hash statistics.....	319
Viewing IPFIX exporter statistics.....	320
Displaying IS-IS system statistics.....	320
Displaying IS-IS interface counters.....	321
Displaying IS-IS interface control packets.....	322
Graphing IS-IS interface counters.....	323
Graphing IS-IS interface sending control packet statistics.....	324
Graphing IS-IS interface receiving control packet statistics.....	325
Displaying SPBM packet drop statistics by port.....	326
Resetting AbsoluteValues counter using EDM.....	326
Glossary	328

Chapter 1: Introduction

Purpose

This document describes conceptual and procedural information about the switch management tools and features that are available to monitor and manage the Avaya Virtual Services Platform 9000. Operations include the following:

- Remote Monitoring (RMON)
- Simple Network Management protocol (SNMP)
- IPFIX
- Chassis performance
- Port performance

This document also provides information about how to prevent faults and improve the performance of the Avaya Virtual Services Platform 9000. Fault management includes procedures for RMON, link state change, Key Health Indicators (KHI), and logs and traps. The fault management function supports tasks related to managing or preventing faults, troubleshooting, and monitoring and improving the performance of the network or product.

Related resources

Documentation

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

Training

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

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

Viewing Avaya Mentor videos

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

About this task

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

Procedure

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

Note:

Videos are not available for all products.

Support

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

Searching a documentation collection

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

Before you begin

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

Procedure

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

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

Chapter 2: New in this release

The following sections detail what is new in *Monitoring Performance on Avaya Virtual Services Platform 9000*, NN46250-701 for Release 4.0.

Features

See the following sections for information about feature changes.

9012QQ-2 I/O module

Release 4.0.1 introduces the second generation 9012QQ-2 Input/Output (I/O) module. The 9012QQ-2 module is a 12-port 40 gigabits per second (Gbps) module that supports 40GBASE-LR4 quad small form factor pluggable plus (QSFP+), 40GBASE-SR4 QSFP+, and QSFP+ to QSFP+ DAC (1 meter Passive DAC, 3 meter Passive DAC, and 5 meter Passive DAC) transceivers in a mixed mode module configuration.

Release 4.0.1 updates information on digital diagnostic monitoring to include QSFP+ support for the 9012QQ-2 I/O module.

Release 4.0.1 updates the `show pluggable-optical-modules detail` command to support the 9012QQ-2 module and adds the `pluggable-optical-module reset {slot/port}` command to reset the quad small form factor pluggable plus (QSFP+). For more information, see:

- [Digital diagnostic monitoring](#) on page 20.
- [Viewing DDI module information](#) on page 24.

Release 4.0.1 updates the **DDI/SFP** tab in Enterprise Device Manager to include QSFP+ support for the 9012QQ-2 I/O module. For more information, see [Viewing DDI information](#) on page 31.

Digital Diagnostic Interface (DDI)

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

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

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

For more information, see:

- [Viewing DDI module information](#) on page 24.
- [Viewing DDI information](#) on page 31.

Digital Diagnostic Monitoring (DDM)

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

- Digital Diagnostic Monitoring (DDM) is now enabled by default. The `pluggable-optical-module ddm-monitor` and `no pluggable-optical-module ddm-monitor` commands to enable and disable DDM are no longer available.
- The `pluggable-optical-module ddm-alarm-portdown` command is also no longer available.
- Release 4.0 also updates the `pluggable-optical-module ddm-monitor-interval <10..40>`. The DDM monitor interval is now 10–40 seconds. The default is 10 seconds.
- The device sends traps if you enable the DDM traps send feature using the `pluggable-optical-module ddm-traps-send` command. The device always generates logs no matter if the DDM traps send feature is enabled or disabled.
- For the `show pluggable-optical-modules basic` command, the device reports qualified optics as Avaya in the type field. The device reports non-qualified best-effort optics as a different manufacturer in the type field. Unsupported optics display as unsupported in the type field, and do not operate in the system.

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

For more information, see:

- [Digital diagnostic monitoring](#) on page 20.
- [Viewing DDI module information](#) on page 24.
- [DDI/SFP field descriptions](#) on page 32.

Interface tab

Release 4.0 updates options in the Interface tab for first generation and second generation modules. For more information, see [Viewing port interface statistics](#) on page 264 .

Key Health Indicators

Release 4.0 updates Key Health Indicators.

- Release 4.0 adds the `show khi forwarding sierra [<3-12>]` command to display Sierra data path health indicators.
- Release 4.0 updates the output for the `show khi forwarding rsp [<3-12>]` command to display differently for first generation and second generation modules.
- Release 4.0 updates the show output for the `show khi forwarding ifp [<3-12>]` command.
- Release 4.0 updates the show output for the `show khi forwarding qe [<3-12>]` command.
- Release 4.0 updates the show output for the `show khi forwarding zagros [<3-12>]` command.

- In Release 4.0, some additional context was added to the KHI performance information for ACLI procedures.

For more information, see

- [Key Health Indicators](#) on page 55.
- [Displaying KHI information](#) on page 55.
- [Viewing KHI forwarding information for first generation modules](#) on page 79.

Key Health Indicators for second generation modules

Release 4.0 adds new tabs to display Key Health Indicator (KHI) information for second generation modules. For more information, see:

- [Viewing KHI forwarding information for second generation modules](#) on page 80.
- [Displaying second generation QE statistics](#) on page 84.
- [Displaying second generation IFP drop information](#) on page 86.
- [Displaying second generation RSP drop information](#) on page 87.
- [Displaying second generation RSP error information](#) on page 89.
- [Displaying second generation RSP forwarding information](#) on page 90.
- [Displaying second generation RSP statistics](#) on page 95.
- [Displaying second generation sierra drop statistics](#) on page 97.
- [Displaying second generation sierra forwarding statistics](#) on page 98.
- [Displaying second generation sierra extended forwarding statistics](#) on page 98
- [Displaying second generation sierra state information](#) on page 100.
- [Displaying second generation zag status](#) on page 102.
- [Displaying second generation zagros error information](#) on page 106.
- [Displaying second generation zag forward statistics](#) on page 107.
- [Displaying second generation zag drop statistics](#) on page 110.
- [Displaying second generation zag drop extended statistics](#) on page 112.

Link-flap parameters update

Information updated for `link-flap-detect` parameters for ACLI and the **Link Flap** tab in EDM. The default for auto-port-down is enabled. The default for frequency is 20. For more information, see:

- [Controlling link state changes](#) on page 114.
- [Controlling link state changes](#) on page 116.

Remote monitoring

Before you delete a remote monitoring (RMON) event, remove all RMON alarms related to the RMON event. For more information, see:

- [Remote monitoring fundamentals](#) on page 154.
- [Configuring RMON alarms and events](#) on page 176.

- [Deleting an event](#) on page 191.

Rmon tab

The **RisingEventIndex** and the **FallingEventIndex** options are updated for the **RMON** tab. For more information, see [Creating an alarm](#) on page 185.

Rmon tab

Release 4.0 updates the Rmon tab to display a different output for first generation and second generation modules. For more information, see [Viewing RMON statistics](#) on page 246.

RSTP status statistics update

Release 4.0 adds the new ForwardTransitions line to the RSTP Stats table. For more information, see [Viewing RSTP status statistics](#) on page 291.

Update to show logging file CPU parameter

Release 4.0 updates the `show logging file` CPU parameter variable to `WORD<0–100>`. For more information, see [Viewing logs](#) on page 132.

SLA Mon

If you configure the SLA Mon agent address under an IP address for a VLAN or brouter, you must remove the SLA Mon address before you can remove the IP address for the VLAN or brouter. For more information, see:

- [SLA Mon server and agent](#) on page 192.
- [Configuring the SLA Mon agent](#) on page 194.
- [Configuring the SLA Mon agent](#) on page 197.

SLAParameter

Release 4.0 removes **SLAParameter** from the **SLA Monitor** tab. For more information, see [SLA Monitor field descriptions](#) on page 198.

SNMP Params field

Release 4.0 updates the SNMP **Params** field. For more information, see [Configuring an SNMP host target address](#) on page 147.

SPBM drop count

Virtual Services Platform 9000 does not support the `show isis drop-stats port unknown-unicast-sa` drop count parameter for second generation modules in this release. The device always displays the ACLI output for second generation modules as 0 for this counter. For more information, see [Displaying SPBM packet drop statistics by port](#) on page 214.

Switch Fabric statistics

Release 4.0 updates Switch Fabric statistics information. The **StatsPortId** field only allows one slot/port. For more information, see [Configuring Switch Fabric statistics capture](#) on page 260.

System tab

Release 4.0 updates options on the **System** tab. For more information, see [System field descriptions](#) on page 251.

Removed trace auto command

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

Unprintable characters in log and trace messages

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

```
CPI [07/31/14 15:53:36.225] 0x000e4609 00000000 GlobalRouter SW INFO <UNPRINTABLE>
rlogind: session 0 *IN USE* via user: \xc3\xb4\xb6\x8d\x0e\xb9+\xb1\xa2aLO~o\x3m\x9\x8c
\x05g\xd2.\xc9pSWP\xc5\xd9 ip 10.139.82.29 <UNPRINTABLE>
```

For more information, see:

- [Log message format](#) on page 119.
- [Viewing logs](#) on page 132.
- [Configuring ACLI logging](#) on page 135.

Other changes

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

Document title change

In Release 4.0, the title of this document changed from *Avaya Virtual Services Platform 9000 Performance Management, NN46250-701* to *Monitoring Performance on Avaya Virtual Services Platform 9000, NN46250-701*.

Document restructured

In Release 4.0, the content for the retired document *Avaya Virtual Services Platform 9000 Fault Management, NN46250-703* merged with *Monitoring Performance on Avaya Virtual Services Platform 9000, NN46250-701*. Chapters and sections within this document have been reorganized to bring conceptual and procedural content together by feature.

Service Level Agreement name change

All occurrences in the documentation for Service Level Agreement (SLA) agent changed to SLA Mon agent. The name for the SLA Mon server changed to Avaya Diagnostic Server.

Chapter 3: Performance and fault management fundamentals

This chapter includes information about the various management tools and features that are available to monitor and manage your routing switch.

Switch management tools

Use Avaya Command Line Interface, Enterprise Device Manager, or Configuration and Orchestration Manager to access, manage, and monitor the Avaya Virtual Services Platform 9000.

Avaya Command Line Interface

To access the Avaya Command Line Interface (ACLI) initially, you need a direct connection to the system from a terminal or PC. After you enable Telnet, you can access the ACLI from a Telnet session on the network.

ACLI contains commands to configure system operations and management access. ACLI has five major command modes with different privileges.

For more information about ACLI, see *Using ACLI and EDM on Avaya Virtual Services Platform 9000*, NN46250-103.

Enterprise Device Manager

Enterprise Device Manager (EDM) is a Web-based graphical user interface (GUI) tool that operates with a Web browser. Use it to access, manage, and monitor a single Virtual Services Platform 9000 system on your network from various locations within the network.

For more information about EDM, see *Using ACLI and EDM on Avaya Virtual Services Platform 9000*, NN46250-103.

Configuration and Orchestration Manager

Configuration and Orchestration Manager (COM) is a Web-based GUI tool that operates with a Web browser. Use it to access, manage, and monitor multiple devices on your network from various locations within the network.

To access the Web management interface, you need a Web browser and an IP address for the switch. For more information, see *Using ACLI and EDM on Avaya Virtual Services Platform 9000*, NN46250-103.

Dynamic network applications

The remote access services supported on Virtual Services Platform 9000, for example the File Transfer Protocol (FTP), Trivial FTP (TFTP), rlogin, and Telnet, use daemons. These remote access daemons are not enabled by default to enhance security.

After you disable a daemon flag, all existing connections abruptly terminate, and the daemon remains idle (accepts no connection requests). Additionally, if Central Processing Unit High Availability (CPU-HA) is on and you disable a daemon, all existing connections, even those to the standby Control Processor (CP) module, immediately terminate.

Use the following dynamic network applications to manage remote access services:

- Access policies
- Port lock
- ACLI access
- SNMP community strings
- Web management interface access

For more information about how to enable remote access services, see *Quick Start Configuration for Avaya Virtual Services Platform 9000*, NN46250-102.

For more information about how to access policies, lock a port, access the ACLI, and configure SNMP community strings, see *Configuring Security on Avaya Virtual Services Platform 9000*, NN46250-601.

For more information about how to access the Web management interface, see *Using ACLI and EDM on Avaya Virtual Services Platform 9000*, NN46250-103.

Digital diagnostic monitoring

Use Digital Diagnostic Monitoring (DDM) to monitor laser operating characteristics such as temperature, voltage, current, and power. DDM is enabled by default. This feature works at any time during active laser operation without affecting data traffic.

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

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

`pluggable-optical-module ddm-monitor-interval <10..40>` command. The DDM monitor interval is 10–40 seconds. The default is 10 seconds.

There are three optical transceivers that support DDM: 40 Gigabit quad small form-factor pluggable plus (QSFP+) transceivers, 10 Gigabit small form-factor pluggable plus (SFP+), and small form-factor pluggable (SFP).

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

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

Layer 2 redundancy and High Availability clarification

Layer 2 redundancy supports the synchronization of virtual local area network (VLAN) and Quality of Service (QoS) software parameters. Layer 3 redundancy, called Central Processor Unit (CPU)-High Availability (HA) supports the synchronization of application configuration data. For full HA applications, the run-time state is also synchronized and used when switchover occurs. Non-full HA applications like BGP and PIM restart on the new Master CP using the synchronized configuration.

When you enable CPU-HA, and save the configuration file, the system writes the configuration to the Master and Standby CP. The Standby CP resets automatically and begins synchronizing with the Master CP.

Local alarms

Avaya Virtual Services Platform 9000 contains a local alarms mechanism. Local alarms are raised and cleared by applications running on the switch. Active alarms are viewed using the `show alarm database` command in ACLI. Local alarms are an automatic mechanism run by the system that do not require any additional user configuration. Check local alarms occasionally to ensure no alarms require additional operator attention. The raising and clearing of local alarms also creates a log entry for each event.

Connectivity Fault Management

The Shortest Path Bridging MAC (SPBM) network needs a mechanism to debug connectivity issues and isolate faults. This function is performed at Layer 2, not Layer 3. Connectivity Fault Management (CFM) operates at Layer 2 and provides an equivalent of the `ping` and `traceroute` commands. To support troubleshooting of the SPBM cloud, Virtual Services Platform 9000 supports

a subset of CFM functionality. For more information about CFM see *Avaya Virtual Services Platform 9000 Configuration — Shortest Path Bridging MAC (SPBM)*, NN46250–510.

Chapter 4: Performance management chassis and port

This chapter includes EDM and ACLI procedures to configure chassis parameters and port performance management.

Chassis performance management using ACLI

You can use ACLI to configure chassis parameters on the Avaya Virtual Services Platform 9000.

Viewing system performance

For information about how to use Key Health Indicators functionality to view system performance, see [Key Health Indicators using EDM](#) on page 78.

Configuring CPU-HA

When you enable CPU-HA, and save the configuration file, the system writes the configuration to the Master and Standby CP. The Standby CP resets automatically and begins synchronizing with the Master CP.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enable HA-CPU:

```
boot config flags ha-cpu
```

Example

```
VSP-9012:1(config)#boot config flags ha-cpu
The config file on the Master will be overwritten with the current active configuration.
-Layer 2/3 features will be enabled in L2/L3 redundancy mode.
Do you want to continue (y/n) ?y
```

Chassis performance management using EDM

Use Enterprise Device Manager (EDM) to configure chassis parameters and to graph chassis statistics on an Avaya Virtual Services Platform 9000.

Viewing system performance

For information about how to use Key Health Indicators functionality to view system performance, see [Key Health Indicators using EDM](#) on page 78.

Configuring CPU-HA for Layer 3 redundancy

When you enable CPU-HA, and save the configuration file, the system writes the configuration to the Master and Standby CP. The Standby CP resets automatically and begins synchronizing with the Master CP.

For more information about HA CPU Layer 3 redundancy, see *Administering Avaya Virtual Services Platform 9000*, NN46250-600.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > Edit**.
2. Click **Chassis**.
3. Click the **System Flags** tab.
4. Beside HaCpu, select **enable**.
5. Click **Apply**.

Port performance management using ACLI

This section contains procedures to configure port performance management in the ACLI.

Viewing DDI module information

Perform this procedure to view basic QSFP+, SFP+, and SFP manufacturing information and characteristics, and the current configuration.

About this task

This command displays information for both DDI and non-DDI QSFP+s, SFP+s, and SFPs.

Note:

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

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

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

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

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View basic QSFP+, SFP+, and SFP manufacturing information and characteristics:

```
show pluggable-optical-modules basic [{slot/port[-slot/port]}[,...]]
```

3. View configuration information:

```
show pluggable-optical-modules config
```

4. View detailed QSFP+, SFP+, and SFP manufacturing information and characteristics:

```
show pluggable-optical-modules detail [{slot/port[-slot/port]}[,...]]
```

Example

Display QSFP+, SFP+, and SFP manufacturing information, characteristics and configuration information.

```
Switch:1#show pluggable-optical-modules basic
```

```
=====
                          Pluggable Optical Module Info
=====
```

PORT NUM	TYPE	DDM ENABLED	VENDOR NAME	PART NUMBER
3/16	10GbCX	FALSE	Amphenol	574790001
3/17	10GbCX	FALSE	Amphenol	574790001
3/32	10GbCX	FALSE	Amphenol	574790001
3/33	10GbCX	FALSE	Amphenol	574790001
8/1	GbicSx	TRUE	Avaya	AA1419048-E6
8/12	GbicSx	TRUE	Avaya	AA1419048-E6
8/14	GbicSx	TRUE	Avaya	AA1419048-E6
8/15	GbicSx	TRUE	Avaya	AA1419048-E6
8/16	GbicSx	TRUE	Avaya	AA1419048-E6

Performance management chassis and port

```

8/18 Unsupported TRUE FINISAR CORP. FTLX8571D3BCL-N2
8/20 GbicSx TRUE Avaya AA1419048-E6
8/24 GbicSx TRUE Avaya AA1419048-E6
8/25 GbicSx TRUE Avaya AA1419048-E6
8/36 GbicSx TRUE Avaya AA1419048-E6
8/37 GbicSx TRUE Avaya AA1419048-E6
8/48 GbicSx TRUE Avaya AA1419048-E6
9/1 GbicSx TRUE Avaya AA1419048-E6
9/13 GbicSx TRUE Avaya AA1419048-E6
9/21 GbicSx TRUE Avaya AA1419048-E6
11/1 GbicSx TRUE Avaya AA1419048-E6
11/9 GbicSx TRUE Avaya AA1419048-E6
11/13 GbicSx TRUE Avaya AA1419048-E6
11/17 GbicSx TRUE Avaya AA1419048-E6
11/21 GbicSx TRUE Avaya AA1419048-E6

```

```
Switch:1#show pluggable-optical-modules config
```

```

=====
Pluggable Optical Module Global Configuration
=====

```

```

          ddm-monitor : enabled
ddm-monitor-interval : 10
          ddm-traps-send : enabled
ddm-alarm-portdown : enabled

```

```
Switch:1>enable
```

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

```
Switch:1(config)#show pluggable-optical-modules detail
```

```

=====
Pluggable Optical Module Info 7/3 Detail
=====

```

```

Port: 7/3
Type: 10GbSR
DDM Enabled : TRUE
Avaya PEC : AA1403015-E6          CLEI      : IPU3ADZF5AA
Vendor   : Avaya                 Vendor PN : AA1403015-E6
Vendor REV : A                   Vendor SN : FNSRMYAJ5022L
Vendor Date : 07/27/10
Wavelength : 850.00 nm

```

```
Digital Diagnostic Interface Supported
```

```

Optics Status           : Ok
Calibration              : Internal
RX Power Measurement    : Average
Auxiliary 1 Monitoring   : Not Implemented
Auxiliary 2 Monitoring   : Not Implemented

```

```

-----
          LOW_ALARM LOW_WARN   ACTUAL   HIGH_WARN HIGH_ALARM THRESHOLD
          THRESHOLD THRESHOLD  VALUE   THRESHOLD THRESHOLD  STATUS
-----
Temp (C)      0.0      5.0      32.6367  73.0      78.0      Normal
Voltage (V)   3.0350   3.1000   3.2694   3.5000   3.5650   Normal
Bias (mA)     4.0       5.0      8.2400  10.8000  11.8000  Normal
TxPower (dBm) -9.3000  -8.3000  -2.7000   0.0       1.0       Normal
RxPower (dBm) -12.9000 -11.9000 -2.2000  -2.0      -1.0      Normal

```

```

--More-- (q = quit)

Switch:1(config)#show pluggable-optical-modules detail

=====
=====
                          Pluggable Optical Module Info 4/1 Detail
=====
=====
Port: 4/1
Type: 40GbSR4
DDM Enabled : TRUE
Avaya PEC   : AA1404005-E6          CLEI       : n/a
Vendor     : AVAGO                 Vendor PN  : AA1404005-E6
Vendor REV : 01                    Vendor SN  : AVAGPH14160000N
Vendor Date: 02/16/14
Wavelength : 850.00 nm

Digital Diagnostic Interface Supported

Optics Status           : Ok
Calibration             : Internal
RX Power Measurement    : Average
Auxiliary 1 Monitoring  : Not Implemented
Auxiliary 2 Monitoring  : Not Implemented

-----
-----
                LOW_ALARM LOW_WARN   ACTUAL  HIGH_WARN HIGH_ALARM
THRESHOLD

--More-- (q = quit)

```

Variable definitions

Use the data in the following table to use the **show pluggable-optical-modules basic** and **show pluggable-optical-modules detail** commands.

Variable	Value
basic	Displays basic QSFP+, SFP+, and SFP manufacturing information and characteristics. The device reports qualified optics as Avaya in the type field. The device reports non-qualified best-effort optics as a different manufacturer in the type field. Unsupported optics display as unsupported in the type field, and do not operate in the system.
config	Displays configuration information.
detail	Displays detailed QSFP+, SFP+, and SFP manufacturing information and characteristics.
{slot/port[-slot/port][,...]}	Specify a port or a range of ports in the format of slot/port. If you do not specify a port list, the system displays the complete detailed output for each port.

Viewing DDI temperature information

Perform this procedure to view SFP and SFP+ temperatures.

About this task

This command displays information for both DDI and non-DDI SFPs and SFP+s.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View SFP and SFP+ temperatures:

```
show pluggable-optical-modules temperature [{slot/port[-slot/port]}
[,...]]
```

Example

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

```
=====
                        Pluggable Optical Module Temperature (C)
=====
```

PORT NUM	LOW ALARM THRESHOLD	LOW WARN THRESHOLD	ACTUAL VALUE	HIGH WARN THRESHOLD	HIGH ALARM THRESHOLD	THRESHOLD STATUS
3/1	0.0	5.0	34.0625	73.0	78.0	Normal
3/8	0.0	5.0	29.3632	73.0	78.0	Normal
3/9	0.0	5.0	31.1679	73.0	78.0	Normal
3/10	0.0	5.0	30.7070	73.0	78.0	Normal
4/1	-10.0	-5.0	29.0859	88.0	93.0	Normal
4/25	-25.0	-20.0	27.7695	90.0	95.0	Normal
4/37	-25.0	-20.0	31.8632	90.0	95.0	Normal
4/48	-25.0	-20.0	31.0585	90.0	95.0	Normal
6/43	-25.0	-20.0	28.9140	90.0	95.0	Normal

Variable definitions

Use the data in the following table to use the **show pluggable-optical-modules temperature** command.

Variable	Value
{slot/port[-slot/port][,...]}	Specify a port or a range of ports in the format of slot/port. If you do not specify a port list, the system displays the complete detailed output for each port.

Viewing DDI voltage information

Perform this procedure to view SFP and SFP+ voltages.

About this task

This command displays information for both DDI and non-DDI SFPs and SFP+s.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View SFP and SFP+ voltages:

```
show pluggable-optical-modules voltage [{slot/port[-slot/port]}
[, ...]]
```

Example

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

```
=====
                        Pluggable Optical Module Voltage (V)
=====
PORT          LOW ALARM  LOW WARN  ACTUAL  HIGH WARN  HIGH ALARM  THRESHOLD
NUM           THRESHOLD THRESHOLD VALUE    THRESHOLD THRESHOLD  STATUS
-----
3/1            3.0350    3.1000   3.2873   3.5000    3.5650    Normal
3/8            3.0350    3.1000   3.3101   3.5000    3.5650    Normal
3/9            3.0350    3.1000   3.2901   3.5000    3.5650    Normal
3/10           3.0350    3.1000   3.2852   3.5000    3.5650    Normal
4/1            3.0350    3.1000   3.2942   3.5000    3.5650    Normal
4/25           2.7000    2.9000   3.2859   3.7000    3.9000    Normal
4/37           2.7000    2.9000   3.2748   3.7000    3.9000    Normal
4/48           2.7000    2.9000   3.2768   3.7000    3.9000    Normal
6/43           2.7000    2.9000   3.2854   3.7000    3.9000    Normal
```

Variable definitions

Use the data in the following table to use the **show pluggable-optical-modules voltage** command.

Variable	Value
{slot/port[-slot/port][, ...]}	Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2). If you do not specify a port list, the system displays the complete detailed output for each port.

Port performance management using EDM

This section describes port performance management functions on an Avaya Virtual Services Platform 9000.

Configuring rate limits

Configure the rate limit of broadcast or multicast packets to determine the total bandwidth limit on the port.

Procedure

1. On the Device Physical View, select a port or multiple ports.
2. In the navigation pane, expand the following folders: **Configuration > Edit > Port**.
3. Click **General**.
4. Click the **Rate Limiting** tab.
5. Configure the parameters as required.
6. Click **Apply**.

Rate Limiting field descriptions

Use the data in the following table to use the **Rate Limiting** tab.

Name	Description
Index	The port number.
TrafficType	The type of traffic being rate limited, either broadcast or multicast traffic. The default is broadcast.
AllowedRatekPps	<p>This variable is the allowed traffic rate limit for the port in packets per second.</p> <p>For the Avaya Virtual Services Platform 9000, 1 to 25 sets the limit in a percentage of the total bandwidth on the port from 1–25 percent.</p> <p>On gigabit ports and MDAs, there can be up to a 2 percent difference between the configured and actual rate limiting values.</p> <p>For the Avaya Virtual Services Platform 9000, 1–65535 sets the limit in packets for each second.</p>
Enable	Double-click in the field and select to enable (True) or disable (False) rate limiting. The default is false.

Enabling learning limits on a port

Limit MAC address learning to limit the number of forwarding database (FDB) entries learned on a particular port to a user-specified value. After the number of learned forwarding database entries reaches the maximum limit, packets with unknown source MAC addresses are flooded to all member ports.

Procedure

1. In the Device Physical View tab, select a port or multiple ports.
2. In the navigation pane, expand the following folders: **Configuraton** > **Edit** > **Port**.
3. Click **General**.
4. Click the **Limit-Learning** tab.
5. Configure the parameters as required.

Limit-Learning field descriptions

Use the data in the following table to use the **Limit-Learning** tab.

Name	Description
PortNum	Shows the slot and port number to configure.
MaxMacCount	Configures the number of entries in the MAC table for the port that causes learning to stop. The default is 1024.
MinMacCount	Configures the number of entries in the MAC table for the port at which learning can resume. The default is 512.
CurrentMacCount	Shows the number of entries currently in the MAC table for the port.
Enable	Enables or disables limit learning for the port.
MacLearning	Shows if MAC learning is enabled or disabled for the port.
ViolationLogTrap	Configures the system to send a trap to the management station after a MAC address violation is detected on the port. The default is disable.
ViolationDownPort	Configures the system to disable the port after a MAC address violation is detected. The default is disable.

Viewing DDI information

You can view DDI information (such as module information, temperature, and voltages) for transceivers on the interface modules.

Note:

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

To minimize false warnings or alarms due to this inaccuracy, the thresholds for low and high TxPower and for low RxPower are offset by this tolerance. High RxPower thresholds are not

offset due to the potential for receiver saturation and damage that can result over the long-term, however, this increases the possibility of false alerts.

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

Procedure

1. In the Physical Device view, select a port.
2. In the navigation pane, expand the following folders: **Configuration > Edit > Port**.
3. Click **General**.
4. Select the **DDI/SFP** tab.

DDI/SFP field descriptions

Use the data in the following table to use the **DDI/SFP** tab.

Name	Description
DdmStatus	Indicates if DDM is enabled. DDM is enabled by default.
Calibration	Indicates if the calibration is internal or external.
PowerMeasure	Indicates Rx power measurement as average or OMA.
ConnectorType	Indicates the type of connector.
VendorName	Indicates the name of the manufacturer.
VendorPartNumber	Indicates the Avaya PEC.
VendorRevNumber	Indicates the manufacturer revision level.
VendorSN	Indicates the manufacturer serial number.
VendorDateCode	Indicates the manufacturer date code.
CLEI	Indicates the Telcordia register assignment Avaya CLEI code.
SupportsDDM	Indicates if DDM is supported.
Aux1Monitoring	Indicates if auxiliary monitoring is implemented.
Aux2Monitoring	Indicates if auxiliary monitoring is implemented.
Wavelength	Indicates the wavelength in nm.
Temperature	Indicates the current temperature in degrees Celsius.
TemperatureHighAlarmThreshold	Indicates the high alarm threshold in degrees Celsius.
TemperatureLowAlarmThreshold	Indicates the low alarm threshold in degrees Celsius.
TemperatureHighWarningThreshold	Indicates the high warning threshold in degrees Celsius.
TemperatureLowWarningThreshold	Indicates the high warning threshold in degrees Celsius.
TemperatureStatus	Indicates if any temperature thresholds were exceeded.
Voltage	Indicates the current voltage in volts.
VoltageHighAlarmThreshold	Indicates the high alarm threshold in volts.

Name	Description
VoltageLowAlarmThreshold	Indicates the low alarm threshold in volts.
VoltageHighWarningThreshold	Indicates the high warning threshold in volts.
VoltageLowWarningThreshold	Indicates the high warning threshold in volts.
VoltageStatus	Indicates if any voltage thresholds were exceeded.
Bias	Indicates the laser bias current in mA.
BiasHighAlarmThreshold	Indicates the bias current high alarm threshold in mA.
BiasLowAlarmThreshold	Indicates the bias current low alarm threshold in mA.
BiasHighWarningThreshold	Indicates the bias current high warning threshold in mA.
BiasLowWarningThreshold	Indicates the bias current high warning threshold in mA.
BiasStatus	Indicates if any bias thresholds were exceeded.
TxPower	Indicates the current Tx power in mW.
TxPowerHighAlarmThreshold	Indicates the high alarm threshold in mW for the Tx power.
TxPowerLowAlarmThreshold	Indicates the low alarm threshold in mW for the Tx power.
TxPowerHighWarningThreshold	Indicates the high warning threshold in mW for the Tx power.
TxPowerLowWarningThreshold	Indicates the high warning threshold in mW for the Tx power.
TxPowerStatus	Indicates if any Tx power thresholds were exceeded.
RxPower	Indicates the current Rx power in mW.
RxPowerHighAlarmThreshold	Indicates the high alarm threshold in mW for the Rx power.
RxPowerLowAlarmThreshold	Indicates the low alarm threshold in mW for the Rx power.
RxPowerHighWarningThreshold	Indicates the high warning threshold in mW for the Rx power.
RxPowerLowWarningThreshold	Indicates the high warning threshold in mW for the Rx power.
RxPowerStatus	Indicates if any Rx power thresholds were exceeded.
Aux1	Indicates the current auxiliary 1 reading.
Aux1HighAlarmThreshold	Indicates the high alarm threshold auxiliary 1 reading.
Aux1LowAlarmThreshold	Indicates the low alarm threshold auxiliary 1 reading.
Aux1HighWarningThreshold	Indicates the high warning threshold auxiliary 1 reading.
Aux1LowWarningThreshold	Indicates the high warning threshold auxiliary 1 reading.
Aux1Status	Indicates if any auxiliary 1 thresholds were exceeded.
Aux2	Indicates the current auxiliary 2 reading.
Aux2rHighAlarmThreshold	Indicates the high alarm threshold auxiliary 2 reading.
Aux2LowAlarmThreshold	Indicates the low alarm threshold auxiliary 2 reading.
Aux2HighWarningThreshold	Indicates the high warning threshold auxiliary 2 reading.
Aux2LowWarningThreshold	Indicates the high warning threshold auxiliary 2 reading.

Name	Description
Aux2rStatus	Indicates if any auxiliary 2 thresholds were exceeded.

Chapter 5: IPFIX

Internet Protocol Flow Information eXport (IPFIX) is an Internet Engineering Task Force (IETF) standard to export IP flow information. Virtual Services Platform 9000 supports the Netflow V9 7 format.

Use the procedures in this chapter to configure IPFIX to monitor IP flows for the Avaya Virtual Services Platform 9000 using Enterprise Device Manager (EDM) and Avaya command line interface (ACLI).

Internet Protocol Flow Information eXport

Internet Protocol Flow Information eXport (IPFIX) is an Internet Engineering Task Force (IETF) standard to export IP flow information. Virtual Services Platform 9000 supports the Netflow V9 format.

Note:

Virtual Services Platform 9000 does not support IPFIX on IS-IS interfaces.

An IP flow is a set of packets, with the following common properties, that are sent over a period of time:

- source IP address
- destination IP address
- protocol type
- source protocol port
- destination protocol port
- ingress VLAN ID
- ingress port and observation point (VLAN or port)

You can view the flows and you can export the flow information periodically to one or more third-party collectors. A collector can store a large number of flow records from several devices in the network. The IPFIX standard specifies the protocol for exporting the flows to collector, including the formatting of flow records and the underlying transport protocols, such as UDP.

Use the collected information for network planning, troubleshooting a live network, and monitoring security threats. To use the information for accurate billing requires continuous monitoring of every packet in a flow. In the case of hash collisions, the packets are counted in a separate statistic (hash

collisions) and not attributed to a particular flow. Avaya recommends that you do not use this information for billing purposes.

However, for a second generation I/O modules, chances of hash collisions are less likely because the IPFIX flow table is much larger. Therefore, the IPFIX hash collision statistics are not maintained for second generation I/O modules.

Applications

Applications, such as Avaya IP Flow Manager (IPFM), collect flow information from the collector and use it to monitor network flow volume. IPFM can process information it receives to generate textual or graphical displays of traffic patterns. The following list provides examples of the information the application can collect:

- top 10 conversations
- top 10 applications
- top 10 hosts
- top 10 ports
- top 10 protocols
- top 10 subnets

IPFIX specifies requirements to meter flows, to export or report flows to a collector, and for the interface between the exporter and collector. Avaya switching platforms run the metering and reporting processes, and the collector runs on a server or an appliance.

Export flow information to a third-party collector or a local collector by exporting as you do for the Netflow V9 format by using the User Datagram Protocol (UDP) as the transport protocol. The collector and report—generating applications are developed by third parties. The following figure shows the IPFIX data collection process.

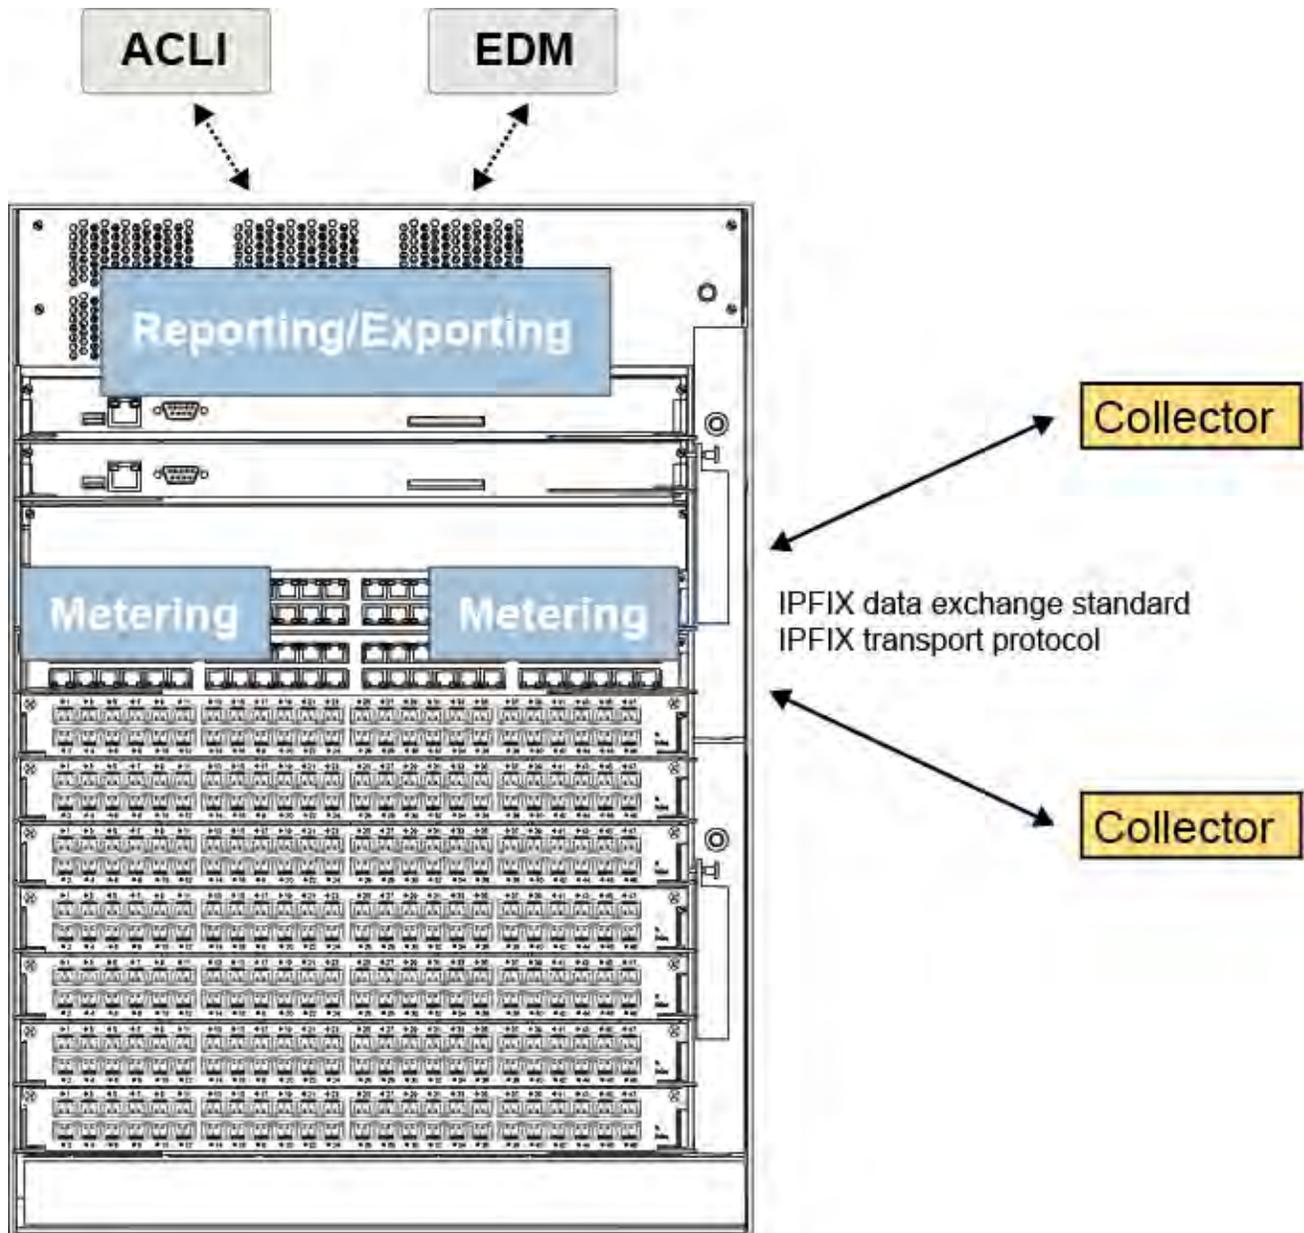


Figure 1: IPFIX collecting and reporting functions for first generation I/O modules

Note:

For second generation I/O modules, Reporting and Exporting are done by the route switch processor (RSP).

IPFIX configuration using ACLI

Use IPFIX on the Virtual Services Platform 9000 to monitor IP flows.

You can use the procedures in this section to perform the following actions:

- Configure filters to match a set of flows and perform ingress metering on those flows.
- Configure IPFIX directly on a port.
- Configure a sampling rate to prevent continuous monitoring.
- View a limited display of flow information using ACLI.
- Export the flow information to a third party collector as in the Netflow V9 format, which uses UDP as transport.
- Maintain the following statistics for each flow:
 - aggregate byte count
 - aggregate packet count
 - time stamp for the start of the flow
 - time stamp for the last time this flow was observed

Enabling IPFIX globally

You must globally enable IPFIX before the software will perform IPFIX functions in filters or on a port.

Note:

Virtual Services Platform 9000 does not support IPFIX on IS-IS interfaces.

The second generation I/O modules do not display flow records or hash statistics. Use an IPFIX Collector for IPFIX flow information.

Note:

The second generation I/O modules does not support `show ip ipfix flows` command.

About this task

After you globally disable IPFIX, even with the exporter enabled or filters configured with IPFIX enabled, the software does not perform IPFIX functions. If IPFIX is enabled on the system and you use a global disable command, a warning message appears before the system disables IPFIX.

The default global state is disabled.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enable IPFIX:

```
ip ipfix enable
```

Example

```
VSP-9012:1(config)#ip ipfix enable
```

Configuring IPFIX metering using filters

Configure IPFIX metering using filters to use IPFIX on selected flows. An ACL can have multiple ACEs and each ACE can have an action of ipfix-enable. The default value of ipfix-action is disable.

Before you begin

- An ACL exists with match criteria.
- You must enable IPFIX globally before you can use it in a filter.

About this task

A packet can match multiple ACEs in an ACL. The system performs the regular actions you configure in the filter. If multiple ACEs have an action of ipfix-enable, the system performs metering only once for a packet. A packet matches multiple ACEs because you configure the ACEs to match overlapping flows. IPFIX metering further categorizes this packet into a flow record based on the unique ipfix-handle. The packet matches n ACEs that correspond to n different ACE flows, but it is still a single IPFIX flow.

The default sampling-rate rate for first generation I/O module is 1 and for second generation I/O module is 50.

Note:

For second generation I/O modules, with IPFIX enabled with a sampling-rate below 50, there can be loss seen at high traffic rates. You can avoid this by assigning the sampling-rate to a value of at least 50.

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

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure the global action to specify packet treatment if a packet matches an ACE:

```
filter acl set <1-2048> global-action ipfix-enable
```

3. Configure ACE actions to meter flows after a packet matches an ACE:

```
filter acl ace action <1-2048> <1-2000> permit ipfix-enable
```

Example

```
VSP-9012:1(config)#filter acl set 3 global-action ipfix-enable
VSP-9012:1(config)#filter acl ace action 2 3 permit ipfix-enable
```

Variable definitions

Use the data in the following table to use the filtering commands for IPFIX.

Variable	Value
<1–2048>	Specifies the ACL ID.
<1–2000>	Specifies the ACE ID. Use ACE IDs 1–1000 for security rules. Use ACE IDs 1001–2000 for QoS rules.
ipfix-enable	Enables IPFIX metering. The default is disabled.
permit	Configures the action mode for security ACEs. The default value is deny. Each ACE has a mode of permit or deny the matched traffic. You can use filters to configure metering of permitted traffic. If you need to enable IPFIX on denied traffic, you must enable it on an individual port basis, which enables IPFIX monitoring on all traffic that enters a port.

Configuring IPFIX on a port

Configure IPFIX on a port to meter IP flows on the port.

Note:

Virtual Services Platform 9000 does not support IPFIX on IS-IS interfaces.

Before you begin

- You must enable IPFIX globally before you can use it on a port.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

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

2. Configure the sampling rate:

```
ip ipfix [port {slot/port[-slot/port][,...]}] sampling-rate <1-100000>
```

The default sampling-rate rate for first generation I/O module is 1 and for second generation I/O module is 50.

Note:

For second generation I/O modules, with IPFIX enabled with a sampling-rate below 50, there can be loss seen at high traffic rates. You can avoid this by assigning the sampling-rate to a value of at least 50.

3. Enable IPFIX on the port:

```
ip ipfix [port {slot/port[-slot/port][,...]}] enable
```

Example

```
VSP-9012:1(config)#interface gigabitEthernet 4/6
```

```
VSP-9012:1(config-if)#ip ipfix port 4/6-4/8 enable sampling-rate 5
```

Variable definitions

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

Variable	Value
sampling-rate <1-100000>	For the first generation I/O module, configure the sampling rate for metering on the port, as one in every n packets. The default value is 1, which configures continuous monitoring. For the second generation I/O module, configure the sampling rate for metering on the port, as fifty in every n packets. The default value is 50, which configures continuous monitoring.
{slot/port[-slot/port][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).

Configuring IPFIX slot parameters

Configure an exporter slot to send data from an exporter to a collector.

Before you begin

- You must enable IPFIX globally before the IPFIX slot configuration will take affect.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure the timeout value for flows:

```
ip ipfix slot {slot[-slot][,...]} active-timeout <1-60>
```

3. Configure the aging interval for flow records:

```
ip ipfix slot {slot[-slot][,...]} aging-interval <10-3600>
```

4. Configure the frequency to export to the collector:

```
ip ipfix slot {slot[-slot][,...]} export-interval <10-3600>
```

5. Enable the slot to export flow information:

```
ip ipfix slot {slot[-slot][,...]} exporter-enable
```

6. Configure the template refresh period based on an interval or number of packets:

```
ip ipfix slot {slot[-slot][,...]} template-refresh-interval <60-3600>
template-refresh-packets <1-600>
```

Note:

The second generation I/O modules support only IPFIX template-refresh-interval, it does not support IPFIX template-refresh-packets. You can configure the template refresh period based on an interval.

Example

```
VSP-9012:1(config)#ip ipfix slot 5 active-timeout 4
```

```
VSP-9012:1(config)#ip ipfix slot 5 export-interval 60
```

```
VSP-9012:1(config)#ip ipfix slot 5 exporter-enable
```

```
VSP-9012:1(config)#ip ipfix slot 5 template-refresh-interval 70
template-refresh-packets 40
```

Variable definitions

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

Variable	Value
active-timeout <1-60>	Configures the flow active timeout. The default is 30 minutes.
aging-interval <10-3600>	Configures the flow record aging interval. The default is 15 seconds.
export-interval <10-3600>	Configure the interval at which to export flow information. The default is 50 seconds.
exporter-enable	Enables the exporter state for the slot. The default is enable.
{slot[-slot][,...]}	Identifies the slot in one of the following formats: a single slot (3), a range of slots (3-4), or a series of slots (3,5,6). Valid slots are 3-12.
template-refresh-interval <60-3600>	Configures the template refresh timeout. The template is sent to the collector every x seconds or after x exported packets, whichever occurs first. The default is 60 seconds.

Variable	Value
template-refresh-packets <1-600>	Configures the template refresh timeout. The template is sent to the collector every x seconds or after x exported packets, whichever occurs first. The default is 20 packets.

Configuring collector parameters

Configure collector parameters to determine to which collector an interface module exports flow information. You can configure up to two collectors for each interface module.

About this task

Specify an exporter IP address to configure the source address in the IPFIX packets the interface module sends to the collectors. If you do not specify an exporter IP address, the source IP address is chosen from virtual IP, management IP, or outgoing interface IP based on the collector IP reachability.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Specify the destination port and exporter IP address:

```
ip ipfix collector {slot[-slot][, ...]} {A.B.C.D} dest-port <1-65535>
[exporter-ip {A.B.C.D}]
```

3. Specify the protocol:

```
ip ipfix collector {slot[-slot][, ...]} {A.B.C.D} protocol udp
```

4. Enable the collector status:

```
ip ipfix collector {slot[-slot][, ...]} {A.B.C.D} enable
```

Example

```
VSP-9012:1(config)#ip ipfix collector 6 47.17.143.146 dest-port 9995
exporter-ip 47.17.159.20
```

```
VSP-9012:1(config)#ip ipfix collector 6 47.17.143.146 enable
```

Variable definitions

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

Variable	Value
{A.B.C.D}	Specifies the IP address of the collector.
dest-port <1-65535>	Specifies the destination port.

Variable	Value
exporter-ip {A.B.C.D}	Specifies the IP address for the exported traffic.
{slot[-slot][,...]}	Identifies the slot in one of the following formats: a single slot (3), a range of slots (3-4), or a series of slots (3,5,6).

Viewing flow information

View flow information to see the flow entries. The flow database is large. The functionality is simple in terms of sorting. The response time can be slow for sorted displays.

About this task

This command displays records from only one interface module at a time. This command uses the following type of optional fields:

- Fields that specify match fields. These fields can be an exact match or an operator like LE, GE, EQ, NE.

Note:

For the second generation I/O modules you cannot display flow records or hash statistics. Thus, second generation I/O modules does not support `show ip ipfix flows` command.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View flow information:

Note:

The `show ip ipfix flows numflows 0` displays the number of IPFIX flows in the IPFIX flow table of the RSP.

```
show ip ipfix flows {slot[-slot][,...]} [byte-count WORD<1-2>
<0-4294967295>] [dest-addr WORD<1-2> {A.B.C.D}] [first-pkt-time
WORD<1-2> <MMddyyyyhhmmss>] [last-pkt-time WORD<1-2>
<MMddyyyyhhmmss>] [monitor <false|true>] [numflows <0-16000>] [pkt-
count WORD<1-2> <0-4294967295>] [port WORD<1-2> {slot/port}]
[protocol WORD<1-2> <0-255>] [source-addr WORD<1-2> {A.B.C.D}] [TCP-
UDP-dest-port WORD<1-2> <0-65535>] [TCP-UDP-src-port WORD<1-2>
<0-65535>] [TOS WORD<1-2> <0-255>] [vlan WORD<1-2> <1-4084>]
```

Example

```
VSP-9012:1#show ip ipfix flows 4
```

```
=====
IPFIX Flows
=====
Slot Number : 4                               Total Number Of Flows : 2
```

```

Port/      SrcIP/DstIP      Src/      Protocol/      DSCP/      Egress      Start/Last
Vlan      Addr              Dst/      Obsv           TcpFlag    Port/       Time
          Port              Port
-----
4/7       16.16.16.1       0         icmp           0          0          Oct 26 14:05:04
7         15.15.15.1       0         Port           none
          15.15.15.1       0         udp            0          0          Oct 26 14:05:04
4/5       15.15.15.1       0         Port           none
5         16.16.16.1       0         Port

```

Total number of Displayed Flows on Slot 4 : 2

```

Port/      SrcMac/DstMac      Byte/Pkt
Vlan                               Count
-----
4/7       00:00:00:00:00:16  953306808
7         00:24:7f:9c:6a:01  951404
4/5       00:00:00:00:00:15  1906423636
5         00:24:7f:9c:6a:00  1902568

```

Total number of Displayed Flows on Slot 4 : 2

```
VSP-9012:1#show ip ipfix flows 3-12 numflows 0
```

```

=====
IPFIX Flows
=====
Slot Number : 4                Total Number Of Flows : 2
Slot Number : 6                Total Number Of Flows : 6
-----
Slot Total: 2                Total Number of Flows on All Selected Slots: 8

```

Variable definitions

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

Variable	Value
byte-count <i>WORD</i> <1-2> <0-4294967295>	Shows the flows that match a number of bytes. Use the format oper{= < = >} and byte-count {0-4294967295}; for example, {>=a}.
dest-addr <i>WORD</i> <1-2> {A.B.C.D}	Shows the flows for a destination address. Use the format oper{= < = >} and ip address {A.B.C.D}; for example, {<=A.B.C.D}.
first-pkt-time <i>WORD</i> <1-2> <MMddyyyyhhmmss>	Shows the flows that match a timestamp for when the flow was first observed. Use the format oper{= < = >} and time {MMddyyyyhhmmss}; for example, {>=a}.
last-pkt-time <i>WORD</i> <1-2> <MMddyyyyhhmmss>	Shows the flows that match a timestamp for when the flow was last observed. Use the format oper{= < = >} and time {MMddyyyyhhmmss}; for example, {>=a}.
monitor <false true>	Monitors the top 10 flows (by byte count) if you configure this variable to true. The maximum number of flows you can monitor is 100.

Variable	Value
numflows <0–16000>	Shows the number of flows you specify. Specify zero (0) to show a flow summary. If you enter 0, the command output contains two extra lines at the bottom. The first line is all dashes and the second line is the total number of flows based on the slots you specify. The <code>show ip ipfix flows numflows 0</code> command displays the number of IPFIX flows in the IPFIX flow table of the RSP.
pkt-count WORD<1-2> <0-4294967295>	Shows the flows that match a packet count. Use the format <code>oper{= != <= >=}</code> and <code>pkt-count {0–4294967295}</code> ; for example, <code>{>=a}</code> .
port WORD<1-2> {slot/port}	Shows the flows for a particular port. Use the format <code>oper{= != <= >=}</code> and <code>{slot/port}</code> ; for example, <code>{=a/b}</code> .
protocol WORD<1-2> <0–255>	Shows the flows for a particular protocol. Use the format <code>oper{= != <= >=}</code> and <code>protocol {0–255}</code> ; for example, <code>{>=a}</code> . The mapping values for some protocol types are: icmp:1, tcp:6, udp:17, ipsecesp:50, ipsecah:51, ospf:89, vrrp:112, snmp:254, undefined:256.
source-addr WORD<1-2> {A.B.C.D}	Shows the flows for a source address. Use the format <code>oper{= != <= >=}</code> and ip address {A.B.C.D}; for example, <code>{<=A.B.C.D}</code> .
TCP-UDP-dest-port WORD<1-2> <0-65535>	Shows the flows for a destination port. Use the format <code>oper{= != <= >=}</code> and <code>port {0–65535}</code> ; for example, <code>{>=a}</code> .
TCP-UDP-src-port WORD<1-2> <0-65535>	Shows the flows for a source port. Use the format <code>oper{= != <= >=}</code> and <code>port {0–65535}</code> ; for example, <code>{>=a}</code> .
TOS WORD<1-2> <0–255>	Shows the flows that match a type of service. Use the format <code>oper{= != <= >=}</code> and <code>TOS{0-255}</code> ; for example, <code>{>=a}</code>
vlan WORD<1-2> <1–4084>	Shows the flows for a particular VLAN. Use the format <code>oper{= != <= >=}</code> and <code>vlan{1–4084}</code> ; for example, <code>{!=10}</code> .

Flushing IPFIX flow information

Flush IPFIX flow information to delete all records that correspond to the port number you specify.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Flush the exporter database:

```
ip ipfix flush port {slot/port[-slot/port][,...]} [export-and-flush]
```

Example

```
VSP-9012:1>ip ipfix flush port 3/1
```

```
VSP-9012:1>ip ipfix flush port 4/1-4/24 export-and-flush
```

Variable definitions

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

Variable	Value
export-and-flush	Optionally, initiates an export of all records, and then deletes the database after the export finishes. in UDP-based transport, the exporter sends out the flow database once, but there is no guarantee that the export reaches the collector. In TCP/SCTP-based transport, the receipt of the export by the collector is guaranteed.
{slot/port[-slot/port][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).

Viewing global IPFIX information

View global IPFIX information to see the global administrative state of IPFIX for the chassis.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View the global information:

```
show ip ipfix
```

Example

```
VSP-9012:1#show ip ipfix
```

```
=====
                                IPFIX Global
=====
                                Global-State : enable
```

Viewing collector information

View collector information to verify the collector configuration.

Procedure

1. Enter Privileged EXEC mode:
enable
2. View the collector information:
show ip ipfix collector [{slot[-slot][, ...]]}

Example

```
VSP-9012:1#show ip ipfix collector
=====
IPFIX Collector-Info
=====
SlotNum      Collector      Enable      Protocol      Dest-Port      Exporter
  IP-Address      State
-----
4             47.17.143.146  true        udp            9995            47.17.159.20
```

Variable definitions

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

Variable	Value
{slot[-slot][, ...]}	Identifies the slot in one of the following formats: a single slot (3), a range of slots (3-4), or a series of slots (3,5,6).

Viewing exporter information

View the exporter configuration to show the following information:

- the administrative state of the exporter
- the template refresh rate
- the export interval
- the aging time
- the active timeout value

Procedure

1. Enter Privileged EXEC mode:
enable
2. View the exporter configuration:
show ip ipfix exporter [{slot[-slot][, ...]]}

Example

```
VSP-9012:1#show ip ipfix exporter
=====
IPFIX Exporter-Info
=====
```

SlotNum	Admin State	Template Refresh-Rate (in sec)	Template Refresh-Rate (# of pkts)	Export Period (in sec)	Aging Period (in sec)	Active Timeout (in mins)
4	enable	60	10000	10	15	2

Variable definitions

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

Variable	Value
{slot[-slot][,...]}	Identifies the slot in one of the following formats: a single slot (3), a range of slots (3-4), or a series of slots (3,5,6).

Viewing IPFIX information for an interface

View IPFIX information for an interface to see the sampling rate and the IPFIX administrative status for the interface.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View interface information:

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

Example

```
VSP-9012:1#show ip ipfix interface gigabitethernet 4/1-4/3
```

```
=====
IPFIX Interface
=====
Port   Sampling   IPFIX
Num    Rate       State
-----
4/1    1          disable
4/2    1          enable
4/3    1          disable
```

Variable definitions

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

Variable	Value
{slot[-slot][,...]}	Identifies the slot in one of the following formats: a single slot (3), a range of slots (3-4), or a series of slots (3,5,6).

IPFIX configuration using EDM

Use IPFIX on the Virtual Services Platform 9000 to monitor IP flows.

You can use the procedures in this section to perform the following actions:

- Configure filters to match a set of flows and perform ingress metering on those flows.
- Configure IPFIX directly on a port.
- Configure a sampling rate to prevent continuous monitoring.
- Export the flow information to a third party collector as in the Netflow V9 format, which uses UDP as transport.
- Maintain the following statistics for each flow:
 - aggregate byte count
 - aggregate packet count
 - time stamp for the start of the flow
 - time stamp for the last time this flow was observed

Enabling IPFIX globally

You must globally enable IPFIX before the software will perform IPFIX functions in filters or on a port.

Note:

Virtual Services Platform 9000 does not support IPFIX on IS-IS interfaces.

About this task

After you globally disable IPFIX, even with the exporter enabled or filters configured with IPFIX enabled, the software does not perform IPFIX functions. If IPFIX is enabled on the system and you use a global disable command, a warning message appears before the system disables IPFIX.

The default global state is disabled.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > Serviceability**.
2. Click **IPFIX**.
3. Click the **Global** tab.
4. In the State box, select **enable**.
5. Click **Apply**.

Configuring collector parameters

Configure collector parameters to determine to which collector an interface module exports flow information. You can configure up to two collectors for each interface module.

About this task

Specify an exporter IP address to configure the source address in the IPFIX packets the interface module sends to the collectors. If you do not specify an exporter IP address, the source IP address is chosen from virtual IP, management IP, or outgoing interface IP based on the collector IP reachability.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > Serviceability**.
2. Click **IPFIX**.
3. Click the **Collectors/Slots** tab.
4. Click **Insert**.
5. Select the slot.
6. Specify the IP address for the collector.
7. Specify the destination port and exporter IP address.
8. Click **Insert**.

Collector/Slots field descriptions

Use the data in the following table to use the **Collector/Slots** tab.

Name	Description
SlotNum	Identifies the slot number. This value provides an index value for the collector entry.
AddressType	Specifies the address type for the collector. Virtual Services Platform 9000 currently supports IPv4.
Address	Specifies the IP address of the collector.
Protocol	Specifies the protocol for export data from the exporter to the collector. Virtual Services Platform 9000 currently supports UDP.
DestPort	Specifies the destination port to which to send flow information.
ExporterIpType	Specifies the address type for the exporter. Virtual Services Platform 9000 currently supports IPv4.
ExporterIp	Specifies the IP address to use as the source IP in the flow data.
Enable	Enables or disables the collector. The default state is enabled (selected).

Configuring slot parameters

Configure an exporter slot to send data from an exporter to a collector.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > Serviceability**.
2. Click **IPFIX**.
3. Click the **Exporters/Slots** tab.

You can edit all fields except the slot number.

4. Double-click a value to change the configuration.
5. Click **Apply**.

Exporters/Slots field descriptions

Use the data in the following table to use the **Exporters/Slots** tab.

Name	Description
SlotNum	Identifies the slot number.
AgingIntv	Configures the flow record aging interval. The default is 15 seconds.
ActiveTimeout	Configures the flow active timeout. The default is 30 minutes.
ExportIntv	Configure the interval at which to export flow information. The default is 50 seconds.
ExportState	Enables the exporter state for the slot. The default is enable.
TempRefIntvSec	Configures the template refresh timeout. The template is sent to the collector every x seconds or after x exported packets, whichever occurs first. The default is 60 seconds.
TempRefIntvPkts	Configures the template refresh timeout. The template is sent to the collector every x seconds or after x exported packets, whichever occurs first. The default is 20 packets.

Configuring IPFIX on a port

Configure IPFIX on a port to meter IP flows on the port.

Note:

Virtual Services Platform 9000 does not support IPFIX on IS-IS interfaces.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > Serviceability**.
2. Click **IPFIX**.
3. Click the **Ports** tab.

You can edit all fields except Id.

4. Double-click a field to change the configuration.
5. Click **Apply**.

Ports field descriptions

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

Name	Description
Id	Identifies the slot and port.
SampleRate	Configures the sampling rate for metering on the port, as one in every n packets. The default value is 1, which configures continuous monitoring.
Flush	Deletes all records stored in the COP or exports all records and deletes the database after the export finishes. The default is none.
AllTraffic	Enables or disables IPFIX on all traffic for the specified port. The default is disabled.

Configuring IPFIX metering using filters

Configure IPFIX metering using filters to use IPFIX on selected flows. An ACL can have multiple ACEs and each ACE can have an action of IPFIX enable. The default state is disable.

Before you begin

- The ACL and ACE exist. For more information about how to configure ACLs and ACEs, see *Virtual Services Platform 9000 Configuration — QoS and ACL-Based Filtering*, NN46250–502.
- IPFIX is enabled globally.

About this task

A packet can match multiple ACEs in an ACL. The system performs the regular actions you configure in the filter. If multiple ACEs have an action of ipfix-enable, the system performs metering only once for a packet. A packet matches multiple ACEs because you configure the ACEs to match overlapping flows. IPFIX metering further categorizes this packet into a flow record based on the unique ipfix-handle. The packet matches *n* ACEs that correspond to *n* different ACE flows, but it is still a single IPFIX flow.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > Security > Data Path**.

2. Click **Advanced Filters (ACE/ACLs)**.
3. Select an ACL.
4. Click **ACE**.
5. Select an ACE.
6. Click **Action**.
7. For **IpfixState**, select **enable**.
8. Click **Apply**.

Chapter 6: Key Health Indicators

The Key Health Indicators (KHI) feature for Avaya Virtual Services Platform 9000 provides a subset of health information that allows for quick assessment of the overall operational state of the device.

Avaya recommends that you capture a snapshot of KHI performance data when the system is running in a normal state of operation to use as a benchmark of a healthy system. You can then compare this snapshot against a new performance snapshot to help troubleshoot fault situations.

Tip:

Take new benchmark snapshots whenever you introduce a new service or alter the network topology.

KHI identifies key information that could lead Avaya support personnel towards discovery of a specific failure. The KHI feature is not intended to provide a comprehensive debugging solution.

Use the procedures in this chapter to view Key Health Indicators (KHI) information and capture KHI performance data snapshots for the Virtual Services Platform 9000 using Enterprise Device Manager (EDM) and Avaya command line interface (ACLI).

Key Health Indicators using ACLI

The Key Health Indicators (KHI) feature of Avaya Virtual Services Platform 9000 provides a subset of health information that allows for quick assessment of the overall operational state of the device.

Note:

The KHI feature is not intended to provide a comprehensive debugging solution. Instead, KHI identifies key information that could lead Avaya support personnel towards discovery of a specific failure. After the technician assesses the KHI information, further debugging is required to determine the specific reason for the fault.

Avaya recommends that you capture KHI information during normal operations to provide a baseline for Avaya support personnel when detecting fault situations.

Displaying KHI information

Use the commands detailed in this section to show KHI information.

About this task

All commands use a slot number as an optional argument. Specifying the slot number limits command output to that slot. Leaving the slot number out of the command displays KHI information for all applicable slots. You can issue the KHI commands from any command mode.

Note:

- The show khi forwarding commands are only valid for I/O slots.
- The caret (^) symbol is used to denote the peak value in the output for the **show khi forwarding zagros** command. An entry that reads 5^ 0 denotes a peak value of 5 and a current value of 0.

Procedure

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

2. Display all KHI statistics:

```
show khi forwarding [<3-12>]
```

3. Display IFP statistics:

```
show khi forwarding ifp [<3-12>]
```

4. Display internal statistics:

```
show khi forwarding k2 [<3-12>]
```

5. Display MAC packet transmit, receive, and drop counts:

```
show khi forwarding mac [<3-12>]
```

Note:

This command is applicable only for first generation modules. It is not supported on second generation modules.

6. Display per lane internal datapath counters:

```
show khi forwarding mac-higig [<3-12>]
```

Note:

This command is applicable only for first generation modules. It is not supported on second generation modules.

7. Display internal QE statistics:

```
show khi forwarding qe [<3-12>]
```

The output displays differently for first generation and second generation modules.

8. Display internal RSP counters:

```
show khi forwarding rsp [<3-12>]
```

The output displays differently for first generation and second generation modules.

9. Display Sierra data path health indicators:

```
show khi forwarding sierra [<3-12>]
```

Note:

This command is only applicable for 9048XS-2 and 9012QQ-2, second generation modules. It is not supported on first generation modules.

10. Display internal Zagros counters:

```
show khi forwarding zagros [<3-12>]
```

11. Run all KHI show commands and capture the output to the file named in the file parameter:

```
show fulltech khi [file WORD<1-99>]
```

Example

The following output displays for first generation modules for the **show khi forwarding ifp [<3-12>]** command.

```
VSP-switch:1#show khi forwarding ifp
=====
Forwarding KHI Details - IFP Statistics - Slot 3
=====
IFP PCV Counters
-----
RuleNo RuleName                Ports 1-8      Ports 9-16     Ports 1
7-24
-----
3      Datapath HB HG0              20440          20441          20440
4      Datapath HB HG1              20441          20441          20440
-----
IFP PCV Counters - MAC-in-MAC Tagged
-----
RuleNo RuleName                Ports 1-8      Ports 9-16     Ports 1
7-24
-----
IFP PCV Counters - MAC-in-MAC Untagged
-----
RuleNo RuleName                Ports 1-8      Ports 9-16     Ports 1
7-24
-----
--More-- (q = quit)
```

The following output displays for second generation modules for the **show khi forwarding ifp [<3-12>]** command.

```
VSP-switch:1#show khi forwarding ifp 9
=====
Forwarding KHI Details - IFP Drop Statistics - Slot 9
=====
```

Key Health Indicators

Counter	Ports 1-8	Ports 9-16
IFP_DROP	6	0
MAC_LOOPBACK_DROP	6	0

Counter	Ports 17-24	Ports 25-32
IFP_DROP	9	6
VLAN_MEM_DROP	25	0
MAC_LOOPBACK_DROP	2	0

Counter	Ports 33-40	Ports 41-48
---------	-------------	-------------

=====

Forwarding KHI Details - IFP Packet Type Statistics - Slot 9

=====

RuleNo	RuleName	Ports 1-8	Ports 9-16
0		9	0
12		2	0
13		14	0
25		121	0
30		6	0
35		27	0

RuleNo	RuleName	Ports 17-24	Ports 25-32
12		1	0
13		14	0
25		80	0
35		15	0

RuleNo	RuleName	Ports 33-40	Ports 41-48
--------	----------	-------------	-------------

The following output displays for the **show khi forwarding k2 [<3-12>]** command.

```
VSP-switch:1#show khi forwarding k2 4
```

=====

Forwarding KHI Details - K2 Statistics - Slot 4

=====

Health Indicator	Ports 1-12	Ports 25-36
MAC->K2 If 0	38944	24959
K2 If 0->Zagros	3356144	3342153
Zagros->K2 If 0	17538	21036
K2 If 0->MAC	13289139	15798
MAC->K2 If 0 Err	0	0
K2 If 0->Zagros Err	0	0
Zagros->K2 If 0->Err	0	0
K2 If 0->MACErr	0	0
IP Multicast Drops		
Ports 1 and 25	0	0
IP Multicast Drops		
Ports 2 and 26	0	0
IP Multicast Drops		
Ports 3 and 27	0	0
IP Multicast Drops		

Ports 4 and 28	0	0
IP Multicast Drops		
Ports 5 and 29	0	0
IP Multicast Drops		
Ports 6 and 30	0	0
IP Multicast Drops		
Ports 7 and 31	0	0
IP Multicast Drops		
Ports 8 and 32	0	0
IP Multicast Drops		
Ports 9 and 33	0	0
IP Multicast Drops		
Ports 10 and 34	0	0
IP Multicast Drops		
Ports 11 and 35	0	0
IP Multicast Drops		
Ports 12 and 36	0	0

Health Indicator	Ports 13-24	Ports 37-48

MAC->K2 If 1	13158276	32160
K2 If 1->Zagros	16496150	3370027
Zagros->K2 If 1	13372000	15900
K2 If 1->MAC	13372000	15900
MAC->K2 If 1 Err	0	0
K2 If 1->Zagros Err	0	0
Zagros->K2 If 1->Err	0	0
K2 If 1->MACErr	0	0
IP Multicast Drops		
Ports 13 and 37	0	0
IP Multicast Drops		
Ports 14 and 38	0	0
IP Multicast Drops		
Ports 15 and 39	0	0
IP Multicast Drops		
Ports 16 and 40	0	0
IP Multicast Drops		
Ports 17 and 41	0	0
IP Multicast Drops		
Ports 18 and 42	0	0
IP Multicast Drops		
Ports 19 and 43	0	0
IP Multicast Drops		
Ports 20 and 44	0	0
IP Multicast Drops		
Ports 21 and 45	0	0
IP Multicast Drops		
Ports 22 and 46	0	0
IP Multicast Drops		
Ports 23 and 47	0	0
IP Multicast Drops		
Ports 24 and 48	0	0

The following output displays for the **show khi forwarding mac [<3-12>]** command.

```
VSP-switch:1#show khi forwarding mac 4
```

Forwarding KHI Details - MAC Statistics - Slot 4				
Ports	Rx OK	Tx OK	Rx Err	Tx Err
4/1	7568	3972	0	0

Key Health Indicators

4/2	0	0	0	0
4/3	0	0	0	0
4/4	0	0	0	0
4/5	0	0	0	0
4/6	4825649	4642006	0	0
4/7	4269079	4220186	0	0
4/8	4260696	4684188	0	0
4/9	0	0	0	0
4/10	0	0	0	0
4/11	0	0	0	0
4/12	0	7137	0	0
4/13	7568	1774	0	0
4/14	0	0	0	0
4/15	0	0	0	0
4/16	0	0	0	0
4/17	0	0	0	0
4/18	0	0	0	0
4/19	0	0	0	0
4/20	7147	1779	0	0
4/21	0	0	0	0
4/22	0	0	0	0
4/23	7145	1783	0	0
4/24	7145	1783	0	0
4/25	7583	1774	0	0
4/26	0	0	0	0
4/27	7145	1783	0	0
4/28	7146	1784	0	0
4/29	0	0	0	0
4/30	0	0	0	0
4/31	0	0	0	0
4/32	0	0	0	0
4/33	0	0	0	0
4/34	0	0	0	0

4/35	0	0	0	0
4/36	0	0	0	0
4/37	7583	1774	0	0
4/38	0	0	0	0
4/39	0	0	0	0
4/40	0	0	0	0
4/41	0	0	0	0
4/42	0	0	0	0
4/43	0	0	0	0
4/44	0	0	0	0
4/45	0	7166	0	0
4/46	0	0	0	0
4/47	8238	1779	0	0
4/48	0	0	0	0

Ports	RDBGCO	RDBGCO MASK	TDBGCO	TDBGCO M
ASK				

4/1	0	0x0	0	0x0
4/2	0	0x0	0	0x0
4/3	0	0x0	0	0x0
4/4	0	0x0	0	0x0
4/5	0	0x0	0	0x0
4/6	40	0xc1	0	0x0
4/7	33	0xc1	0	0x0
4/8	19225	0xd1	0	0x0
4/9	0	0x0	0	0x0
4/10	0	0x0	0	0x0
4/11	0	0x0	0	0x0
4/12	0	0x0	0	0x0
4/13	0	0x0	0	0x0
4/14	0	0x0	0	0x0
4/15	0	0x0	0	0x0
4/16	0	0x0	0	0x0

Key Health Indicators

4/17	0	0x0	0	0x0
4/18	0	0x0	0	0x0
4/19	0	0x0	0	0x0
4/20	0	0x0	0	0x0
4/21	0	0x0	0	0x0
4/22	0	0x0	0	0x0
4/23	0	0x0	0	0x0
4/24	0	0x0	0	0x0
4/25	0	0x0	0	0x0
4/26	0	0x0	0	0x0
4/27	0	0x0	0	0x0
4/28	0	0x0	0	0x0
4/29	0	0x0	0	0x0
4/30	0	0x0	0	0x0
4/31	0	0x0	0	0x0
4/32	0	0x0	0	0x0
4/33	0	0x0	0	0x0
4/34	0	0x0	0	0x0
4/35	0	0x0	0	0x0
4/36	0	0x0	0	0x0
4/37	0	0x0	0	0x0
4/38	0	0x0	0	0x0
4/39	0	0x0	0	0x0
4/40	0	0x0	0	0x0
4/41	0	0x0	0	0x0
4/42	0	0x0	0	0x0
4/43	0	0x0	0	0x0
4/44	0	0x0	0	0x0
4/45	0	0x0	0	0x0
4/46	0	0x0	0	0x0
4/47	1112	0x51	0	0x0
4/48	0	0x0	0	0x0

The following output displays for the **show khi forwarding mac-higig [<3-12>]** command.

```
VSP-switch:1#show khi forwarding mac-higig 4
=====
Forwarding KHI Details - MAC HIGIG Statistics - Slot 4
=====
```

Health Indicator	Ports 1-12	Ports 25-36
MAC->K2 If 0	360228351	346478657
K2 If 0->MAC	13890172	11077
RDGBC0	3	0
RDGBC0 Mask	65	0
TDGBC0	19396	0
TDGBC0 Mask	69	0

Health Indicator	Ports 13-24	Ports 37-48
MAC->K2 If 1	346583396	346471315
K2 If 1->MAC	11082	16558
RDGBC0	0	0
RDGBC0 Mask	0	0
TDGBC0	0	0
TDGBC0 Mask	0	0

Health Indicator	Ports 1-24	Ports 25-48
IFP DOS Drops	0	0

The following output displays for the **show khi forwarding qe [<3-12>]** command.

```
VSP-switch:1#show khi forwarding qe 3
=====
Forwarding KHI Details - QE Statistics - Slot 3
=====
```

Health Indicator	Ports 1-8	Ports 9-16	Ports 17-24
Ingress qm_agr_accepted_pkt_cnt0	19424578	19424596	19424419
Ingress qm_agr_accepted_pkt_cnt1	186	193	188
Ingress qm_agr_dequeued_pkt_cnt	19424764	19424789	19424607
Ingress pm_switch_pkt_cnt	19424764	19424789	19424607
Ingress sr0_rx_p0_pkt_cnt	9712384	9712396	9712306
Ingress srl_rx_p0_pkt_cnt	9712386	9712399	9712306
Egress sv_pkt_cnt	212629146	212505906	212505937
Egress st0_p0_tx_pkt_cnt	641	0	0
Egress st0_p14_tx_pkt_cnt	122632	0	0
Egress st0_p15_tx_pkt_cnt	212484491	212484524	212484556
Egress st1_p15_tx_pkt_cnt	212484492	212484524	212484556

Key Health Indicators

Qm_agr_non_wred_dropped_pkt_cnt 6 6 5

The following output displays for first generation modules for the **show khi forwarding rsp** [**<3-12>**] command.

```
VSP-switch:1#show khi forwarding rsp 4
=====
Forwarding KHI Details - RSP Statistics - Slot 4
=====
```

Health Indicator	Ports 1-12	Ports 25-36
LSM Drops	27886538	27886546
Exception Drops	0	0
Frame Error Drops	0	0
FDIB full drops	0	0
Ingr MLT	0	0
All Ports Down		
Egress mlt all	0	0
Ports Down Drops		
Egress IP Mcast	0	0
Records not found		
Egress IP Mcast	0	0
MLT Wrong Port		
Egress IP Mcast	0	0
Source Knockout		
Ingress DA not	0	0
Found Drops		
Ingress Unknown	0	0
Ingress Discard	0	0
Dest Id Drops		
MAC Learning	0	0
Packet Drops		
Ingress IPMC	0	0
Supression Drops		
Unsupported Feature	0	0
Drops		
ACL Discards	0	0
Ingress IPMC	0	0
Lookup Fails		
IPV4 Dest IP	0	0
Lookup Fails		
IPV4 Source IP	0	0
Lookup Fails		
L3Mirror Drops	0	0

Health Indicator	Ports 13-24	Ports 37-48
LSM Drops	27886543	27886553
Exception Drops	0	0
Frame Error Drops	0	0
FDIB full drops	0	0
Ingr MLT	0	0
All Ports Down		
Egress mlt all	0	0
Ports Down Drops		
Egress IP Mcast	0	0
Records not found		
Egress IP Mcast	0	0
MLT Wrong Port		
Egress IP Mcast	0	0
Source Knockout		
Ingress DA not	0	0

```

Found Drops
Ingress Unknown          0          0
Ingress Discard          0          0
Dest Id Drops
MAC Learning              0          0
Packet Drops
Ingress IPMC              0          0
Supression Drops
Unsupported Feature      0          0
Drops
ACL Discards              0          0
Ingress IPMC              0          0
Lookup Fails
IPV4 Dest IP              0          0
Lookup Fails
IPV4 Source IP           0          0
Lookup Fails
L3Mirror Drops            0          0

```

The following output displays for second generation modules for the **show khi forwarding rsp [<3-12>]** command.

```
VSP-switch:1#show khi forwarding rsp 9
```

```

=====
Forwarding KHI Details - RSP Statistics - Slot 9
=====
-----
RSPng Statistics
-----
Name                               Ports 1-16
-----
HAB0_RX_PACKET_COUNTER_REG         9340956
HAB0_TX_PACKET_COUNTER_REG         9341254
HAB1_RX_PACKET_COUNTER_REG         9342612
-----
Name                               Ports 17-32
-----
HAB0_RX_PACKET_COUNTER_REG         9342563
HAB0_TX_PACKET_COUNTER_REG         9342861
-----
Name                               Ports 33-48
-----
HAB0_RX_PACKET_COUNTER_REG         9341285
-----
RSPng Drop Statistics
-----
Name                               Ports 1-16
-----
GS_IST_NO_UNI_MEM                   154
GS_UNKNOWN_MCAST_BDA                3067
GS_IPMC_V4_DROP                     216
GS_V4_LKUP_DROP                     22
-----
Name                               Ports 17-32
-----
GS_SRC_KNOCK_OUT                    49
GS_IST_NO_UNI_MEM                   78
GS_IPMC_V4_DROP                     10974
GS_V4_LKUP_DROP                     3000
-----
Name                               Ports 33-48
-----
GS_IPMC_V4_DROP                     48
-----

```

Key Health Indicators

RSPng Status Registers	
Name	Ports 1-16
DEVICE_ID_REG	0x0D055A1A14070920
PLL_STATUS_REG	0x00000000000000FF
EDI0_STATUS_REG	0x000000000000000B
EDI1_STATUS_REG	0x000000000000000B
Name	Ports 17-32
DEVICE_ID_REG	0x0D055A1A14070920
PLL_STATUS_REG	0x00000000000000FF
Name	Ports 33-48
DEVICE_ID_REG	0x0D055A1A14070920
PLL_STATUS_REG	0x00000000000000FF
RSPng Error Registers	
Name	Ports 1-16
Name	Ports 17-32
Name	Ports 33-48

The following output displays for first generation modules for the **show khi forwarding sierra** [**<3-12>**] command.

```
VSP-switch:1(config)#show khi forwarding sierra 4
=====
Forwarding KHI Details - SIERRA Statistics - Slot 4
=====
SIERRA STATUS Statistics - Slot 4
=====
SIERRA_DEV_ID          0x0D052A1B
SIERRA_BUILD           0x14012410
SIERRA_BAD_DATA        0xDEADDEAD
SIERRA_SB_REX_EVENT    0x00000000
SIERRA_REX_STATUS      0x0F0F0F0F
SIERRA_SB_STATUS       0x00000000
SIERRA_SLICEX_WINDOW_EVENT 0x00000000
SIERRA_PCIE_CORE_STATUS 0x5F600007
SIERRA_PCIE_EVENT      0x00000000
SIERRA_PSI1_STATUS     0x00030433
SIERRA_PSI0_STATUS     0x00030433
SIERRA_PSI_EVENT       0x00020002
SIERRA_PSI1_LINK_STATUS 0x00000000
SIERRA_PSI0_LINK_STATUS 0x00008200
SIERRA_GE_STATUS       0x00000003
--More-- (q = quit)
```

The following output displays for second generation modules for the **show khi forwarding sierra** [**<3-12>**] command.

```
VSP65:1(config)#show khi forwarding sierra 10
```

```
=====
Forwarding KHI Details - SIERRA Statistics - Slot 10
=====

Sierra Status Registers - Slot 10
=====
SIERRA_DEV_ID                0x0D052A1B
SIERRA_BUILD                 0x14070813
SIERRA_BAD_DATA              0xDEADDEAD
SIERRA_SB_REX_EVENT          0x00000000
SIERRA_REX_STATUS            0x0B0B0F0F
SIERRA_SB_STATUS             0x00000000
SIERRA_SLICEX_WINDOW_EVENT   0x00000000
SIERRA_PCIE_CORE_STATUS      0xDF600007
SIERRA_PCIE_EVENT            0x00000000
SIERRA_PSII_STATUS           0x00030433
SIERRA_PSIO_STATUS           0x00030433
SIERRA_PSI_EVENT             0x00020002
SIERRA_PSII_LINK_STATUS      0x00000104
SIERRA_PSIO_LINK_STATUS      0xC108C311
SIERRA_GE_STATUS             0x00000003

SIERRA_GE_EVENT              0x00000003
SIERRA_PKT_EVENT             0x00000000
SIERRA_ZIP_STATUS0           0x001C37FF
SIERRA_ZIP_STATUS1           0x001C37FF
SIERRA_ZIP_STATUS2           0x004C27FF
SIERRA_ZIP_STATUS3           0x001C37FF
SIERRA_ZIP_STATUS4           0x001C37FF
SIERRA_ZIP_STATUS5           0x001C27FF
SIERRA_ZIP_EVENT0            0x00000000
SIERRA_ZIP_EVENT1            0x00000000
SIERRA_ZIP_EVENT2            0x00000000
SIERRA_ZIP_EVENT3            0x00000000
SIERRA_ZIP_EVENT4            0x00000000
SIERRA_ZIP_EVENT5            0x00000000
SIERRA_SBM_STATUS0           0x000F0007
SIERRA_SBM_STATUS1           0x000F0007
SIERRA_SBM_STATUS2           0x000F0007
SIERRA_SBM_STATUS3           0x000F0007
SIERRA_SBM_STATUS4           0x000F0007
SIERRA_SBM_STATUS5           0x000F0007
SIERRA_SBM_EVENT0            0x00000000
SIERRA_SBM_EVENT1            0x00000000
SIERRA_SBM_EVENT2            0x00000000

SIERRA_SBM_EVENT3            0x00000000
SIERRA_SBM_EVENT4            0x00000000
SIERRA_SBM_EVENT5            0x00000000
SIERRA_SEP_EVENT             0x00000000

=====
Sierra Forwarding Statistics - Slot 10
=====
SIERRA_GE_CNT_RX_OK          1696
SIERRA_GE_CNT_TX_OK          1682
SIERRA_GE_CNT_RX_PAUSE       0
SIERRA_PKT_CNT_GE_GE         0
```

Key Health Indicators

```

SIERRA_PKT_CNT_GE_Z0          292
SIERRA_PKT_CNT_GE_Z1          280
SIERRA_PKT_CNT_GE_Z2          282
SIERRA_PKT_CNT_GE_Z3          280
SIERRA_PKT_CNT_GE_Z4          281
SIERRA_PKT_CNT_GE_Z5          281
SIERRA_PKT_CNT_Z0_GE_OK      1682
SIERRA_PKT_CNT_Z1_GE_OK        0
SIERRA_PKT_CNT_Z2_GE_OK        0
SIERRA_PKT_CNT_Z3_GE_OK        0
SIERRA_PKT_CNT_Z4_GE_OK        0

```

--More-- (q = quit)

The following output displays for first generation modules for the **show khi forwarding zagros** [**<3-12>**] command.

```
VSP-switch:1(config)#show khi forwarding zagros 3
```

```

=====
Forwarding KHI Details - Zagros Statistics - Slot 3
=====
Health Indicator                Ports 1-12                Ports 25-36
-----
K2 If 1->Zagros                 519353                    519353
Zagros->RSP                     27841744                  27844449
Zagros->QE If 1                 1035031                   1035034
QE If 1->Zagros                 26815334                  26812615
Zagros->K2 If 1                 1631                      2785
ZAP Tx Ctl                      7622240                   7628830
ZAP Tx Data                     6570                      0
ZAP Rx Ctl                      7622240                   7628830
ZAP Tx HBE                      6359901                   6362598
Egress Esb6Count                0                          1155
PMM output Drop count          26807126                  26809827
PMM Admission RSP
Drop Count                      26807126                  26809827
PMM RSP rx count               27841744                  27844448
PMM RSP tx count               1036662                   1037819
PMM HAB bus rx                 27841744                  27844449

PMM CIF request count          27838478                  27841183
PMM CIF response count        27838479                  27841183
PMM RSP PLC Threshold          170                       170
PMM RE PLC Threshold           48                         48
PMM Free Page Count (OPA)      17                         17
PMM Free Page Count (RPA)      17                         17
PMM Free Page Count (FPM)      2009                       2009
PMM RSP PLC Packet Count       5^ 0                       5^ 0
PMM Egress OOB                 5^ 0                       5^ 0
PMM Ingress Heartbeat          2^ 0                       2^ 0
PMM Ingress COP Insertion      1^ 0                       1^ 0
PMM Egress L2BC/UC             0^ 0                       2^ 0
PMM number of pools            9                           9

-----
Health Indicator                Ports 13-24                Ports 37-48
-----
K2 If 1->Zagros                 524327                    523824
Zagros->RSP                     27846969                  27850388
Zagros->QE If 1                 1040007                   1039508
QE If 1->Zagros                 26847401                  26830807
Zagros->K2 If 1                 30068                     18992
Zagros EHP All Port down

```

IST counter	9952	517
ZAP Tx Ctl	7641293	7628520
ZAP Rx Ctl	7641293	7628520
ZAP Tx HBE	6360149	6364066
Egress Esb3Count	825	0
Egress Esb4Count	825	0
Egress Esb5Count	26783	0
PMM output Drop count	26807377	26811292
PMM Admission RSP		
Drop Count	26807375	26811292
PMM RSP rx count	27846969	27850388
PMM RSP tx count	1070075	1058500
PMM HAB bus rx	27846969	27850388
PMM CIF request count	27838731	27842654
PMM CIF response count	27838731	27842654
PMM RSP PLC Threshold	170	170
PMM RE PLC Threshold	48	48
PMM Free Page Count (OPA)	17	17
PMM Free Page Count (RPA)	18	16
PMM Free Page Count (FPM)	2008	2010
PMM RSP PLC Packet Count	5^ 0	5^ 0
PMM Egress OOB	5^ 0	5^ 0
PMM Ingress Heartbeat	2^ 0	2^ 0
PMM Ingress COP Insertion	2^ 0	2^ 0
PMM Egress L2BC/UC	2^ 0	1^ 0
PMM number of pools	9	9

The following output displays for second generation modules for the **show khi forwarding zagros [<3-12>]** command.

```
VSP-9012:1(config)#show khi forwarding zagros 4
```

```
=====
Forwarding KHI Details - Zagros Statistics - Slot 4
=====
-----
Forwarding Statistics
-----
--
Name                               Ports 1-8                Ports 9-16
-----
--
PMM_DP_RX_COUNT                     36570605                 36570614
PMM_RSP_RX_COUNT                    255643119                255643137
PMM_RSP_TX_COUNT                    255643119                255643128
PMM_FP0_RX_COUNT                    219130934                219130934
PMM_FP0_TX_COUNT                    36629025                 36629025
PMM_FLOP_TX_COUNT                   0                         9
PMM_INGRESS_LSM_ADMIT_COUNT         36512185                 36512185
PMM_INGRESS_COP_ADMIT_COUNT         58420                     58438
PMM_EGRESS_OOB_ADMIT_COUNT          219130934                219130934
PMM_EGRESS_EXPAND_ADMIT_COUNT       58420                     58420
FP_RX_PACKET_COUNT_REGO             219130934                219481448
FP_TX_PACKET_COUNT_REGO             36629025                 36629025
HAB_RX_PACKET_COUNT                 255643119                255643137
HAB_TX_PACKET_COUNT                 255643119                255643128
--More-- (q = quit)
```

Variable definitions

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

Variable	Value
<3–12>	Specifies the slot number.
file <i>WORD</i> <1-99>	Specifies the filename and location, from 1–99 characters, in one of the following formats: <ul style="list-style-type: none"> • /intflash/<file> • /extflash/<file> • /usb/<file>

Clearing KHI information

You can clear KHI information for a specific slot or across the whole device. Use the command to clear the statistics.

About this task

Specify a slot number to clear statistics for a specific slot or leave it absent to clear information for the whole device.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Clear forwarding statistics:

```
clear khi forwarding [slot <3-12>]
```

If you clear the forwarding statistics, the IFP rules do not appear in the show command output again until the specific traffic hits the rule again.

3. Clear CPP statistics:

```
clear khi cpp <iocop-statistics|port-statistics|protocol-drops>
```

Variable definitions

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

Variable	Value
<3–12>	Specifies the slot number.

Displaying KHI performance information

Use the following commands to display information about the performance of the Key Health Indicators feature.

About this task

KHI performance commands can be used on all slots. If you do not specify a slot, the system shows information for all slots.

Tip:

Avaya recommends that you capture a snapshot of KHI performance data when the system is running in a normal state of operation to use as a benchmark of a healthy system.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display buffer performance and utilization statistics for KHI on the specified slot or all slots:

```
show khi performance buffer-pool [{slot[-slot][, ...]}
```

Important:

If the free buffers show a lack of resources, this can indicate a system under stress.

Buffer pool is not supported on interface or Switch Fabric (SF) slots.

3. Show current utilization, 5-minute average utilization, and 5-minute high water mark with date and time of event:

```
show khi performance cpu [{slot[-slot][, ...]}
```

4. Display memory performance and utilization statistics for KHI on the specified slot or all slots:

```
show khi performance memory [{slot[-slot][, ...]}
```

5. Display process performance and utilization statistics for KHI on the specified slot or all slots:

```
show khi performance process [{slot[-slot][, ...]}
```

Note:

Use this information to identify processes that are running outside their normal operational boundaries when compared to a healthy system snapshot.

6. Display thread performance and utilization statistics for KHI on the specified slot or all slots:

```
show khi performance pthread [{slot[-slot][, ...]}
```

Note:

This process displays pthread CPU information for each process thread by slot. Identify runaway processes or processes that consume too much CPU time when compared to a healthy system snapshot.

7. Display internal memory management resource performance and utilization statistics for KHI on the specified slot or all slots:

```
show khi performance slabinfo [{slot[-slot][, ...]}
```

Note:

Slabinfo displays various system information on a slot-by-slot basis, such as queues and caches on the system. Identify systemic issues on the system when compared to a healthy system snapshot.

Example

```
VSP-9012:1#show khi performance buffer-pool 1
Slot:1
  CPP:
    UsedFBuffs: 0
    FreeFBuffs: 1600
    NoFbuff: 0

  Network stack system:
    UsedMbuf: 120
    FreeMbuf: 47730
    SocketMbuf: 15

  Network stack data:
    UsedMbuf: 4
    FreeMbuf: 5372

  Letter API message queue:
    QHigh: 0
    QNormal: 0
    FreeQEntries: 51200
```

```
VSP-9012:1#show khi performance cpu 4
Slot:4
  Current utilization: 13
  5-minute average utilization: 12
  5-minute high water mark: 16 (02/08/11 09:41:04)
```

```
VSP-9012:1#show khi performance memory 4
Slot:4
  Used: 139164 (KB)
  Free: 345276 (KB)
  Current utilization: 28 %
  5-minute average utilization: 28 %
  5-minute high water mark: 28 (02/08/11 09:41:24)
```

```
VSP-9012:1#show khi performance process 1
Slot:1
```

PID	PPID	PName	VmSize	VmLck	VmRss	VmData	VmStk	VmExe	VmLib
1	0	init	2948	0	680	344	84	496	1764
2	1	ksoftirqd/0	0	0	0	0	0	0	0
3	1	events/0	0	0	0	0	0	0	0
4	1	khelper	0	0	0	0	0	0	0
5	1	kthread	0	0	0	0	0	0	0
30	5	kblockd/0	0	0	0	0	0	0	0
55	5	pdflush	0	0	0	0	0	0	0
56	5	pdflush	0	0	0	0	0	0	0
57	1	kswapd0	0	0	0	0	0	0	0
58	5	aio/0	0	0	0	0	0	0	0
624	1	mtdblockd	0	0	0	0	0	0	0
652	1	init	2948	0	244	344	84	496	1764
653	652	rcS	2868	0	1308	184	84	792	1496
674	5	wdd	0	0	0	0	0	0	0
925	1	ntpd	3536	0	1156	496	84	584	1896
929	1	portmap	1924	0	488	160	84	80	1348

931	1	inetd	3056	0	620	344	84	496	1804
933	1	syslogd	2948	0	508	344	84	496	1764
935	1	klogd	2948	0	496	344	84	496	1764
1333	5	khubd	0	0	0	0	0	0	0
1498	1	nfsd	0	0	0	0	0	0	0
1499	1	nfsd	0	0	0	0	0	0	0
1500	1	nfsd	0	0	0	0	0	0	0
1501	1	lockd	0	0	0	0	0	0	0
1502	5	rpciod/0	0	0	0	0	0	0	0
1518	1	nfsd	0	0	0	0	0	0	0
1519	1	nfsd	0	0	0	0	0	0	0
1520	1	nfsd	0	0	0	0	0	0	0
1521	1	nfsd	0	0	0	0	0	0	0
1522	1	nfsd	0	0	0	0	0	0	0
1526	1	rpc.mountd	1996	0	668	180	84	128	1352
1534	1	sshd	4612	0	928	368	84	468	2760
1560	1	tar	0	0	0	0	0	0	0
1561	653	rc.appfs	2868	0	1352	184	84	792	1496
1683	1	ckrm_cpud/0	0	0	0	0	0	0	0
1819	5	oxide	0	0	0	0	0	0	0
1897	1561	start	2868	0	1296	184	84	792	1496
1935	1897	lifecycle	75880	0	3348	68512	84	216	4900
1983	1935	patchAgent	4276	0	1372	632	180	148	2776
1984	1935	sockserv	4752	0	1548	416	84	80	3168
1985	1935	externalcf	3932	0	956	192	84	76	2256
1986	1935	oom95	109120	0	104176	102860	84	120	4672
1987	1935	oom90	109120	0	104176	102860	84	120	4672
1988	1935	imgsync.x	41236	0	3168	34196	84	164	4888
2001	1	gzip	0	0	0	0	0	0	0
2011	1	gzip	0	0	0	0	0	0	0
2014	1935	cbbcm-main.x	52716	0	12044	36416	84	9300	4232
2015	1935	cbcp-main.x	745168	32896	604720	705096	84	29352	3184
2016	1935	coreManager.x	33088	0	2872	26892	84	136	4664
2017	1935	patchmanager.x	33280	0	3384	25876	84	148	5448
2018	1935	patchAgtIf.x	33152	0	2848	26888	84	128	4676
2019	1935	remCmdAgent.x	33144	0	2844	26888	84	124	4668
2020	1935	bg_threads	40928	0	2532	35004	84	120	4552
2021	1935	logrotate	5688	0	2200	436	84	792	3376
2504	2021	sleep	3296	0	684	356	84	496	1372

```
VSP-9012:1#show khi performance pthread 1
Slot:1
```

TID	PID	PName	CPU(%)	5MinAvg CPU(%)	5MinHiWater	CPU(%(time stamp))
1	1	init	0.0	0.0		
2	2	ksoftirqd/0	0.0	0.0		
3	3	events/0	0.0	0.0		
4	4	khelper	0.0	0.0		
5	5	kthread	0.0	0.0		
30	30	kblockd/0	0.0	0.0		
55	55	pdflush	0.0	0.0		
56	56	pdflush	0.0	0.0		
57	57	kswapd0	0.0	0.0		
58	58	aio/0	0.0	0.0		
624	624	mtdblockd	0.0	0.0		
652	652	init	0.0	0.0		
653	653	rcS	0.0	0.0		
674	674	wdd	0.0	0.0		
925	925	ntpd	0.0	0.0		
929	929	portmap	0.0	0.0		
931	931	inetd	0.0	0.0		
933	933	syslogd	0.0	0.0		
935	935	klogd	0.0	0.0		

Key Health Indicators

1333	1333	khubd	0.0	0.0
1498	1498	nfsd	0.0	0.0
1499	1499	nfsd	0.0	0.0
1500	1500	nfsd	0.0	0.0
1501	1501	lockd	0.0	0.0
1502	1502	rpciod/0	0.0	0.0
1518	1518	nfsd	0.0	0.0
1519	1519	nfsd	0.0	0.0
1520	1520	nfsd	0.0	0.0
1521	1521	nfsd	0.0	0.0
1522	1522	nfsd	0.0	0.0
1526	1526	rpc.mountd	0.0	0.0
1534	1534	sshd	0.0	0.0
1561	1561	rc.appfs	0.0	0.0
1683	1683	ckrm_cpud/0	0.0	0.0
1819	1819	oxide	0.0	0.0
1897	1897	start	0.0	0.0
1935	1935	lifecycle	0.0	0.0
1975	1935	dpmXportRxMonit	0.0	0.0
1976	1935	dpmXportTxMonit	0.0	0.0
1977	1935	ltrBulkTimerThr	0.0	0.0
1978	1935	lc_wd_exception	0.0	0.0
1979	1935	lc_hwwd_feed	0.0	0.0
1980	1935	lc_swwd_feed	0.0	0.0
1981	1935	worker_thread	0.0	0.0
1982	1935	lc_master	0.0	0.0
1983	1983	patchAgent	0.0	0.0
1984	1984	sockserv	0.0	0.0
1985	1985	externalcf	0.0	0.0
1986	1986	oom95	0.0	0.0
1987	1987	oom90	0.0	0.0
1988	1988	imgsync.x	0.0	0.0
1989	1988	dpmXportRxMonit	0.0	0.0
1990	1988	dpmXportTxMonit	0.0	0.0
1991	1988	ltrBulkTimerThr	0.0	0.0
2014	2014	cbbcm-main.x	0.0	0.0
2040	2014	tUsrRoot	0.0	0.0
2043	2014	tExcTask	0.0	0.0
2044	2014	bcmtty	0.0	0.0
2045	2014	dpmXportRxMonit	0.0	0.0
2046	2014	dpmXportTxMonit	0.0	0.0
2047	2014	ltrBulkTimerThr	0.0	0.0
2048	2014	bcmtMainTask	0.0	0.0
2050	2014	bcmtLcdTask	0.0	0.0
2051	2014	bcmDPC	0.0	0.0
2052	2014	_interrupt_thre	0.0	0.0
2063	2014	bcmCNTR.1	0.0	0.0
2070	2014	bcmLINK.1	0.0	0.0
2076	2014	bcmScoreboard.0	0.0	0.0
2077	2014	bcmCNTR.0	0.0	0.0
2079	2014	bcmLINK.0	0.0	0.0
2098	2014	bcmtSfAgentTime	0.0	0.0
2099	2014	bcmtSfAgentMsgT	0.0	0.0
2015	2015	cbcp-main.x	0.0	0.0
2053	2015	tUsrRoot	0.0	0.0
2054	2015	tExcTask	0.0	0.0
2055	2015	tExcJobTask	0.0	0.0
2056	2015	tNetTask	0.0	0.0
2057	2015	traceOutput	0.0	0.0
2058	2015	profile_cmd	0.0	0.0
2059	2015	tLoggerTask	0.0	0.0
2064	2015	tTelnetd	0.0	0.0
2065	2015	tTelnetV6d	0.0	0.0
2066	2015	tRlogind	0.0	0.0
2067	2015	tRshd	0.0	0.0

```

2068 2015 tTftpdTask      0.0  0.0
2069 2015 tFtpdTask      0.0  0.0
2071 2015 dpmXportRxMonit 0.0  0.0
2072 2015 dpmXportTxMonit 0.0  0.0
2074 2015 BootpServer    0.0  0.0
2075 2015 tSioMsgRx      0.0  0.0
2078 2015 tUsrRoot       0.0  0.0
2080 2015 ch_heartbeat_cp 0.0  0.0
2081 2015 chEvmTask      0.0  0.0
2082 2015 chFsmTask      0.0  0.0
2083 2015 chServiceTask  0.0  0.0
2085 2015 CpuHATask     0.0  0.0
2086 2015 tHAQTask      0.0  0.0
2087 2015 tSnmpTmr      0.0  0.0
2088 2015 tSnmpd        0.0  0.0
2089 2015 haTick        0.0  0.0
2090 2015 tMainTask     0.4  0.3      0.4 (02/08/11 09:48:02)
2091 2015 rtMainTask    0.0  0.0
2092 2015 tCppSend      0.0  0.0
2093 2015 tCppInterruptTa 0.1  0.0      0.1 (02/08/11 11:59:41)
2094 2015 tTrapd        0.0  0.0
2095 2015 tOspf6SpfTimer 0.0  0.0
2096 2015 tSpfTimer     0.0  0.0
2097 2015 tBgpTask      0.0  0.0
2100 2015 tTrapd        0.0  0.0

```

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

```
VSP-9012:1#show khi performance slabinfo SF4
Slot:SF4
```

Name	Active Objs	Num Objs	Objsize	Objper slab	Pageper slab	Active Slabs	Num Slabs
TIPC	12	12	320	12	1	1	1
tipc_queue_items	0	0	16	203	1	0	0
rpc_buffers	8	8	2048	2	1	4	4
rpc_tasks	20	8	192	20	1	1	1
rpc_inode_cache	18	10	416	9	1	2	2
merc_sock	0	0	352	11	1	0	0
UNIX	11	2	352	11	1	1	1
tcp_bind_bucket	203	6	16	203	1	1	1
inet_peer_cache	59	1	64	59	1	1	1
ip_fib_alias	113	34	32	113	1	1	1
ip_fib_hash	113	29	32	113	1	1	1
ip_dst_cache	0	0	256	15	1	0	0
arp_cache	0	0	128	30	1	0	0
RAW	9	2	448	9	1	1	1
UDP	16	10	480	8	1	2	2
tw_sock_TCP	0	0	96	40	1	0	0
request_sock_TCP	0	0	64	59	1	0	0
TCP	8	6	960	4	1	2	2
cfq_ioc_pool	0	0	48	78	1	0	0
cfq_pool	0	0	96	40	1	0	0
crq_pool	0	0	44	84	1	0	0
deadline_drq	0	0	48	78	1	0	0
as_arq	0	0	60	63	1	0	0
rcfs_inode_cache	72	62	320	12	1	6	6
nfs_write_data	40	36	480	8	1	5	5
nfs_read_data	36	32	448	9	1	4	4
nfs_inode_cache	63	59	560	7	1	9	9
nfs_page	0	0	64	59	1	0	0
ext2_inode_cache	252	247	432	9	1	28	28
ext2_xattr	0	0	44	84	1	0	0
inotify_event_cache	0	0	28	127	1	0	0

Key Health Indicators

inotify_watch_cache	0	0	36	101	1	0	0
kioclx	0	0	160	24	1	0	0
kiocb	0	0	128	30	1	0	0
fasync_cache	0	0	16	203	1	0	0
shmem_inode_cache	660	652	400	10	1	66	66
posix_timers_cache	40	1	96	40	1	1	1
uid_cache	59	1	64	59	1	1	1
relayfs_inode_cache	12	2	312	12	1	1	1
blkdev_ioc	0	0	28	127	1	0	0
blkdev_queue	30	24	380	10	1	3	3
blkdev_requests	0	0	152	26	1	0	0
biovec-(256)	54	54	3072	2	2	27	27
biovec-128	110	109	1536	5	2	22	22
biovec-64	220	218	768	5	1	44	44
biovec-16	220	218	192	20	1	11	11
biovec-4	236	218	64	59	1	4	4
biovec-1	406	218	16	203	1	2	2
bio	295	256	64	59	1	5	5
file_lock_cache	0	0	96	40	1	0	0
sock_inode_cache	44	39	352	11	1	4	4
skbuff_fclone_cache	26	26	288	13	1	2	2
skbuff_head_cache	336	336	160	24	1	14	14
proc_inode_cache	312	311	320	12	1	26	26
sigqueue	26	4	148	26	1	1	1
radix_tree_node	924	923	276	14	1	66	66
bdev_cache	9	2	416	9	1	1	1
sysfs_dir_cache	2024	1985	40	92	1	22	22
mnt_cache	40	22	96	40	1	1	1
inode_cache	910	910	304	13	1	70	70
dentry_cache	2639	2579	136	29	1	91	91
filp	384	384	160	24	1	16	16
names_cache	1	1	4096	1	1	1	1
idr_layer_cache	116	97	136	29	1	4	4
buffer_head	624	568	48	78	1	8	8
mm_struct	28	20	544	7	1	4	4
vm_area_struct	1012	928	88	44	1	23	23
fs_cache	113	19	32	113	1	1	1
files_cache	27	20	448	9	1	3	3
signal_cache	44	36	352	11	1	4	4
sighand_cache	36	35	1312	3	1	12	12
task_struct	77	72	1056	7	2	11	11
anon_vma	339	318	8	339	1	1	1
size-131072 (DMA)	0	0	131072	1	32	0	0
size-131072	0	0	131072	1	32	0	0
size-65536 (DMA)	0	0	65536	1	16	0	0
size-65536	0	0	65536	1	16	0	0
size-32768 (DMA)	0	0	32768	1	8	0	0
size-32768	2	2	32768	1	8	2	2
size-16384 (DMA)	0	0	16384	1	4	0	0
size-16384	0	0	16384	1	4	0	0
size-8192 (DMA)	0	0	8192	1	2	0	0
size-8192	1	1	8192	1	2	1	1
size-4096 (DMA)	0	0	4096	1	1	0	0
size-4096	42	42	4096	1	1	42	42
size-2048 (DMA)	0	0	2048	2	1	0	0
size-2048	304	292	2048	2	1	152	152
size-1024 (DMA)	0	0	1024	4	1	0	0
size-1024	40	37	1024	4	1	10	10
size-512 (DMA)	0	0	512	8	1	0	0
size-512	136	136	512	8	1	17	17
size-256 (DMA)	0	0	256	15	1	0	0
size-256	90	90	256	15	1	6	6
size-192 (DMA)	0	0	192	20	1	0	0
size-192	40	27	192	20	1	2	2
size-128 (DMA)	0	0	128	30	1	0	0

size-128	900	893	128	30	1	30	30
size-96 (DMA)	0	0	96	40	1	0	0
size-96	560	537	96	40	1	14	14
size-64 (DMA)	0	0	64	59	1	0	0
size-32 (DMA)	0	0	32	113	1	0	0
size-64	472	448	64	59	1	8	8
size-32	1582	1582	32	113	1	14	14
kmem_cache	120	120	96	40	1	3	3

Variable definitions

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

Variable	Value
{slot[-slot][,....]}	Identifies the slot in one of the following formats: a single slot (3), a range of slots (3-4), or a series of slots (3,5,6). Valid slots are 1-12, SF1-SF6, or all.

Displaying KHI control processor information

Use the following commands to display key health information about the packets generated by interface modules, the type of packets and protocols received on a port, and protocol drops.

About this task

You can use KHI commands on all interface slots. If you do not specify a slot, information for all slots is shown.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display KHI statistics for packets generated by the interface modules and sent to the control processor:

```
show khi cpp iocop-statistics [<3-12>]
```

3. Display statistics for control packets that go to the control processor:

```
show khi cpp port-statistics [{slot/port[-slot/port][,....]}]
```

4. Display KHI information about packets dropped due to CP Limit violations:

```
show khi cpp protocol-drops
```

Example

```
VSP-9012:1#show khi cpp iocop-statistics 4
```

```
=====
      KHI CPP Details - IO COP Statistics - Slot 4
=====
Slot      IO Generated Packets      Rx Packets
-----
4          MAC_MGMT                    4
```

```
VSP-9012:1#show khi cpp port-statistics 4/1-4/6
```

```
=====
      KHI CPP Details - Port Statistics
```

Key Health Indicators

```
=====
Ports      Packet Type                               Rx Packets  Tx Packets
-----
-
4/1        Ether2_EAP(140)                           0           2
4/1        LLC_BPDU(456)                             9882        0
4/1        LLC_TDP(464)                              0          2312
4/6        Ether2_IPv4_TTL_EXP(50)                   13          0
4/6        Ether2_ARP_Other(129)                     3           0
4/6        LLC_BPDU(456)                             2312        0
4/6        LLC_TDP(464)                              0          2314
=====
```

```
VSP-9012:1#show khi cpp protocol-drops
```

```
=====
KHI CPP Details - Protocol Drop Counters
=====
Protocol ID  Discard Count
-----
LACP        12914288
=====
```

Variable definitions

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

Variable	Value
<3-12>	Specifies the slot number.
slot/port[-slot/port][,...]	Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).

Key Health Indicators using EDM

The Key Health Indicators (KHI) feature of Avaya Virtual Services Platform 9000 provides a subset of health information that allows for quick assessment of the overall operational state of the device.

Note:

The KHI feature is not intended to provide a comprehensive debugging solution. Instead, KHI identifies key information that could lead Avaya support personnel towards discovery of a specific failure. After the technician assesses the KHI information, further debugging is required to determine the specific reason for the fault.

Avaya recommends that you capture KHI information during normal operations to provide a baseline for Avaya support personnel when detecting fault situations.

Clearing KHI statistics

Clear KHI statistics.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > Edit**.
2. Click **Chassis**.
3. Click the **CPP Stats Control** tab.
4. Select the statistics you want to clear.
5. Click **Apply**.

CPP Stats Control field descriptions

Use the data in the following table to use the **CPP Stats Control** tab.

Name	Description
PortStatsClear	Clears port statistics.
IoCopStatsClear	Clears statistics for the processors on the interface modules.
ProtocolDropsClear	Clears dropped protocol statistics.

Viewing KHI forwarding information for first generation modules

View KHI forwarding information for first generation modules to see statistics and counters for each lane.

About this task

Use these statistics to know which IFP rules are hit and to understand why packets are dropped.

Procedure

1. In the Device Physical View, select a module.
2. In the navigation pane, expand the following folders: **Configuration > Edit**.
3. Click **Card**.
4. Click the **IFP** tab.

IFP field descriptions

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

Name	Description
Index	Shows the index number for the IFP rule.
Name	Shows the name of the IFP rule.
Slice0Ports	Shows the range of ports 1-24 or 1-8 based on the module type.
Slice0Cnt	Shows the counter for the range of ports 1-24 or 1-8 based on the module type.

Name	Description
Slice1Ports	Shows the range of ports 25-48 or 9-16 based on the module type.
Slice1Cnt	Shows the counter for the range of ports 25-48 or 9-16 based on the module type.
Slice2Ports	Shows the range of ports 17 - 24 based on the module type. This field appears only for first generation modules.
Slice2Cnt	Shows the counter for ports 17 - 24 based on the module type. This field appears only for 24-port modules. This field appears only for first generation modules.

Viewing KHI forwarding information for second generation modules

View KHI forwarding information to see statistics and counters for each lane in second generation modules.

Note:

The **IFP** tab only displays for second generation modules.

About this task

Use these statistics to know which IFP rules are hit and to understand why packets are dropped.

Procedure

1. In the Device Physical View, select a module.
2. In the navigation pane, expand the following folders: **Configuration > Edit**.
3. Click **Card**.
4. Click the **IFP** tab.

IFP field descriptions

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

Name	Description
Slot	Specifies the slot number.
Slice	Specifies the slice number. Second generation modules are three slice modules, with slice 0, slice 1, and slice 2.
Lane	Specifies the lane number.
TblType	Specifies the type of table. Ingress Filter Processor packet type (IFP PT) or Ingress Filter Processor pre-

Name	Description
	classification vector (IFP PCV) for second generation modules.
Index	Specifies the index of the Ingress Filter Processor (IFP) rule.
Name	Specifies the name of the IFP rule.
StatsPorts	Specifies the range of ports in a slice for the slot. Port ranges for the 9048XS-2 module are 1-16, 17-32, and 33-48.
StatsCnt	Specifies the counter for slice 0.

Viewing protocol drop counters

View protocol drop counters to see the number of packets dropped due to CP Limit violations.

Procedure

1. In the Device Physical View, select the chassis.
2. In the navigation pane, expand the following folders: **Configuration > Graph**.
3. Click **Chassis**.
4. Click the **Protocol Drop** tab.

Protocol Drop field descriptions

The **Protocol Drop** tab shows the number of packets dropped for the following protocol-violation counters:

- **DataExpCnt**
- **TtlExpCnt**
- **IpmcDataCnt**
- **MacLearningCnt**
- **IsIsCnt**
- **BgpCnt**
- **RipV1Cnt**
- **RipV2Cnt**
- **OspfMcCnt**
- **FtpCnt**
- **TftpCnt**
- **SnmpCnt**
- **TelnetCnt**
- **SshCnt**

- **RshCnt**
- **IstCtlCnt**
- **RadiusCnt**
- **NtpCnt**
- **DhcpCnt**
- **IcmpV4Cnt**
- **IcmpV6Cnt**
- **IgmpCnt**
- **PimMcCnt**
- **VrrpCnt**
- **ArpReqCnt**
- **ArpOtherCnt**
- **RarpReqCnt**
- **RarpOtherCnt**
- **SlppCnt**
- **BpduCnt**
- **TdpCnt**
- **EapCnt**
- **LacpCnt**
- **VlaccCnt**
- **MldV2Cnt**
- **LldpCnt**
- **HttpCnt**
- **PimUcCnt**
- **OspfUcCnt**
- **DnsCnt**
- **IpfixCnt**
- **TestPktCnt**
- **IcmpV4BcCnt**
- **OspfV6UcCnt**
- **OspfV6McCnt**
- **HopByHopCnt**
- **CfmCnt**
- **RipV6Cnt**
- **VrrpV6Cnt**

- **NdMcV6Cnt**
- **NdUcV6Cnt**
- **IcmpMcV6Cnt**
- **IcmpUcV6Cnt**
- **FragUcV6Cnt**
- **FragMcV6Cnt**
- **RloginCnt**

Displaying KHI port information

Use the following commands to display key health information about the types of control packets and protocols received on a port and sent to the control processor.

Procedure

1. In the Device Physical View, select a port.
2. In the navigation pane, expand the following folders: **Configuration > Graph**.
3. Click **Port**.
4. Click the **CPP Stats** tab.

CPP Stats field descriptions

Use the data in the following table to use the **CPP Stats** tab.

Name	Description
Port	Identifies the slot and port.
Packet	Shows the packet type.
PacketName	Shows the name of the packet.
RxPackets	Indicates the number of received packets on the port for the packet type.
TxPackets	Indicates the number of transmitted packets on the port for the packet type.

Viewing COP statistics

View COP statistics for packets generated on interface modules.

Procedure

1. In the Device Physical View, select a module.
2. In the navigation pane, expand the following folders: **Configuration > Graph**.
3. Click **Card**.

4. Click the **COP Stats** tab.
5. To graph the statistics, select the information you want to graph, and then click the type of graph you want to create.

COP Stats field descriptions

Use the data in the following table to use the **COP Stats** tab.

Name	Description
MacMgmtRxPackets	Shows the number of received MAC management packets.
IpFixRxPackets	Shows the number of received IPFIX packets.
AbsoluteValue	Displays the counter value.
Cumulative	Displays the total value since you opened the COP Stats tab.
Average/sec	Displays the average value for each second.
Minimum/sec	Displays the minimum value for each second.
Maximum/sec	Displays the maximum value for each second.
LastVal/sec	Displays the last value for each second.

Displaying second generation QE statistics

Use the following procedure to display drop QE statistics Key Health Indicator (KHI) information for second generation modules.

Note:

The **Qe** tab only displays for second generation modules.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit**.
2. Click **Card**.
3. Click the **Qe** tab.

Qe field descriptions

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

Name	Description
QeSlot	Specifies the slot number.
QeSlice	Specifies the slice number.
Hg0RxPkts	Specifies the packets received through the HiGig0, specifies the proprietary interconnect consisting of XAUI electrical protocol with additional headers.

Name	Description
Hg1RxPkts	Specifies the packets received through the HiGig1, specifies the proprietary interconnect consisting of XAUI electrical protocol with additional headers.
Hg2RxPkts	Specifies the packets received through the HiGig2, specifies the proprietary interconnect consisting of XAUI electrical protocol with additional headers.
Hg3RxPkts	Specifies the packets received through the HiGig3, specifies the proprietary interconnect consisting of XAUI electrical protocol with additional headers.
Hg0TxPkts	Specifies the packets transmitted through the HiGig0, specifies the proprietary interconnect consisting of XAUI electrical protocol with additional headers.
Hg1TxPkts	Specifies the packets transmitted through the HiGig1, specifies the proprietary interconnect consisting of XAUI electrical protocol with additional headers.
Hg2TxPkts	Specifies the packets transmitted through the HiGig2, specifies the proprietary interconnect consisting of XAUI electrical protocol with additional headers.
Hg3TxPkts	Specifies the packets transmitted through the HiGig3, specifies the proprietary interconnect consisting of XAUI electrical protocol with additional headers.
RbIF1DropCnt	Specifies the packets dropped on the IF1 (HG0).
RbIF2DropCnt	Specifies the packets dropped on the IF1 (HG1).
RbIF3DropCnt	Specifies the packets dropped on the IF1 (HG2).
RbIF4DropCnt	Specifies the packets dropped on the IF1 (HG3).
FdPktDropCount	Specifies the packets dropped in the FDM block.
FrRxPktschannelA	Specifies the packets received from the fabric through channel A.
FrRxPktschannelB	Specifies the packets received from the fabric through channel B.
NumDP0PktsAccepted	Specifies the number of data port 0 packets accepted.
NumDP1PktsAccepted	Specifies the number of data port 1 packets accepted.
NumDP2PktsAccepted	Specifies the number of data port 2 packets accepted.
NumDP3PktsAccepted	Specifies the number of DP3 packets accepted.
NumDP0PktsDroppedDueToWRED	Specifies the DP0 packets dropped due to Weighted Random Early Detection (WRED).

Name	Description
NumDP1PktsDroppedDueToWRED	Specifies the DP1 packets dropped due to Weighted Random Early Detection (WRED).
NumDP2PktsDroppedDueToWRED	Specifies the DP2 packets dropped due to Weighted Random Early Detection (WRED).
NumDP3PktsDroppedDueToWRED	Specifies the DP3 packets dropped due to Weighted Random Early Detection (WRED).
NumDP0PktsMarkedDueToWRED	Specifies the DP0 packets marked due to Weighted Random Early Detection (WRED).
NumDP1PktsMarkedDueToWRED	Specifies the DP1 packets marked due to Weighted Random Early Detection (WRED).
NumDP2PktsMarkedDueToWRED	Specifies the DP2 packets marked due to Weighted Random Early Detection (WRED).
NumDP3PktsMarkedDueToWRED	Specifies the DP3 packets marked due to WRED.
NumDP0PktsDroppedNonWRED	Specifies the DP0 packets dropped non-Weighted Random Early Detection (WRED).
NumDP1PktsDroppedNonWRED	Specifies the DP1 packets dropped non-Weighted Random Early Detection (WRED).
NumDP2PktsDroppedNonWRED	Specifies the DP2 packets dropped non-Weighted Random Early Detection (WRED).
NumDP3PktsDroppedNonWRED	Specifies the DP3 packets dropped non-Weighted Random Early Detection (WRED).
NumPacketsDiscardedForBadQueueNum	Specifies the number of packets dropped due to a bad queue number.
NumQueuesAgedOut	Specifies the queues deleted due to age.
NumPktsDequeuedFromQM	Specifies the number of packets QM dequeued.
NumHcfcMsgsDroppedDueToCrError	Specifies the number of messages dropped due to a cyclic redundancy check (CRC) error.

Displaying second generation IFP drop information

Use the following procedure to display IFP drop statistics Key Health Indicator (KHI) information for second generation modules.

Note:

The **IfpDrop** tab only displays for second generation modules.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit**.
2. Click **Card**.
3. Click the **IfpDrop** tab.

IfpDrop field descriptions

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

Name	Description
Slot	Specifies the slot.
Slice	Specifies the slice.
Lane	Specifies the lane.
McRateLimitDrop	Specifies a multicast rate limit drop.
BcRateLimitDrop	Specifies a broadcast rate limit drop.
Ipv4TcpMcDrop	Specifies an IPv4 address Transmission Control Protocol (TCP) multicast drop.
IfpDrop	Specifies an Internet Filtering Protocol (IFP) drop.
VlanMemDrop	Specifies a VLAN member drop.
BlacklistDrop	Specifies a blacklist drop.
BadIpv4AddrDrop	Specifies a bad IPv4 address drop.
MacLoopbackDrop	Specifies a MAC loopback drop.
1dotPDrop	Specifies a tagged frame with VID=0 drop.
CfiDrop	Specifies a canonical format identifier (CFI) drop.
UntagDrop	Specifies an untag drop.
TagDrop	Specifies a tag drop.
IstFilterDrop	Specifies an IST filter drop.
IpVerDrop	Specifies an IP version drop.
Vid0Drop	Specifies a VLAN ID 0 (Vid0) drop.
UnsupportFrameDrop	Specifies an unsupported frame drop.
Ipv4MacTtl0Drop	Specifies an IPv4 MAC time-to-live (TTL) drop.

Displaying second generation RSP drop information

Use the following procedure to display drop RSP statistics Key Health Indicator (KHI) information for second generation modules.

Note:

The **RspDrop** tab only displays for second generation modules.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit**.
2. Click **Card**.
3. Click the **RspDrop** tab.

RspDrop field descriptions

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

Name	Description
Slot	Specifies the slot number.
Slice	Specifies the slice number.
Reg	Specifies the register.
Hab0RxErrCounterReg	Specifies the count of errored packets received through the high speed ASIC bus (HAB) interface.
Hab1RxErrCounterReg	Specifies the count of errored packets received through the HAB interface.
Cif0RxErrCounterReg	Specifies the count of errored packets received through the core interface (CIF).
Cif1RxErrCounterReg	Specifies the count of errored packets received through the CIF interface.
GsFloodingMeterRead	Specifies the flooding meter read.
GsMacLkupErrDA	Specifies the MAC destination address lookup error.
GsFilterActDeny	Specifies the filter act deny.
GsIstDstNni	Specifies the InterSwitch Trunk (IST) destination Network-to-Network Interface (NNI).
GsSrcMacLkup	Specifies the source MAC lookup.
GSUnknownSA	Specifies the unknown source address.
GsDstMacLkup	Specifies the destination MAC lookup.
GsSrcKnockOut	Specifies the source knock out.
GsNoUniNniMem	Specifies the no user-to-network (UNI) and network-to-network interface (NNI) member.
GsIstNoUniMem	Specifies no IST and UNI member.
GsCidDiscard	Specifies CID discard
GsLrnDisableUnkwnSA	Specifies the learn disable unknown source address.
GsInstMemParityExp	Specifies the InstMemParityExp.
GsPcOver4095Exp	Specifies the PcOver4095Exp.
GsPacketActiveExp	Specifies the packet active exp.
GsNonAlignedAddrExp	Specifies the non-aligned address exp.
GsIllegalRegionExp	Specifies the illegal region exp.
GsSplitOperActiveExp	Specifies the split operation active exp.
GsSplitOperNotDoneExp	Specifies the split operation not done exp.
GsGabTimeoutExp	Specifies the Global Access Bus (GAB) timeout exp.
GsInvalidOpTypeExp	Specifies the invalid op type exp.
GsIllegalArgExp	Specifies an illegal argument exp.

Name	Description
GsEccErrorExp	Specifies the Ecc error exp.
GsSearchError	Specifies a search error.
GsRadixSearchRunOnExp	Specifies a RADIX search run on exp.
GsSearchMgmtError	Specifies a search management error.
GsHwFailureExp	Specifies a hardware failure exp.
GsDontKnowExp	Specifies a do not know exp.
GsMacLkupErrSA	Specifies the MAC source address lookup error.
Gslpv4OcsErr	Specifies the IPv4 address OCS error.
GslsidLkupFail	Specifies the Service Instance Identifier (I-SID) lookup failure.
GsNniDstSpbDrop	Specifies the Network-to-Network (NNI) destination Shortest Path Bridging (SPB) drop.
GsNniIstExtraUniCpy	Specifies the NNI IST extra UNI copy.
GsNniUnkwNBDA	Specifies the NNI unknown NBDA.
GsNniIstDown	Specifies the NNI IST is down.
GsPcketErr	Specifies a packet error.
GslsidLkupMiss	Specifies an I-SID lookup miss.
GsNniIstPeerUpBvid2Drop	Specifies an NNI IST peer up backbone VLAN ID 2 drop.
GsUnknownMcastBda	Specifies the unknown multicast BDA.
GsNniCfmDrop	Specifies an NNI connectivity fault management (CFM) drop.
GslpNotEnDrop	Specifies an IP not En drop.
Gslpv6LkupFail	Specifies an IPv6 lookup fail.
Gslpv6OcsErr	Specifies an IPv6 Office Communications Server (OCS) error.
GslpmcV4Drop	Specifies an IP multicast version 4 drop.
GsV4LkupDrop	Specifies a version 4 lookup drop.

Displaying second generation RSP error information

Use the following procedure to display RSP error statistics Key Health Indicator (KHI) information for second generation modules.

Note:

The **RspError** tab only displays for second generation modules.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit**.
2. Click **Card**.
3. Click the **RspError** tab.

RspError field descriptions

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

Name	Description
Slot	Specifies the slot number.
Slice	Specifies the slice number.
Reg	Specifies the register.
HifErrorStatusReg	Specifies the HIF decoder error status register.
Hab0ErrorStatusReg	Specifies the high-speed ASIC bus (HAB) interface error status register.
Hab1ErrorStatusReg	Specifies the high-speed ASIC bus (HAB) interface error status register.
Erdi0EccSoftErrorCntReg to Erdi5EccSoftErrorCntReg	Specifies the Enhanced Remote Defect Indication (ERDI) memory interface soft error counter register.
Erdi0EccErrorCounterReg to Erdi5EccErrorCounterReg	Specifies the Enhanced Remote Defect Indication (ERDI) memory interface error counter register.
Cif0DropInfoReg	Specifies the core interface dropped information register.
Cif1DropInfoReg	Specifies the core interface dropped information register.
Hab0DropInfoReg	Specifies the high-speed ASIC bus (HAB) interface dropped information register.
Hab1DropInfoReg	Specifies the high-speed ASIC bus (HAB) interface dropped information register.

Displaying second generation RSP forwarding information

Use the following procedure to display RSP forwarding information for second generation modules.

Note:

The **RspForw** tab only displays for second generation modules.

Procedure

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

3. Click the **RspForw** tab.

RspForw field descriptions

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

Name	Description
Slot	Specifies the slot.
Slice	Specifies the slice.
Reg	Specifies the register.
Hab0RxPktCounterReg	Specifies the count of high-speed ASIC bus (HAB) interface packets received.
Hab0TxPktCounterReg	Specifies the count of high-speed ASIC bus (HAB) interface packets transmitted.
Hab1RxPktCounterReg	Specifies the count of high-speed ASIC bus (HAB) interface packets received.
Hab1TxPktCounterReg	Specifies the count of high-speed ASIC bus (HAB) interface packets transmitted.
Cif0RxPktCntReg	Specifies the core interface received packet count register.
Cif0TxPktCntReg	Specifies the core interface transmitted packet count register.
Cif1RxPktCntReg	Specifies the core interface received packet count register.
Cif1TxPktCntReg	Specifies the core interface transmitted packet count register.
PbmHab0RxPktCntReg	Specifies the packet buffer memory (PBM) base high-speed ASIC bus (HAB) interface received packet count register.
PbmHab0TxPktCntReg	Specifies the packet buffer memory (PBM) base high-speed ASIC bus (HAB) interface transmitted packet count register.
PbmHab1RxPktCntReg	Specifies the packet buffer memory (PBM) base high-speed ASIC bus (HAB) interface received packet count register.
PbmHab1TxPktCntReg	Specifies the packet buffer memory (PBM) base high-speed ASIC bus (HAB) interface transmitted packet count register.
PbmHab0Hab1SopEopCntReg	Specifies the number of EOP and SOP fragments received on the high-speed ASIC bus (HAB) interfaces for the packet buffer memory (PBM) base.
SamPacketInCntReg	Specifies the count of packets into the counter register for the streamlined analysis machine packet control global access bus (PC global access bus) base.

Name	Description
SamPacketOutCntReg	Specifies the count of packets out of the counter register for the streamlined analysis machine packet control global access bus base.
SamAe0GabTransRegTx	Specifies the number of transactions transmitted on the Ae0 global access bus interface for the streamlined analysis machine packet control global access bus base.
SamAe0GabTransRegRx	Specifies the number of transactions received on the Ae0 global access bus interface for the streamlined analysis machine packet control global access bus base.
SamAe1GabTransRegTx	Specifies the number of transactions transmitted on the Ae1 global access bus interface for the streamlined analysis machine packet control global access bus base.
SamAe1GabTransRegRx	Specifies the number of transactions received on the Ae0 global access bus interface for the streamlined analysis machine packet control global access bus base.
SamAe2GabTransRegTx	Specifies the number of transactions transmitted on the Ae0 global access bus interface for the streamlined analysis machine packet control global access bus base.
SamAe2GabTransRegRx	Specifies the number of transactions received on the Ae0 global access bus interface for the streamlined analysis machine packet control global access bus base.
SamAe3GabTransRegTx	Specifies the number of transactions transmitted on the Ae0 global access bus interface for the streamlined analysis machine packet control global access bus base.
SamAe3GabTransRegRx	Specifies the number of transactions received on the Ae0 global access bus interface for the streamlined analysis machine packet control global access bus base.
SamR6Le0GabTransRegTx	Specifies the number of transactions transmitted on the R6Le0 global access bus interface for the streamlined analysis machine packet control global access bus base.
SamR6Le0GabTransRegRx	Specifies the number of transactions received on the R6Le0 global access bus interface for the streamlined analysis machine packet control global access bus base.
SamR6Le1GabTransRegTx	Specifies the number of transactions transmitted on the R6Le1 global access bus interface for the

Name	Description
	streamlined analysis machine packet control global access bus base.
SamR6Le1GabTransRegRx	Specifies the number of transactions received on the R6Le1 global access bus interface for the streamlined analysis machine packet control global access bus base.
SamR6Le2GabTransRegTx	Specifies the number of transactions transmitted on the R6Le2 global access bus interface for the streamlined analysis machine packet control global access bus base.
SamR6Le2GabTransRegRx	Specifies the number of transactions transmitted on the R6Le2 global access bus interface for the streamlined analysis machine packet control global access bus base.
SamR6Le3GabTransRegTx	Specifies the number of transactions transmitted on the R6Le0 global access bus interface for the streamlined analysis machine packet control global access bus base.
SamR6Le3GabTransRegRx	Specifies the number of transactions received on the R6Le0 global access bus interface for the streamlined analysis machine packet control global access bus base.
SamR5LeGabTransRegTx	Specifies the number of transactions transmitted on the R5Le global access bus interface for the streamlined analysis machine packet control global access bus base.
SamR5LeGabTransRegRx	Specifies the number of transactions received on the R5Le global access bus interface for the streamlined analysis machine packet control global access bus base.
SamCifGabTransRegTx	Specifies the number of transactions transmitted on the core interface global access bus for the streamlined analysis machine packet control global access bus base.
SamCifGabTransRegRx	Specifies the number of transactions received on the core interface global access bus interface for the streamlined analysis machine packet control global access bus base.
SanHle0GabTransRegTx	Specifies the number of transactions transmitted through the hash-based lookup engine 0 (Hle0) for the streamlined analysis machine packet control global access bus (PC Gab) base.
SanHle0GabTransRegRx	Specifies the number of transactions received through the hash-based lookup engine (Hle0) for the

Name	Description
	streamlined analysis machine packet control global access bus (PC Gab) base.
SamAIsGabTransRegTx	Specifies the number of transactions transmitted through the Atomic Load and Store Engine for the streamlined analysis machine packet control global access bus (PC Gab) base.
SamAIsGabTransRegRx	Specifies the number of transactions received through the Atomic Load and Store Engine for the streamlined analysis machine packet control global access bus (PC Gab) base.
AmPaGabPacketInReg	Tracks the number of packets received on the PA global access bus from the packet buffer memory (PBM) base.
AmPrGabPacketOutReg	Tracks the number of packets sent from the Analysis Machine on the PR global access bus to the packet buffer memory (PBM) base.
AmAe0GabTransRegTx	Specifies the number of transactions transmitted out on the atomic engine base0 global access bus (Ae0 Gab) interface.
AmAe0GabTransRegRx	Specifies the number of transactions received on the atomic engine base0 global access bus (Ae0 Gab) interface.
AmAe1GabTransRegTx	Specifies the number of transactions transmitted out on the atomic engine base1 global access bus (Ae1 Gab) interface.
AmAe1GabTransRegRx	Specifies the number of transactions received out on the atomic engine base1 global access bus (Ae1 Gab) interface.
AmAe2GabTransRegTx	Specifies the number of transactions transmitted out on the atomic engine base2 global access bus (Ae2 Gab) interface.
AmAe2GabTransRegRx	Specifies the number of transactions received out on the atomic engine base2 global access bus (Ae2 Gab) interface.
AmAe3GabTransRegTx	Specifies the number of transactions transmitted out on the atomic engine base3 global access bus (Ae3 Gab) interface.
AmAe3GabTransRegRx	Specifies the number of transactions received out on the atomic engine base3 global access bus (Ae3 Gab) interface.
AmR6Le0GabTransRegTx	Specifies the number of transactions transmitted on the R6le0 Gab interface.
AmR6Le0GabTransRegRx	Specifies the number of transactions received on the R6le0 global access bus interface.

Name	Description
AmR6Le1GabTransRegTx	Specifies the number of transactions transmitted on the R6le1 global access bus interface.
AmR6Le1 GabTransRegRx	Specifies the number of transactions received on the R6le1 global access bus interface.
AmR6Le2GabTransRegTx	Specifies the number of transactions transmitted on the R6le2 global access bus interface.
AmR6Le2 GabTransRegRx	Specifies the number of transactions received on the R6le2 global access bus interface.
AmR6Le3GabTransRegTx	Specifies the number of transactions transmitted on the R6le3 global access bus interface.
AmR6Le3GabTransRegRx	Specifies the number of transactions received on the R6le0 global access bus interface.
AmR5LeGabTransRegTx	Specifies the number of transactions transmitted on the R5le global access bus interface.
AmR5LeGabTransRegRx	Specifies the number of transactions received on the R5le global access bus interface.
AmCifGabTransRegTx	Specifies the number of transactions transmitted on the core interface global access bus interface.
AmCifGabTransRegRx	Specifies the number of transactions received on the core interface global access bus interface.
AmAlseGabTransRegTx	Specifies the number of transactions transmitted on the Atomic Load and Store Engine global access bus (AlseGab) interface.
AmAlseGabTransRegRx	Specifies the number of transactions received on the Atomic Load and Store Engine global access bus (Alse Gab) interface.
AmHle0GabTransRegTx	Specifies the number of transactions transmitted on the hash-based lookup engine 0 global access bus (Hle0 Gab) interface.
AmHle0GabTransRegRx	Specifies the number of transactions received on the hash-based lookup engine 0 global access bus (Hle0 Gab) interface.

Displaying second generation RSP statistics

Use the following procedure to display RSP statistics for second generation modules.

Note:

The **RspStats** tab only displays for second generation modules.

Procedure

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

2. Click **Card**.
3. Click the **RspStats** tab.

RspStats field descriptions

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

Name	Description
Slot	Specifies the slot.
Slice	Specifies the slice.
Reg	Specifies the register.
DeviceldReg	Specifies the device ID register.
PIIStatusReg	Specifies the PII status register.
Edi0StatusReg to Edi5StatusReg	Specifies the EDI0 to ED5 memory controller status.
AeStatusReg	Specifies the Atomic Engine base status.
Alse0StatusReg	Specifies the Atomic Load and Store Engine 0 (ALSE0) status.
Hab0StatusReg	Specifies the high-speed ASIC bus (HAB) 0 interface status register.
Hab1StatusReg	Specifies the high-speed ASIC bus (HAB) 1 interface status.
Cif0StatusReg	Specifies the core 0 interface status.
Cif1StatusReg	Specifies the core 1 interface status.
HleStatusReg	Specifies the hash-based lookup engine status.
PbmStatusReg	Specifies the packet buffer memory (PBM) status.
SamThreadActive0Reg to SamThreadActive7Reg	Specifies the status of a thread that is enabled and actively processing a packet or background maintenance task.
SamThreadCountReg	Specifies the total count of available threads.
AmPcGabPbmThreadStartsReg	Specifies the number of threads that were started from a packet control 1 global access bus (Pc1Gab) request.
AmPcGabBkgrdThreadReg	Specifies the number of background threads initiated from the packet Control Global Access Bus (PcGab).
AmThreadActiveReg	Specifies the status that indicates that a thread is enabled and actively processing a packet or background maintenance task.
AmThreadCountReg	Specifies the total count of available threads.

Displaying second generation sierra drop statistics

Use the following procedure to display sierra drop statistics Key Health Indicator (KHI) information for second generation modules.

Note:

The **SierraDropStats** tab only displays for second generation modules.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit**.
2. Click **Card**.
3. Click the **SierraDropStats** tab.

SierraDropStats field descriptions

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

Name	Description
DropStatsSlot	Specifies the slot number.
GeCntRxErr	Specifies the count of error packets received by the receiving RGMII logic.
GeCntTxErr	Specifies the count of packets transmitted by the GE RGMII logic that were marked as error beforehand.
PktCntGeDrop	Specifies the count of packets received by the GE interface, and then dropped.
PktCntZ0GeErr to PktCntZ5GeErr	Specifies the count of error packets sent to the GE interface from the ZIP interfaces, numbered 0 to 5.
ZipRxFragDatErrCnt0 to ZipRxFragDatErrCnt5	Specifies the count of errored data fragments received from the ZIP interface, numbered from 0 to 5.
ZipRxFragRspErrCnt0 to ZipRxFragRspErrCnt5	Specifies the count of errored Route Switch Processor (RSP) sub command fragments received across the ZIP interface.
ZipRxFragPcieErrCnt0 to ZipRxFragPcieErrCnt5	Specifies the count of error SB/Peripheral Component Interconnect Express (PCIe) sub command fragments received across the ZIP interface.
ZipRxFragSephErrCnt0 to ZipRxFragSephErrCnt5	Specifies the count of errored high priority SEP high priority sub-command fragments received across the ZIP interface.
ZipRxFragSepLErrCnt0 to ZipRxFragSepLErrCnt5	Specifies the count of error SEP low priority sub-command fragments received.

Name	Description
ZipRxFragDropCnt0 to ZipRxFragDropCnt5	Specifies the count of dropped fragments (usually due to bad headers).
ZipRxFragInvctlCnt0 to ZipRxFragInvctlCnt5	Specifies the count of dropped fragments (usually due to bad headers).

Displaying second generation sierra forwarding statistics

Use the following procedure to display Sierra forwarding statistics for second generation modules.

Note:

The **SierraForwStats** tab only displays for second generation modules.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit**.
2. Click **Card**.
3. Click the **SierraForwStats** tab.

SierraForwStats field descriptions

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

Name	Description
ForwStatsSlot	Specifies the forward statistics slot.
GeCntRxOk	Counts the GE interface received packets.
GeCntTxOk	Counts the GE interface transmitted packets.
GeCntRxPause	Counts the GE interface valid received pause packets.
PkCntGeGe	Counts the loopback GE interface packets.
PktCntGeZ0 to PktCntGeZ5	Counts the packets sent to ZIP0 to ZIP5.
PktCntZ0GeOk to PktCntZ5GeOk	Counts the OK packets sent from ZIPs 1 to 5 to the GE interface.
SepPcieCnt0 to SepPcieCnt6	Counts the number of context 0 to context6 RetSubCmd packets written back into the host memory.

Displaying second generation sierra extended forwarding statistics

Use the following procedure to display extended Sierra forwarding statistics for second generation modules.

Note:

The **SierraExtendedForwStats** tab only displays for second generation modules.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit**.
2. Click **Card**.
3. Click the **SierraExtendedForwStats** tab.

SierraExtendedForwStats field descriptions

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

Name	Description
ExtendedForwStatsSlot	Specifies the extended forward statistics slot.
ZipTxFragDatCnt0 to ZipTxFragDatCnt5	Counts the data fragments transmitted across the ZIP interface from either the complex event processing (CEP), Diag 0, or Diag 1 interfaces.
ZipTxFragRFU0Cnt0 to ZipTxFragRFU0Cnt5	This option is reserved for future use.
ZipTxFragRspCnt0 to ZipTxFragRspCnt5	Counts the route switch processor (RSP) SubCmdRet control fragments transmitted across the ZIP interface.
ZipTxFragPcieCnt0 to ZipTxFragPcieCnt5	Counts the peripheral component interconnect express (PCIe) interface SubCmdRet control fragments transmitted across the ZIP interface.
ZipTxFragSephCnt0 to ZipTxFragSephCnt5	Counts the high priority script engine SubCmdRet transmitted across the ZIP interface.
ZipTxFragSep1Cnt0 to ZipTxFragSep1Cnt5	Counts the low priority script engine SubCmdRet control fragments transmitted across the ZIP interface.
ZipTxFragRFU1Cnt0 to ZipTxFragRFU1Cnt5	This option is reserved for future use.
ZipTxFragRFU2Cnt0 to ZipTxFragRFU2Cnt5	This option is reserved for future use.
ZipRxFragDatOkCnt0 to ZipRxFragDatOkCnt5	Counts the non-errored data fragments received across the ZIP interface.
ZipRxFragRspOkCnt0 to ZipRxFragRspOkCnt5	Counts the non-errored route switch processor (RSP) SubCmd data fragments received across the ZIP interface.
ZipRxFragPcieOkCnt0 to ZipRxFragPcieOkCnt5	Counts the non-errored peripheral component interconnect express (PCIe) SubCmd control fragments received across the ZIP interface.
ZipRxFragSephOkCnt0 to ZipRxFragSephOkCnt5	Counts the non-errored high priority script engine SubCmd control fragments received across the ZIP interface.

Name	Description
ZipRxFragSep1OkCnt0 to ZipRxFragSep1OkCnt5	Counts the non-errored low priority script engine SEP1 SubCmd control fragments received across the ZIP interface.

Displaying second generation sierra state information

Use the following procedure to display sierra state information for second generation modules.

Note:

The **SierraState** tab only displays for second generation modules.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit**.
2. Click **Card**.
3. Click the **SierraState** tab.

SierraState field descriptions

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

Name	Description
StatusSlot	Specifies the slot.
DevId	Specifies the development ID.
Build	Specifies the build.
BadData	Specifies bad data.
SbRexEvent	Specifies various events that have occurred related to the SBus and (Register EXtention) RexBus interfaces.
RexStatus	<p>Specifies the Register EXtension (REX) status. The contents of this register are updated after every RexBus access. When hardware detects an error it will set either the PIM0_REX_ERR or PIM1_REX_ERR bits in the SB_REX_EVENT register. At that time the software should look at the appropriate REX status fields to determine the type of error that occurred.</p> <p>The *_REX_TO bit indicates a timeout occurred.</p> <p>*_REX_ERRCODE being nonzero is also obvious. If *_REX_ERRCODE is zero, then *_REX_CRPAR not equal to *_REX_PAR indicates a parity error on a read has occurred (only check for this if the</p>

Name	Description
	PIM*_REX_ERR bit in the SB_REX_EVENT register is set).
SbStatus	Specifies the SB status.
SliceXWindowEvent	Specifies SliceX window event information. This register captures error information on failed accesses to and from the SliceX Window.
PcieCoreStatus	Specifies the peripheral component interconnect express (PCIe) core status. This register is used to show 32 of the 512 available status signals of the Altera PCIe core. It is used for debug purposes.
PcieEvent	Specifies the peripheral component interconnect express (PCIe) event. This register indicates when various PCIe related events have occurred.
PSI1Status	Specifies port interface module (PIM) status interfaces (PSI) 1.
PSI0Status	Specifies port interface module (PIM) status interfaces (PSI) 0.
PSIEvent	Specifies port interface module (PIM) status interfaces (PSI) event information. This register determines which bits of the PSI event register can set the PSI event bit in the main event register.
PSI1LinkStatus	Specifies port interface module (PIM) status interfaces (PSI) 1 link status information. This register holds the contents of a link status message received by PSI 1. This data is remapped into the per lane port interface module (PIM) LS register using the PSI LS REMAP memory.
PSI0LinkStatus	Specifies port interface module (PIM) status interfaces (PSI) 0 link status information. This register holds the contents of a link status message received by PSI 0. This data is remapped into the per lane PIM LS register using the PSI LS REMAP memory.
GeStatus	Specifies the GE status information. This register shows status for the GE GRMII interface.
GeEvent	Specifies the GE event information. This register indicates which GE related events have occurred since last cleared. Events can be cleared by writing a 1 to their bit positions.
PktEvent	Specifies the Pkt event information. This register indicates which packet switching related events have occurred since last cleared. Events can be cleared by writing a 1 to their bit positions. The DROP_NO_ENDST and DROP_NO_DST indicate that a packet was not sent anywhere.

Name	Description
	The DROP_ZIP_DIS and DROP_ZIP_DOWN bitfields are for packets that were sent to at least one ZIP interface, but were also intended to be sent out others. Such packets are not counted in the PKT_CNT_GE_DROP counter and indicate either a software misconfiguration, such as sending to a disabled ZIP interface, or a serious error, such as the ZIP interface is downHi .
ZipStatus0 to ZipStatus5	Displays status for ZIP interfaces.
ZipEvent0 to ZipEvent5	Displays which events occurred for ZIP interfaces.
SbmStatus0 to SbmStatus5	Specifies the SideBandMessage (SBM) status register related for ZIP interfaces.
SbmEvent0 to SbmEvent5	Specifies the SideBandMessage event register related for ZIP interfaces.
SepEvent	Specifies various events that occurred in the script engine parsing logic section.

Displaying second generation zag status

Use the following procedure to display zag status for second generation modules.

Note:

The **ZagStatus** tab only displays for second generation modules.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit**.
2. Click **Card**.
3. Click the **ZagStatus** tab.

ZagStatus field descriptions

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

Name	Description
StatusSlot	Specifies the slot.
StatusSlice	Specifies the slice.
StatusLane	Specifies the lane.
Zag2DeviceIdBuild	Specifies the zagros 2 device ID build.
Zag2Build	Specifies the zagros 2 build.
Zag2Status	Specifies the zagros 2 status.
Zag2Interrupt	Specifies the zagros 2 interrupt register.

Name	Description
ZipStatus	Specifies the ZIP status register.
ZipSbmStatus	Specifies the ZIP SideBandMessage (SBM) status.
CifStatus	Specifies the core interface (CIF) status.
DpStatus0 to DpStatus7	Specifies the data port (DP) status.
DpInterrupt0 to DpInterrupt7	Specifies the data port (DP) interrupt mask.
FpStatus0 to FpStatus1	Specifies the fabric port (FP) status register 0 to FP status register 1.
HabStatus	Specifies the high-speed ASIC bus (HAB) interface status.
PmmInterrupt	Specifies the Packet Memory Manager (PMM) interrupt register.
PmmIngDataPageCnt	Specifies the Packet Memory Manager (PMM) ingress data page count of the number of pages currently in use.
PmmIngDataPeakPageCnt	Captures the highest value reached by the Packet Memory Manager ingress data page count register since the register was last set to zero.
PmmIngScPageCnt	Specifies the Packet Memory Manager (PMM) ingress SC page count of the number of pages currently in use.
PmmIngScOrdinaryLmt	Specifies the maximum number of pages reserved for ordinary (non-control) packets.
PmmIngScMaxLmt	Specifies the maximum number of pages reserved for this page pool.
PmmIngLsmPageCnt	Specifies the Packet Memory Manager (PMM) ingress LSM page count of the number of pages currently in use.
PmmIngLsmMaxLmt	Specifies the maximum number of pages reserved for this page pool.
PmmIngCopPageCnt	Specifies the number of pages currently in use.
PmmIngCopMaxLmt	Specifies the maximum number of pages reserved for this page pool.
PmmIngExpandPageCnt	Specifies the count of the number of pages currently in use.
PmmIngExpandMaxLmt	Specifies the maximum number of pages reserved for this page pool. The pool size is calculated by subtracting the total of that max limits for pools 0 to 3 from the data port page pool size in the page pool size register.
PmmEgrUcPageCnt	Specifies the Packet Memory Manager (PMM) egress UC page count of the number of pages currently in use.

Name	Description
PmmEgrUcPeakPageCnt	Specifies the Packet Memory Manager egress UC peak page count. This register captures the highest value reached by the egress UC page count register since this register was last reset to zero.
PmmEgrUcMaxLmt	Specifies the maximum number of pages reserved for this page pool.
PmmEgrScPageCnt	Specifies the Packet Memory Manager (PMM) egress SC page count of the number of pages currently in use.
PmmEgrScMaxLmt	Specifies the maximum number of pages reserved for this page pool.
PmmEgrMcPageCnt	Specifies the Packet Memory Manager (PMM) egress MC page count of the number of pages currently in use.
PmmEgrMcPeakPageCnt	Specifies the peak page count. This register captures the highest value reached by the egress MC page count register since this register was last reset to zero.
PmmEgrMcRootLmt	Specifies the Packet Memory Manager Egress MC root limit register. This is the maximum number of pages reserved for MC root packets.
PmmEgrMcMaxLmt	Specifies the Packet Memory Manager egress MC maximum limit register. Specifies the maximum number of pages reserved for this page pool.
PmmEgrOobPageCnt	Specifies the Packet Memory Manager (PMM) egress OOB page count of the number of pages currently in use.
PmmEgrOobMaxLmt	Specifies the Packet Memory Manager Egress OOB maximum limit register. This is the maximum number of pages reserved for this page pool.
PmmEgrExpandPageCnt	Specifies the Packet Memory Manager (PMM) egress expand page count of the number of pages currently in use.
PmmEgrExpandMaxLmt	Specifies the Packet Memory Manager (PMM) egress expand maximum limit, which is the maximum number of pages reserved for this page pool. This pool size is calculated by subtracting the total of the maximum limits for pools 0 to 3 from the data port page pool size in the page pool size register.
PmmReplicaEngPktLmt	Specifies the current number of outstanding root packets sent to the Packet Memory Manager (PMM) replica engine packet count register.

Name	Description
PmmReplicaEngPktCnt	Specifies the current number of outstanding root packets sent to the Packet Memory Manager (PMM) replica engine packet count register.
PmmReplicaEngPeakPktCnt	Specifies the Packet Memory Manager (PMM) replica engine peak packet count. This register captures the highest value reached by the PMM replica engine packet count register since this register was last reset to zero.
IdpStatus0 to IdpStatus1	Specifies the IDP Status0 to IDP Status1.
EhpStatus	Specifies the EHP control register status.
EhpFifoStatus	Specifies the EHP FIFO status register.
IhpCtlReg0	Specifies the IHP CTL register0.
IhpStatReg0	Specifies the IHP STAT register0.
IhpStatus1	Specifies the IHP status1. This register selects whether or not a port will insert zagros headers on all traffic.
IhpCtlReg1	Specifies the IHP CTL Reg1, which controls functions in the IHP.
IhpGenCfgReg	Specifies the IHP GEN CFG REG.
ZfaFifoStatusReg	Specifies the FIFO STATUS REG.
HbmEvents	Specifies HBM events.
HbmCreditCnt	Specifies the HBM credit counter register. Returns the value of the current available credits. A write to this register will update the credit counter to the value set in bits.
LsmLinkState	Specifies the link state of each port in zagros 2. The state is based on the status of the link, but is not the same as the status. Link state is included in all of the HBM LSM messages.
LsmHwLinkState	Specifies the hardware link state. This registers shows the hardware link state of each port in zagros2. The state is based on the status of the link, but it is not the same as the status. Hardware link state is used by the TXB in conjunction with the per port discard on down settings.
TxbDebug0 to TxbDebug3	Specifies the TXB debug0 to TXB debug3 fill level.
TxbEvent	Specifies if any of the TXB FIFOs have reached the FULL level.
TxbfifoFull	Specifies if the TXB FIFO is full or close to full.
DpExtStatus0 to DpExtStatus7	Specifies the data port (DP) ext status0 to data port ext status7.

Displaying second generation zagros error information

Use the following procedure to display zagros error information for second generation modules.

Note:

The **ZagrosError** tab only displays for second generation modules.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit**.
2. Click **Card**.
3. Click the **ZagrosError** tab.

ZagrosError field descriptions

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

Name	Description
ErrSlot	Specifies the slot.
ErrSlice	Specifies the slice.
ErrLane	Specifies the lane.
PmmRspErr	Specifies the Packet Memory Manager (PMM) route switch processor (RSP) error.
PmmDpErr	Specifies the Packet Memory Manager (PMM) data port (DP) error.
PmmFpErr	Specifies the Packet Memory Manager (PMM) fabric port (FP) error.
PmmPmErr	Specifies the Packet Memory Manager (PMM) data port PM error.
FpErr0	Specifies the fabric port (FP) error 0.
FpErr1	Specifies the fabric (FP) port error 1.
DpPortErr0 to DpPortErr7	Specifies the data port (DP) error 0 to data port (DP) error 7.
Zag2HapErr	Specifies the zagros2 host access port (HAP) error.
EtpEgrFifoErr	Specifies the ETP egress FIFO error information.
CifErrEvent	Specifies core interface (Cif) error event information.
HabErrEvent	Specifies high-speed ASIC bus (Hab) error event information.
CifRxDropInfo	Specifies core interface (CIF) received dropped information.
HabRxDropInfo	Specifies high-speed ASIC bus (HAB) interface received dropped information.

Name	Description
DpExtErr0 to DpExtErr7	Specifies data ports with errors.

Displaying second generation zag forward statistics

Use the following procedure to display zag drop extended statistics for second generation modules.

Note:

The **ZagForwStats** tab only displays for second generation modules.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit**.
2. Click **Card**.
3. Click the **ZagForwStats** tab.

ZagForwStats field descriptions

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

Name	Description
StatsSlot	Specifies the slot.
StatsSlice	Specifies the slice.
StatsLane	Specifies the lane.
PmmDpRxCnt	Specifies the packet count from the data port (DP) to the packet memory manager (PMM), which is the packet switch internal to the Zagros2 field programmable gate array (FPGA). This counts all packets.
PmmDpTxCnt	Specifies the packet count from the packet memory manager (PMM) to the data port (DP).
PmmRspRxCnt	Specifies the packet count from the route switch processor (RSP) to the packet memory manager (PMM).
PmmRspTxCnt	Specifies the packet count from the packet memory manager (PMM) to the route switch processor (RSP).
PmmFp0RxCnt	Specifies the packet count from fabric port (FP) 0 to the packet memory manager (PMM).
PmmFp0TxCnt	Specifies the packet count from the packet memory manager (PMM) to fabric port (FP) 0.
PmmFp1RxCnt	Specifies the packet count from fabric port (FP) 1 to the packet memory manager (PMM).

Name	Description
PmmFp1TxCnt	Specifies the packet count from the packet memory manager (PMM) to fabric port (FP) 1.
PmmFlopTxCnt	Specifies the packet count from the packet memory manager (PMM) to the fabric local port (FLOP).
PmmIngDataAdmitCnt	Specifies the aggregate count of the number of packets admitted across all ports.
PmmIngScAdmitCnt	Specifies the count of the number of packets admitted.
PmmIngLsmAdmitCnt	Specifies the count of the number of pages currently in use.
PmmIngCopAdmitCnt	Specifies the count of the number of packets admitted.
PmmIngExpAdmitCnt	Specifies the count of the number of packets admitted that were successfully expanded across all page pools. This occurs when the packet offset becomes less than 128 on return from the route switch processor.
PmmEgrUcAdmitCnt	Specifies the count of the number of pages currently in use.
PmmEgrScAdmitCnt	Specifies the count of the number of packets admitted.
PmmEgrScAdmitCnt	Specifies the count of the number of packets admitted.
PmmEgrMcRootAdmitCnt	Specifies the count of the number of MC root packets admitted.
PmmEgrOobAdmitCnt	Specifies the count of the number of packets admitted.
PmmEgrExpandAdmitCnt	Specifies the count of the number of packets admitted that were successfully expanded across all page pools.
FpRxPacketcount0 to FpRxPacketcount1	Counts the increments by one for each RX packet received (good or bad).
FpTxPacketCnt0 to FpTxPacketCnt1	Counts the increments by one for each TX packet transmitted.
HabRxPacketCnt	Specifies the number of packets received.
HabTxPacketCnt	Specifies the number of packets transmitted.
CifRxOperationCnt	Specifies the count of words received with errors.
CifTxOperationCnt	Specifies the count of words transmitted.
ZpktRxOkCnt	Counts the non-errored packets received across the ZIP interface and sent to the packet memory manager (PMM).

Name	Description
ZpktTxOkCnt	Counts the number of non-errored co-processor (COP) extraction packets sent across the ZIP interface, and also counts the non-errored diag loopback packets. The co-processor is the local processor in the linecard.
ZipTxfragDatCnt	Counts the data fragments transmitted across the ZIP interface from either complex event processing (CEP), diag 0, or diag 1 interfaces.
ZipTxfragRspCnt	Counts the route switch processor (RSP) subcmdret control fragments transmitted across the ZIP interface.
ZipTxfragPcieCnt	Counts the peripheral component interconnect express (PCIE) subcmdret control fragments transmitted across the ZIP interface.
ZipTxfragSephCnt	Counts the high priority script engine subcmdret control fragments transmitted across the ZIP interface.
ZipTxfragSepICnt	Counts the low priority script engine subcmdret control fragments transmitted across the ZIP interface.
ZipRxfragDatOkCnt	Counts the non-errored data fragments received across the ZIP interface.
ZipRxfragRspOkCnt	Counts the non-errored route switch processor (RSP) subcmd data fragments received across the ZIP interface.
ZipRxfragPcieOkCnt	Counts the non-errored peripheral component interconnect express (PCIE) subcmd control fragments received across the ZIP interface.
ZipRxfragSephOkCnt	Counts the non-errored high priority script engine subcmd control ragments received across the ZIP interface.
ZipRxfragSepIOkCnt	Counts the non-errored low priority script engine subcmd control ragments received across the ZIP interface.
PmmRspRxTestCount	Specifies the packet count from route switch processor (RSP) to packet memory manager (PMM).
PmmRspTxTestCount	Specifies the packet count from packet memory manager (PMM) to route switch processor (RSP).
PmmDpRxTestCount	Specifies the packet count from data port (DP) to packet memory manager (PMM). This counts all of the packets.
PmmDpTxTestCount	Specifies the packet count from packet memory manager (PMM) to data port (DP).

Name	Description
PmmFP0RxTestCount	Specifies the packet count from fabric port (FP) 0 to packet memory manager (PMM).
PmmFP0TxTestCount	Specifies the packet count from packet memory manager (PMM) to fabric port (FP) 0.
PmmFP1RxTestCount	Specifies the packet count from fabric port (FP) 1 to packet memory manager (PMM).
PmmFP1TxTestCount	Specifies the packet count from packet memory manager (PMM) to fabric port (FP) 1.
PmmFlopTxTestCount	Specifies the packet count from packet memory manager (PMM) to fabric local port (FLOP).

Displaying second generation zag drop statistics

Use the following procedure to display zag drop statistics for second generation modules.

Note:

The **ZagDropStats** tab only displays for second generation modules.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit**.
2. Click **Card**.
3. Click the **ZagDropStats** tab.

ZagDropStats field descriptions

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

Name	Description
StatsSlot	Specifies the slot.
StatsSlice	Specifies the slice.
StatsLane	Specifies the lane.
PmmRspCmdDropCnt	Specifies the Packet Memory Manager (PMM) command drop count. This is a count of the number of packets dropped as a result of the PDP=2'b11 (SP) and LDST=2'b1X in the zagros header of a packet returning from the route switch processor (RSP).
PmmFabricPort0FfeDrpCnt to PmmFabricPort1FfeDrpCnt	Counts the number of packets dropped in the fabric port 0 DMA or fabric port 1 DMA after FFE processing.
PmmPmCmdDropCnt	Counts the number of packets dropped as a result of the PM detecting a packet which if modified would

Name	Description
	cause either a RX_MIN_LEN_ERR or a RX_MAX_LEN_ERR.
PmmDpRxDropCnt	Specifies the packet drop count.
PmmIngDataCtlDropCnt	Specifies the control packet drop count.
PmmIngDataPort0DropCnt to PmmIngDataPort15DropCnt	Counts the number of pages currently in use.
PmmIngCopDropCnt	Specifies the packet drop count.
PmmIngExpandDropCnt	Counts the number of pages currently in use.
PmmEgrUcDropCnt	Specifies the packet drop count.
PmmEgrScDropCnt	Specifies the packet drop count.
PmmEgrMcRootT0DropCnt to PmmEgrMcRootT7DropCnt	Specifies the root limit, which is the maximum number of pages reserved for MC root packets.
PmmEgrMcCopyDropCnt	Specifies the MC copy drop count.
PmmEgrOobDropCnt	Specifies the packet drop count.
PmmEgrExpandDropCnt	Specifies the count of the number of pages currently in use.
PmmMpvMpidDropCnt	Specifies the count of the number of packets dropped because the MPID value of the copy exceeds the number of MPV FIFO ports configured.
ZfaInvalidGpidDropCnt	Specifies the count of packets dropped in the ZFA due to invalid GPID.
CifRxOperErrCnt	Specifies the count of words received with errors.
EhpStgDropCnt	Specifies the packet drop count.
EhpNoRcvrDropCnt	Specifies the packet drop count.
HcfcRxErrCnt	Specifies the RX HCFC frame error count.
ZipRxfragDatErrCnt	Counts the errored data fragments received across the ZIP interface.
ZipRxfragRspErrCnt	Counts the errored route switch processor (RSP) subcmd data fragments received across the ZIP interface.
ZipRxfragPcieErrCnt	Counts the errored peripheral component interconnect express (PCIE) subcmd control fragments received across the ZIP interface.
ZipRxfragSephErrCnt	Counts the errored high priority script engine subcmd control fragments received across the ZIP interface.
ZipRxfragSepLerrCnt	Counts the errored low priority script engine subcmd control fragments received across the ZIP interface.
ZipRxfragDropCnt	Counts the dropped fragments.
ZipRxfragInvctlCnt	Counts the control fragments with an invalid subcmd DST field.

Name	Description
ZpktRxErrCnt	Counts the errored packets received across the ZIP interface and sent to the packet memory manager lane 0.
ZpktTxErrCnt	Counts the non-errored co-processor (COP) extraction packets sent across the ZIP interface. Also counts the errored diag loopback packets.

Displaying second generation zag drop extended statistics

Use the following procedure to display zag drop extended statistics for second generation modules.

Note:

The **ZagDropExtendedStats** tab only displays for second generation modules.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit**.
2. Click **Card**.
3. Click the **ZagDropExtendedStats** tab.

ZagDropExtendedStats field descriptions

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

Name	Description
StatsSlot	Specifies the slot.
StatsSlice	Specifies the slice.
StatsLane	Specifies the lane.
FpRxPacketErrCnt0 and FpRxPacketErrCnt1	Specifies the count increments by one for each packet received with errors.
FpTxPacketErrCnt0 and FpTxPacketErrCnt1	Specifies the count increments by one for each packet transmitted with errors. These packets are transmitted with an inverted cyclic redundancy check (CRC).
FpRxJunkErrCnt0 and FpRxJunkErrCnt1	Specifies the count increments by one for each received junk detected. Junk is considered any control code that is not /SOP/ following a string of IDLEs in IPG.
PmmMpvFifoFullDropCnt0 to PmmMpvFifoFullDropCnt7	Counts the number of packets dropped because the MPV FIFO full drop condition is detected.
PmmMpvSkoDropCnt0 to PmmMpvSkoDropCnt7	Counts the number of packets dropped because the MPV SKO drop conditions is detected.

Name	Description
PmmMpvMitDropCnt0 to PmmMpvMitDropCnt7	Counts the number of packets dropped because the MPV MLT drop condition is detected.
PmmMpvFfeDropCnt0 to PmmMpvFfeDropCnt7	Counts the number of packets dropped because the MPV Ffe drop condition is detected.
PmmMpvStgDropCnt0 to PmmMpvStgDropCnt7	Counts the number of packets dropped because the MPV STG drop condition is detected.
TxbDropCnt0 to TxbDropCnt7	Counts the number of packets dropped because the DROP bit was set in the zagros header. It does not count packets dropped due to discard or drop-on-down conditions. These counters are not affected by reset and should be read-cleared at least once before you use them.
TxbDiscardCnt0 to TxbDiscardCnt7	Counts the number of packets discarded because the DISCARD_EN bit was set or because the DROP_ON_DOWN bit was set and the LinkState for that port was down. These counters are not affected by reset and should be read-cleared at least once before you use them.

Chapter 7: Link state change control

Rapid fluctuation in a port link state is called link flapping.

Link flapping is detrimental to network stability because it can trigger recalculation in spanning tree and the routing table.

If the number of port down events exceeds a configured limit during a specified interval, the system forces the port out of service.

You can configure link flap detection to control link state changes on a physical port. You can set thresholds for the number and frequency of changes allowed.

You can configure the system to take one of the following actions if changes exceed the thresholds:

- send a trap
- bring down the port

If changes exceed the link state change thresholds, the system generates a log entry.

Link state change control using ACLI

Detect and control link flapping to bring more stability to your network.

Controlling link state changes

Configure link flap detection to control state changes on a physical port.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure the interval for link state changes:

```
link-flap-detect interval <2-600>
```

3. Configure the number of changes allowed during the interval:

```
link-flap-detect frequency <1-9999>
```

4. Enable automatic port disabling:

```
link-flap-detect auto-port-down
```

5. Enable sending a trap:

```
link-flap-detect send-trap
```

Example

Enable automatic disabling of the port. Configure the link-flap-detect interval to 10. Enable sending traps.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#link-flap-detect auto-port-down
Switch:1(config)#link-flap-detect interval 10
Switch:1(config)#link-flap-detect send-trap
```

Variable definitions

Use the data in the following table to use the `link-flap-detect` command.

Variable	Value
<auto-port-down>	Automatically disables the port if state changes exceed the link-flap threshold. By default, auto-port-down is enabled. Use the no operator to remove this configuration. To set this option to the default value, use the default operator with the command.
frequency <1-9999>	Configures the number of changes that are permitted during the time specified by the interval command. The default is 20. To set this option to the default value, use the default operator with the command.
interval <2-600>	Configures the link-flap-detect interval in seconds. The default value is 60. To set this option to the default value, use the default operator with the command.
send-trap	Activates traps transmission. The default setting is activated. Use the no operator to remove this configuration. To set this option to the default value, use the default operator with the command.

Displaying link state changes

Displays link flap detection state changes on a physical port.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display link state changes:

```
show link-flap-detect
```

Example

```
VSP-9012:1>enable
VSP-9012:#show link-flap-detect

Auto Port Down : enable
Send Trap      : enable
Interval       : 60
Frequency      : 20
```

Link state change control using EDM

Detect and control link flapping to bring more stability to your network.

Controlling link state changes

Configure link flap detection to control link state changes on a physical port.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **General**.
3. Click the **Link Flap** tab.
4. Configure the parameters as required.
5. Click **Apply**.

Link Flap field descriptions

Use the data in the following table to use the **Link Flap** tab.

Name	Description
AutoPortDownEnable	Enables or disables Link Flap Detect. If you enable Link Flap Detect, the system monitors the number of times a port goes down during a designated interval. If the number of drops exceeds a specified limit, the system forces the port out-of-service. The default is enabled.
SendTrap	Specifies that a trap is sent if the port is forced out-of-service.
Frequency	Specifies the number of times the port can go down. The default is 20.
Interval	Specifies the interval (in seconds) between port failures. The default is 60.

Chapter 8: Log and trap information

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

Log and trap fundamentals

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

Simple Network Management Protocol

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

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

Important:

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

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

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

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

Overview of traps and logs

SNMP traps

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

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

System log messaging

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

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

The remote UNIX management workstation performs the following actions:

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

Log consolidation

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

The system log captures messages for the following components:

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

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

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

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

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

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

System log client over IPv6 transport

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

Log message format

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

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

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

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

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

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

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

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

The following table describes the system message severity levels.

Table 1: Severity levels

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

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

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

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

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

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

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

Table 2: Default and system log severity level mapping

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

Log files

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

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

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

Log file naming conventions

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

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

Log file transfer

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

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

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

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

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

- The maximum size of the log file.

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

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

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

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

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

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

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

Three parameters exist to configure the log file:

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

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

Log configuration using ACLI

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

Configuring a UNIX system log and syslog host

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

About this task

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

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enable the system log:

```
syslog enable
```

3. Specify the IP header in syslog packets:

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

4. Configure the maximum number of syslog hosts:

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

5. Create the syslog host:

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

6. Configure the IP address for the syslog host:

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

7. Enable the syslog host:

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

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

8. View the configuration to ensure it is correct:

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

Example

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

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

```

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

```

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

```

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

```

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

```

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

```

```

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

```

Variable definitions

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

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

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

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

Table 3: Variable definitions

Variable	Value
1–10	Creates and configures a host instance. Use the no operator before this parameter, no syslog host to delete a host instance.
address WORD <0–46>	Configures a host location for the syslog host. WORD <0–46> is the IPv4 or IPv6 address of the UNIX system syslog host in the format A.B.C.D or x:x:x:x:x:x. You can log system log messages to external system log hosts with both IPv4 and IPv6 addresses with no difference in functionality or configuration using ACLI.
enable	Enables the syslog host. Use the no operator before this parameter, no syslog host enable to disable syslog host. The default is disabled.
facility {local0 local1 local2 local3 local4 local5 local6 local7}	Specifies the UNIX facility in messages to the syslog host. {local0 local1 local2 local3 local4 local5 local6 local7} is the UNIX system syslog host facility. The default is local7.
maperror {emergency alert critical error warning notice info debug}	Specifies the syslog severity to use for error messages. The default is error.
mapfatal {emergency alert critical error warning notice info debug}	Specifies the syslog severity to use for fatal messages. The default is emergency.
mapinfo {emergency alert critical error warning notice info debug}	Specifies the syslog severity level to use for information messages. The default is info.
mapwarning {emergency alert critical error warning notice info debug}	Specifies the syslog severity to use for warning messages. The default is warning.
severity <info warning error fatal> [<info warning error fatal>] [<info warning error fatal>] [<info warning error fatal>]	Specifies the severity levels for which to send syslog messages for the specified modules. The default is info.
udp-port <514-530>	Specifies the User Datagram Protocol port number on which to send syslog messages to the syslog host. This value is

Variable	Value
	the UNIX system syslog host port number from 514–530. The default is 514.

Configuring logging

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

About this task

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

Procedure

1. Enter Global Configuration mode:


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


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


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


```
logging screen
```

Example

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

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

Variable definitions

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

Variable	Value
level <0-4>	Shows and configures the logging level. The level is one of the following values: <ul style="list-style-type: none"> • 0: Information — Records all messages.

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

Configuring the remote host address for log transfer

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

Before you begin

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

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure the remote host address for log transfer:

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

Example

Configure the remote host address for log transfer to 192.0.2.10:

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

Variable definitions

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

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

Configuring system logging to external storage

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

Before you begin

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

Caution:

Risk of data loss

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

About this task

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

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

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

Procedure

- Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable system logging to a CF card:

```
boot config flags logging
```

3. Configure the logfile parameters:

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

Example

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

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

Variable definitions

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

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

Configuring system message control

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

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure system message control action:

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

3. Configure the maximum number of messages:

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

4. Configure the interval:

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

5. Enable message control:

```
sys msg-control
```

Example

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

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

Variable definitions

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

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

Extending system message control

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

About this task

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

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure the force message control option:

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

Example

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

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

Variable definitions

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

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

Viewing logs

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

About this task

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

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

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

Procedure

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

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

Example

Display log information:

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

```

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

```

Variable definitions

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

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

Configuring ACLI logging

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

Note:

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

About this task

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

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

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

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

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

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

Procedure

1. Enter Global Configuration mode:


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


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


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

```
show clilog
```

5. View the ACLI log:

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

```
show logging file module clilog
```

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

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

Example

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

Variable definitions

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

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

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

Note:

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

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

Log configuration using EDM

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

Configuring the system log

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

Procedure

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

System Log field descriptions

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

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

Configuring the system log table

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

About this task

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

Procedure

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

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

System Log Table field descriptions

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

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

SNMP trap configuration using ACLI

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

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

Configuring an SNMP host

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

About this task

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

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure an SNMPv1 host:

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

3. Configure an SNMPv2c host:

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

4. Configure an SNMPv3 host:

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

5. Ensure that the configuration is correct:

```
show snmp-server host
```

Example

1. Configure the target table entry:

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

2. Configure an SNMPv3 host:

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

Variable definitions

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

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

Configuring an SNMP notify filter table

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

Before you begin

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

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Create a new notify filter table:

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

3. Ensure that the configuration is correct:

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

Example

Create a new notify filter table:

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

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

Variable definitions

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

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

Configuring SNMP interfaces

Configure an interface to send SNMP traps.

About this task

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

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

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

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

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

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

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

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

5. Activate the generation of authentication traps:

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

Example

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

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

Variable definitions

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

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

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

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

Enabling SNMP trap logging

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

Before you begin

- You must configure and enable the syslog server.

About this task

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

Procedure

1. Enter Global Configuration mode:


```
enable
configure terminal
```
2. Enable SNMP trap logging:


```
snmplog enable
```
3. Disable SNMP trap logging:


```
no snmplog enable
```
4. View the contents of the SNMP log:
 - a. View the SNMP log files generated for Release 3.2 and greater:


```
show logging file module snmplog
```
 - b. View the SNMP log files generated by releases prior to Release 3.2:


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

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
VSP-9012:1(config)#snmplog enable
VSP-9012:1(config)#show logging file module snmplog
CP1 [04/01/13 17:16:04.130] 0x002c4600 00000000 GlobalRouter SNMPLOG INFO 1
  ver=v2c private rcSysActionL1.0 = 7
CP1 [04/02/13 10:50:41.122] 0x002c4600 00000000 GlobalRouter SNMPLOG INFO 2
  ver=v2c public rcVrfRpTrigger.3 = 1
CP1 [04/02/13 14:22:11.620] 0x002c4600 00000000 GlobalRouter SNMPLOG INFO 3
  ver=v2c private rcSysActionL1.0 = 7
CP1 [04/03/13 11:03:35.991] 0x002c4600 00000000 GlobalRouter SNMPLOG INFO 4
  ver=v2c public pingCtlRowStatus.11.111.119.110.101.114.105.110.100.101.120.4
  9.8.116.101.115.116.105.112.118.52 = 4
CP1 [04/03/13 11:03:35.992] 0x002c4600 00000000 GlobalRouter SNMPLOG INFO 5
  ver=v2c public pingCtlTargetAddressType.11.111.119.110.101.114.105.110.100.1
  01.120.49.8.116.101.115.116.105.112.118.52 = 1
CP1 [04/03/13 11:03:35.992] 0x002c4600 00000000 GlobalRouter SNMPLOG INFO 6
  ver=v2c public pingCtlTargetAddress.11.111.119.110.101.114.105.110.100.101.1
  20.49.8.116.101.115.116.105.112.118.52 = 1.1.1.1
CP1 [04/03/13 11:03:35.993] 0x002c4600 00000000 GlobalRouter SNMPLOG INFO 7
  ver=v2c public pingCtlAdminStatus.11.111.119.110.101.114.105.110.100.101.120
  .49.8.116.101.115.116.105.112.118.52 = 2
CP1 [04/03/13 11:03:35.993] 0x002c4600 00000000 GlobalRouter SNMPLOG INFO 8
  ver=v2c public pingCtlDataSize.11.111.119.110.101.114.105.110.100.101.120.49
  .8.116.101.115.116.105.112.118.52 = 16
--More-- (q = quit)
```

Variable definitions

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

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

SNMP trap configuration using EDM

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

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

Configuring an SNMP host target address

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

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit > SnmpV3**.
2. Click **Target Table**.
3. In the **Target Table** tab, click **Insert**.
4. In the **Name** box, type a unique identifier.
5. In the **TDomain** box, select the transport type of the address. Select either **ipv4Tdomain** or **ipv6Tdomain**.
6. In the **TAddress** box, type the transport address and User Datagram Protocol (UDP) port.

7. In the **Timeout** box, type the maximum round trip time.
8. In the **RetryCount** box, type the number of retries to be attempted.
9. In the **TagList** box, type the list of tag values.
10. In the **Params** box, click the ellipsis (...) to select **TparamV1** or **TparamV2**.
11. Click **OK**.
12. In the **TMask** box, type the mask.
13. In the **MMS** box, type the maximum message size.
14. Click **Insert**.

Target Table field descriptions

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

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

Name	Description
	Although the maximum MMS is 2147483647, the device supports the maximum SNMP packet size of 8192.

Configuring target table parameters

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

Procedure

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

Target Params Table field descriptions

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

Name	Description
Name	Identifies the target table.
MPModel	Specifies the message processing model to use to generate messages: SNMPv1, SNMPv2c, or SNMPv3/USM.
SecurityModel	Specifies the security model to use to generate messages: SNMPv1, SNMPv2c, or USM. You can receive an <code>inconsistentValue</code> 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: <code>noAuthNoPriv</code> , <code>authNoPriv</code> , or <code>authPriv</code> .

Configuring an SNMP notify table

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

Procedure

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

Notify Table field descriptions

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

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

Configuring SNMP notify filter profiles

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

Procedure

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

3. Click the **Notify Filter Table** tab.
4. Click **Insert**.
5. In the **NotifyFilterProfileName** box, type a name for the notify filter profile.
6. In the **Subtree** box, type subtree location information in x.x.x.x.x.x.x.x format.
7. In the **Mask** box, type the mask location in hex string format.
8. From the **Type** options, select **included** or **excluded**.
9. Click **Insert**.

Notify Filter Table field descriptions

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

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

Configuring SNMP notify filter profile table parameters

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

Before you begin

- The notify filter profile exists.

Procedure

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

Notify Filter Profile Table field descriptions

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

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

Enabling SNMP trap logging

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

Procedure

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

Error field descriptions

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

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

Viewing the trap sender table

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

Procedure

1. On the Device Physical View, select a chassis.

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

Trap Sender Table field descriptions

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

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

Chapter 9: Remote monitoring

This chapter provides information and procedures about Remote network monitoring (RMON).

You can view all RMON information using ACLI, EDM, or COM. Alternatively, you can use any management application that supports SNMP traps to view RMON trap information.

Remote monitoring fundamentals

Remote network monitoring (RMON) is a management information base (MIB) or a group of management objects that you use to obtain or configure values using the Simple Network Management Protocol (SNMP). Use ACLI, EDM, or COM, you can globally enable RMON on the system. After you globally enable RMON, you then enable monitoring for individual devices on a port-by-port basis.

RMON has four major functions:

- configure alarms for user-defined events
- collect Ethernet statistics
- log events
- send traps for events

Within EDM, you can configure RMON alarms that relate to specific events or variables by selecting these variables from a list. Specify events associated with alarms to trap or log-and-trap. In turn, tripped alarms are trapped or logged.

You can view all RMON information using ACLI, EDM, or COM. Alternatively, you can use any management application that supports SNMP traps to view RMON trap information.

This section describes RMON alarms, RMON history, RMON events, and RMON statistics.

RMON alarms

You can configure alarms to alert you if the value of a variable goes out of range. You can define RMON alarms on any MIB variable that resolves to an integer value. You cannot use string variables (such as system description) as alarm variables.

All alarms share the following characteristics:

- a defined upper and lower threshold value
- a corresponding rising and falling event
- an alarm interval or polling period

The alarm variable is polled and the result is compared against upper and lower limit values selected when the alarm is created. If either limit is reached or crossed during the polling period, then the alarm fires and generates an event that you can view in the event log or the trap log. You can configure the alarm to either create a log, or have the alarm send a Simple Network Management Protocol (SNMP) trap to a Network Management System (NMS). You can view the activity in a log or a trap log, or you can create a script to alert you by beeping at a console, sending an e-mail, or calling a pager.

The upper limit of the alarm is the rising value, and the lower limit is the falling value. RMON periodically samples data based upon the alarm interval. During the first interval that the data passes above the rising value, the alarm fires as a rising event. During the first interval that the data drops below the falling value, the alarm fires as a falling event.

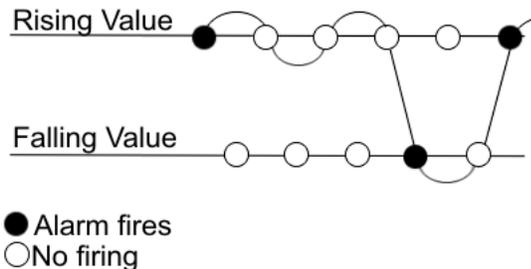


Figure 2: How alarms fire

The alarm fires during the first interval that the sample goes out of range. No additional events generate for that threshold until the opposite threshold is crossed. Therefore, it is important to carefully define the rising and falling threshold values for alarms. Incorrect thresholds cause an alarm to fire at every alarm interval, or never at all.

A general rule is to define one threshold value to an expected, baseline value, and then define the opposite threshold as the out-of-bounds limit. Because of sample averaging, the value is equal to ± 1 baseline unit. For example, assume you define an alarm with octets leaving a port as the variable. The intent of the alarm is to notify you if excessive traffic occurs on that port. You enable spanning tree, and then 52 octets transmit from the port every 2 seconds, which is equivalent to baseline traffic of 260 octets every 10 seconds. This alarm notifies you if the lower limit of exiting octets is defined at 260 and the upper limit is defined at 320 (or at any value greater than $260 + 52 = 312$).

The first time outbound traffic other than spanning tree Bridge Protocol Data Units (BPDUs) occurs, the rising alarm fires. After outbound traffic other than spanning tree ceases, the falling alarm fires. This process provides the time intervals of any nonbaseline outbound traffic.

If you define the alarm with a falling threshold less than 260 (assuming the alarm polling interval is 10 seconds), for example, 250, then the rising alarm can fire only once, as shown in the following example. The falling alarm (the opposite threshold) must fire for the rising alarm to fire a second time. The falling alarm cannot fire unless the port becomes inactive or you disable spanning tree (which causes the value for outbound octets to drop to zero) because the baseline traffic is always greater than the value of the falling threshold. By definition, the failure of the falling alarm to fire prevents the rising alarm from firing a second time.

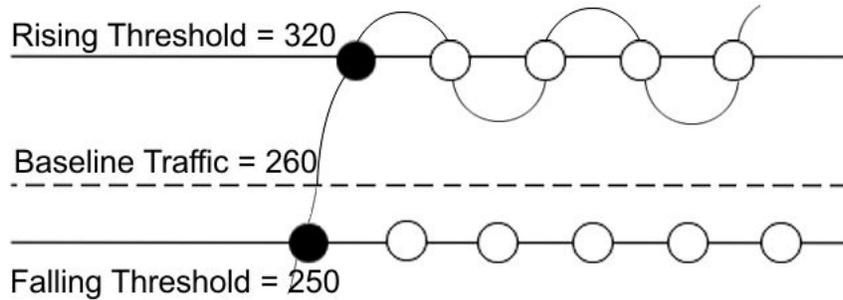


Figure 3: Alarm example, threshold less than 260

When you create an alarm, you select a variable from the variable list and a port, or another system component to which it connects. Some variables require port IDs, card IDs, or other indexes (for example, spanning tree group IDs). You then select a rising and a falling threshold value. The rising and falling values compare to the actual value of the variable that you choose. If the variable falls outside of the rising or falling value range, an alarm triggers and an event is logged or trapped.

When you create an alarm, you also select a sample type, which can be either absolute or delta. Define absolute alarms for alarms based on the cumulative value of the alarm variable. An example of an absolute alarm value is card operating status. Because this value is not cumulative, but instead represents states, such as card up (value 1) and card down (value 2), you configure it as the absolute value. Therefore, you can create an alarm with a rising value of 2 and a falling value of 1 to alert you whether the card is up or down.

Configure most alarm variables related to Ethernet traffic as a delta value. Define delta alarms for alarms based on the difference in the value of the alarm variable between the start of the polling period and the end of the polling period. Delta alarms are sampled twice for each polling period. For each sample, the last two values are added and compared to the threshold values. This process increases precision and detects threshold crossings that span the sampling boundary. Therefore, if you track the current values of a delta-valued alarm and add them, the result is twice the actual value. This result is not an error in the software.

RMON history

The RMON history group records periodic statistical samples from a network. A sample is a history and is gathered in time intervals referred to as buckets. You enable and create histories to establish a time-dependent method to gather RMON statistics on a port. The following are the default values for history:

- Buckets are gathered at 30-minute intervals.
- The number of buckets gathered is 50.

You can configure both the time interval and the number of buckets. However, after the last bucket is reached, bucket 1 is dumped and recycled to hold a new bucket of statistics. Then bucket 2 is dumped, and so forth.

RMON events

RMON events and alarms work together to notify you when values in your network go out of a specified range. After a value passes the specified range, the alarm fires. The event specifies how the activity is recorded.

An event specifies whether a trap, a log, or both a trap and a log generates to view alarm activity. After you globally enable RMON, two default events generate:

- RisingEvent
- FallingEvent

The default events specify that after an alarm goes out of range, both a trap and a log track the firing of the alarm. For example, after an alarm fires at the rising threshold, the rising event specifies to send this information to both an SNMP trap to the NMS, and a log on the Virtual Services Platform 9000. Likewise, after an alarm passes the falling threshold, the falling event specifies to send this information to a trap and a log.

Note:

Before you delete an RMON event, remove all RMON alarms related to the RMON event.

RMON statistics

You can use EDM to gather and graph Ethernet statistics in a variety of formats, or you can save them to a file and export them to a third-party presentation or graphing application.

This implementation of RMON requires a control row for Ethernet statistics. This control row appears as port 0/1 when you choose RMON, Control, Ethernet Statistics. The row ID is reserved for the control row. Therefore, some automated tests, such as ANVL, can fail when the test attempts to create a row 1.

RMON alarm variables

RMON alarm variables are divided into the following three categories:

- Security
- Errors
- Traffic

Each category has subcategories.

The following table lists the alarm variable categories and provides a brief variable description.

Table 4: RMON alarm variables

Category	Subcategory	Variable	Definition
Security		rcCliNumAccessViolations.0	The number of CLI access violations detected by the system.
		rcWebNumAccessBlocks.0	The number of accesses the Web server blocked.
		snmpInBadCommunityNames.0	The total number of SNMP messages delivered to the SNMP protocol entity that represented an SNMP operation not allowed

Category	Subcategory	Variable	Definition
			by the SNMP community named in the message.
Errors	Interface	ifInDiscards	The number of inbound packets discarded even though no errors were detected to prevent the packets being deliverable to a higher-layer protocol. One possible reason for discarding a packet is to free buffer space.
		ifInErrors	For packet-oriented interfaces, the number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol. For character-oriented or fixed-length interfaces, the number of inbound transmission units that contained errors, preventing them from being deliverable to a higher-layer protocol.
		ifOutDiscards	The number of outbound packets discarded even though no errors were detected to prevent the packets being transmitted. One possible reason for discarding such a packet is to free buffer space.
		ifOutErrors	For packet-oriented interfaces, the number of outbound packets that were not transmitted because of errors. For character-oriented or fixed-length interfaces, the number of outbound transmission units that were not transmitted because of errors.
	Ethernet	dot3StatsAlignmentErrors	A count of frames received on a particular interface that are not an integral number of octets in length and do not pass the Frame Check Sequence (FCS) check. The count represented by an instance of this object increments when the alignmentError status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions exist are, according to the conventions of IEEE 802.3

Category	Subcategory	Variable	Definition
			Layer Management, counted exclusively according to the error status presented to the LLC.
		dot3StatsFCSErrors	A count of frames received on a particular interface that are an integral number of octets in length but do not pass the FCS check. The count represented by an instance of this object increments when the frameCheckError status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions occur are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.
		dot3StatsSingleCollisionFrames	A count of successfully transmitted frames on a particular interface where transmission is inhibited by exactly one collision. A frame that is counted by an instance of this object is also counted by the corresponding instance of either the ifOutUcastPkts object, the ifOutMulticastPkts object, or the ifOutBroadcastPkts object, and is not counted by the corresponding instance of the dot3StatsMultipleCollisionFrames object.
		dot3StatsMultipleCollisionFrames	A count of successfully transmitted frames on a particular interface where transmission is inhibited by more than one collision. A frame that is counted by an instance of this object is also counted by the corresponding instance of either the ifOutUcastPkts, ifOutMulticastPkts, or ifOutBroadcastPkts object, and is not counted by the corresponding instance of the

Category	Subcategory	Variable	Definition
			dot3StatsSingleCollisionFrames object.
		dot3StatsSQETestErrors	A count of times that the SQE TEST ERROR message is generated by the PLS sublayer for a particular interface. The SQE TEST ERROR message is defined in section 7.2.2.2.4 of ANSI/IEEE 802.3-1985 and its generation is described in section 7.2.4.6 of the same document.
		dot3StatsDeferredTransmissions	A count of frames for which the first transmission attempt on a particular interface is delayed because the medium is busy. The count represented by an instance of this object does not include frames involved in collisions.
		dot3StatsLateCollisions	The number of times that a collision is detected on a particular interface later than 512 bit-times into the transmission of a packet; 512 bit-times corresponds to 51.2 microseconds on a 10 Mb/s system. A (late) collision included in a count represented by an instance of this object is also considered as a (generic) collision for purposes of other collision-related statistics.
		dot3StatsExcessiveCollisions	A count of frames where the transmission on a particular interface fails due to excessive collisions.
		dot3StatsInternalMacTransmitErrors	A count of frames where the transmission on a particular interface fails due to an internal MAC sublayer transmit error. A frame is only counted by an instance of this object if it is not counted by the corresponding instance of either the dot3StatsLateCollisions object, the dot3StatsExcessiveCollisions object, or the dot3StatsCarrierSenseErrors object.

Category	Subcategory	Variable	Definition
			The precise meaning of the count represented by an instance of this object is implementation specific. In particular, an instance of this object can represent a count of transmission errors on a particular interface that are not otherwise counted.
		dot3StatsCarrierSenseErrors	The number of times the carrier sense condition was lost or never asserted when the switch attempted to transmit a frame on a particular interface. The count represented by an instance of this object increments at most once for each transmission attempt, even if the carrier sense condition fluctuates during a transmission attempt.
		dot3StatsFrameTooLongs	A count of frames received on a particular interface that exceeds the maximum permitted frame size. The count represented by an instance of this object increments when the frameTooLong status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions obtained are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.
		dot3StatsInternalMacReceiveErrors	A count of frames where the transmission on a particular interface fails due to an internal MAC sublayer transmit error. A frame is counted by an instance of this object only if it is not counted by the corresponding instance of either the dot3StatsLateCollisions object, the dot3StatsExcessiveCollisions object, or the dot3StatsCarrierSenseErrors object.

Category	Subcategory	Variable	Definition
			The precise meaning of the count represented by an instance of this object is implementation-specific. In particular, an instance of this object can represent a count of transmission errors on a particular interface that are not otherwise counted.
	IP	ipInHdrErrors.0	The number of input datagrams discarded due to errors in the datagram IP headers, including bad checksums, version number mismatch, other format errors, time-to-live exceeded, and errors discovered in processing IP options.
		ipInDiscards.0	The number of discarded input IP datagrams where no problems were encountered to prevent continued processing. An example of why they were discarded can be lack of buffer space. This counter does not include any datagrams discarded while awaiting reassembly.
		ipOutDiscards.0	The number of output IP datagrams where no problems were encountered to prevent transmission to the destination, but that were discarded (for example, for lack of buffer space). This counter includes datagrams counted in ipForwDatagrams if packets meet this (discretionary) discard criterion.
		ipFragFails.0	The number of IP datagrams discarded because they needed to be fragmented at this entity but were not, for example, because the Don't Fragment flag was set.
		ipReasmFails.0	The number of failures detected by the IP reassembly algorithm (for whatever reason: timed out, errors, and so forth). This is not necessarily a count of discarded IP fragments because some algorithms (notably the algorithm in RFC 815) can lose track of the

Category	Subcategory	Variable	Definition
			number of fragments by combining them as they are received.
		icmpInParmProbs.0	The number of ICMP In parameter problem messages received.
		icmpOutParmProbs.0	The number of ICMP Out parameter problem messages received.
	MLT	rcStatMltEtherAlignmentErrors	The number of frames received on an MLT that are not an integral number of octets in length, but do not pass the FCS check.
		rcStatMltEtherFCSErrors	The number of frames received on an MLT that are an integral number of octets in length, but do not pass the FCS check.
		rcStatMltEtherSingleCollFrames	The number of successfully transmitted frames on a particular MLT where transmission is inhibited by exactly one collision.
		rcStatMltEtherMultipleCollFrames	The number of successfully transmitted frames on a particular MLT where transmission is inhibited by more than one collision.
		rcStatMltEtherSQETestError	A count of times that the SQE TEST ERROR message is generated by the PLS sublayer for a particular MLT.
		rcStatMltEtherDeferredTransmiss	A count of frames where the first transmission attempt on a particular MLT is delayed because the medium is busy. The count represented by an instance of this object.
		rcStatMltEtherLateCollisions	The number of times that a late collision is detected on a particular MLT later than 512 bit-times into the transmission of a packet; 512-bit-times corresponds to 51.2-microseconds on a 10 Mb/s system.
		rcStatMltEtherExcessiveCollis	The number of times that excessive collisions are detected on a particular MLT later than 512

Category	Subcategory	Variable	Definition
			bit-times into the transmission of a packet; 512 bit-times corresponds to 51.2 microseconds on a 10-Mb/s system.
		rcStatMltEtherMacTransmitError	A count of frames where the transmission on a particular MLT fails due to an internal MAC sublayer transmit error. A frame is only counted by an instance of this object if it is not counted by the corresponding instance of either the LateCollisions object, the ExcessiveCollisions object, or the CarrierSenseErrors object.
		rcStatMltEtherCarrierSenseError	The number of times the carrier sense condition was lost or never asserted when attempting to transmit a frame on a particular MLT. The count represented by an instance of this object increments at most once for each transmission attempt, even if the carrier sense condition fluctuates during a transmission attempt.
		rcStatMltEtherFrameTooLong	A count of frames received on a particular MLT that exceeds the maximum permitted frame size. The count represented by an instance of this object increments when the frameTooLong status is returned by the MAC service to the LLC (or other MAC user).
		rcStatMltEtherMacReceiveError	A count of frames for which reception on a particular MLT fails due to an internal MAC sublayer receive error. A frame is only counted by an instance of this object if it is not counted by the corresponding instance of either the FrameTooLongs object, the AlignmentErrors object, or the FCSErrors object.
	Other	rcTblArNoSpace	The number of entries not added to the address translation table due to lack of space.

Category	Subcategory	Variable	Definition
		snmpInAsnParseErrs.0	The total number of ASN.1 or BER errors encountered by the SNMP protocol entity when it decodes received SNMP messages.
		rcStgPortInBadBpdus	The number of bad BPDUs received by this port.
		dot1dTpPortInDiscards	Count of valid frames received that were discarded (that is, filtered) by the forwarding process.
		rip2IfStatRcvBadPackets	The number of routes in valid RIP packets that were ignored for any reason.
		rip2IfStatRcvBadRoutes	The number of RIP response packets received by the RIP process that were subsequently discarded for any reason.
		rcStatOspfBufferAllocFailures.0	The number of times that OSPF failed to allocate buffers.
		rcStatOspfBufferFreeFailures.0	The number of times that OSPF failed to free buffers.
Traffic	Interface	ifInOctets	The total number of octets received on the interface, including framing characters.
		ifInMulticastPkts	The number of packets, delivered by this sublayer to a higher sublayer, that are addressed to a multicast address at this sublayer. For a MAC layer protocol, this number includes both Group and Functional addresses.
		ifInBroadcastPkts	The number of packets, delivered by this sublayer to a higher (sub) layer, that are addressed to a broadcast address at this sublayer.
		ifInUnkownProtos	For packet-oriented interfaces, the number of packets received through the interface that are discarded because of an unknown or unsupported protocol. For character-oriented or fixed-length interfaces that support protocol multiplexing, the number of transmission units received

Category	Subcategory	Variable	Definition
			through the interface that are discarded because of an unknown or unsupported protocol. For any interface that does not support protocol multiplexing, this counter is always 0.
		ifOutOctets	The total number of octets transmitted from the interface, including framing characters.
		ifOutMulticastPkts	The total number of packets that higher-level protocols requested be transmitted, and that are addressed to a multicast address at this sublayer, including those that are discarded or not sent. For a MAC layer protocol, this includes both Group and Functional addresses.
		ifoutBroadcastPkts	The total number of packets that higher level protocols requested transmitted, and that were addressed to a broadcast address at this sublayer, including those discarded or not sent.
		ifLastChange	The value of sysUpTime at the time the interface entered its current operational state. If the current state was entered prior to the last reinitialization of the local network management subsystem, this object contains a value of zero.
	RmonEther Stats	etherStatsOctets	The total number of octets of data (including those in bad packets) received on the network (excluding framing bits but including FCS octets). Use this object as a reasonable estimate of Ethernet utilization. For greater precision, sample the etherStatsPkts and etherStatsOctets objects before and after a common interval.
		etherStatsPkts	The total number of packets (including bad packets, broadcast packets, and multicast packets) received.

Category	Subcategory	Variable	Definition
		etherStatsBroadcastPkts	The total number of good packets received that are directed to the broadcast address. This number does not include multicast packets.
		etherStatsMulticastPkts	The total number of good packets received that are directed to a multicast address. This number does not include packets directed to the broadcast address.
		etherStatsCRCAAlignErrors	The total number of packets received that had a length (excluding framing bits, but including FCS octets) of 64 to 1518 octets, inclusive, but had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).
		etherStatsUndersizePkts	The total number of packets received that are less than 64 octets long (excluding framing bits, but including FCS octets) and were otherwise well formed.
		etherStatsOversizePkts	The total number of packets received that are longer than 1518 octets (excluding framing bits, but including FCS octets) and were otherwise well formed.
		etherStatsFragments	<p>The total number of packets received that are less than 64 octets in length (excluding framing bits but including FCS octets) and had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).</p> <p>It is entirely normal for etherStatsFragments to increment because it counts both runts (which are normal occurrences due to collisions) and noise hits.</p>

Category	Subcategory	Variable	Definition
		etherStatsCollisions	The best estimate of the total number of collisions on this Ethernet segment.
	IP	ipInReceives.0	All incoming IP packets.
		ipInAddrErrors.0	The number of bad IP destination addresses.
		ipForwDatagrams.0	IP packets forwarded.
		ipInUnknownProtos.0	Number of unsupported IP protocols.
		ipInDelivers.0	The number of IP In packets delivered.
		ipOutRequests.0	The total number of IP datagrams that local IP user protocols supplied to IP in request for transmission.
		ipOutNoRoutes.0	The number of IP datagrams discarded because no route was found to transmit to the destination.
		ipFragOKs.0	The number of IP datagrams successfully fragmented.
		ipFragCreates.0	The number of IP datagram fragments generated as a result of fragmentation.
		ipReasmReqds.0	The number of requests to reassemble fragments.
		ipReasmOKs.0	The number of fragments reassembled successfully.
	ICMP	icmpInSrcQuenchs.0	The number of ICMP Source Quench messages received.
		icmpInRedirects.0	The number of ICMP redirect messages.
		icmpInEchos.0	The number of ICMP Echo requests messages received.
		icmpInEchosReps.0	The number of ICMP Echo reply messages received.
		icmpInTimeStamps.0	The number of ICMP timestamp request messages received.
		icmpInTimeStampsReps.0	The number of ICMP timestamp reply messages received.
		icmpInAddrMasks.0	The number of ICMP mask request messages reviewed.

Category	Subcategory	Variable	Definition
		icmpInAddrMasksReps.0	The number of ICMP mask reply messages reviewed.
		icmpInDestUnreachs.0	The number of ICMP destinations unreachable messages received.
		icmpInTimeExcds.0	The number of ICMP Time Exceeded messages received.
		icmpOutSrcQuenchs.0	The number of ICMP Source Quench messages sent.
		icmpOutRedirects.0	The number of ICMP redirect messages sent.
		icmpOutEchos.0	The number of ICMP Echo request messages sent.
		icmpOutEchosReps.0	The number of ICMP Echo reply messages sent.
		icmpOutTimeStamps.0	The number of ICMP Timestamp request messages sent.
		icmpOutTimeStampsReps.0	The number of ICMP Timestamp reply messages sent.
		icmpOutAddrMasks.0	The number of ICMP Address mask messages sent.
		icmpOutAddrMasksReps.0	The number of ICMP Address mask reply messages sent.
		icmpOutDestUnreachs.0	The number of ICMP destination unreachable messages sent.
		icmpOutTimeExcds.0	The number of ICMP time exceeded messages sent.
	Snmp	snmpInPkts.0	The total number of messages delivered to the SNMP entity from the transport service.
		snmpOutPkts.0	The total number of SNMP messages passed from the SNMP protocol entity to the transport service.
		snmpInBadVersions.0	The total number of SNMP messages delivered to the SNMP protocol entity that were intended for an unsupported SNMP version.
		snmpInBadCommunityUses.0	The total number of SNMP messages delivered to the SNMP protocol entity that represented an SNMP operation that was not

Category	Subcategory	Variable	Definition
			allowed by the SNMP community named in the message.
		snmpInTooBigs.0	The total number of SNMP PDUs delivered to the SNMP protocol entity and for which the value of the error-status field is tooBig.
		snmpInNoSuchNames.0	The total number of SNMP PDUs delivered to the SNMP protocol entity and for which the value of the error-status field is noSuchName.
		snmpInBadValues. 0	The total number of SNMP PDUs received that were generated by the SNMP protocol entity and for which the value of the error-status field is badValue.
		snmpInReadOnlys.0	The total number of valid SNMP PDUs delivered to the SNMP protocol entity and for which the value of the error-status field is readOnly. It is a protocol error to generate an SNMP PDU that contains the value readOnly in the error-status field; as such, this object is provided as a means of detecting incorrect implementations of the SNMP.
		snmpInGenErrs.0	The total number of SNMP PDUs delivered to the SNMP protocol entity and for which the value of the error-status field is genErr.
		snmpInTotalReqVars.0	The total number of MIB objects retrieved successfully by the SNMP protocol entity as the result of receiving valid SNMP Get-Request and Get-Next PDUs.
		snmpInTotalSetVars.0	The total number of MIB objects altered successfully by the SNMP protocol entity as the result of receiving valid SNMP Set-Request PDUs.
		snmpInGetRequests.0	The total number of SNMP Get-Request PDUs accepted and processed by the SNMP protocol entity.

Category	Subcategory	Variable	Definition
		snmpInGetNexts.0	The total number of SNMP Get-Next PDUs accepted and processed by the SNMP protocol entity.
		snmpInSetRequests.0	The total number of SNMP Set-Request PDUs accepted and processed by the SNMP protocol entity.
		snmpInGetResponses.0	The total number of SNMP Get-Response PDUs accepted and processed by the SNMP protocol entity.
		snmpInTraps.0	The total number of SNMP Trap PDUs accepted and processed by the SNMP protocol entity.
		snmpOutTooBig.0	The total number of SNMP PDUs generated by the SNMP protocol entity and for which the value of the error-status field is tooBig.
		snmpOutNoSuchNames.0	The total number of SNMP PDUs generated by the SNMP protocol entity and for which the value of the error-status field is noSuchName.
		snmpOutBadValues.0	The total number of SNMP PDUs sent that were generated by the SNMP protocol entity and for which the value of the error-status field is badValue.
		snmpOutGenErrs.0	The total number of SNMP PDUs generated by the SNMP protocol entity and for which the value of the error-status field is genErr.
		snmpOutGetRequests.0	The total number of SNMP Get-Request PDUs generated by the SNMP protocol entity.
		snmpOutGetNexts.0	The total number of SNMP Get-Next PDUs generated by the SNMP protocol entity.
		snmpOutSetRequests.0	The total number of SNMP Set-Request PDUs generated by the SNMP protocol entity.
		snmpOutGetResponses.0	The total number of SNMP Get-Response PDUs generated by the SNMP protocol entity.

Category	Subcategory	Variable	Definition
		snmpOutTraps.0	The total number of SNMP Trap PDUs generated by the SNMP protocol entity.
	Bridge	rcStgTimeSinceTopologyChange	The time (in hundredths of a second) since the last topology change was detected by the bridge entity.
		rcStgTopChanges	The total number of topology changes detected by this bridge since the management entity was last reset or initialized.
		rcStgMaxAge	The maximum age of Spanning Tree Protocol information learned from the network on any port before it is discarded, in hundredths of a second. This is the actual value that this bridge is currently using.
		rcStgPortForwardTransitions	The number of times this port transitioned from the Learning state to the Forwarding state.
		rcStgPortInConfigBpdus	The number of Config BPDUs received by this port.
		rcStgPortInTcnBpdus	The number of Topology Change Notification BPDUs received by this port.
		rcStgPortOutConfigBpdus	The number of Config BPDUs transmitted by this port.
		rcStgPortOutTcnBpdus	The number of Topology Change Notification BPDUs transmitted by this port.
		dot1dTpPortInFrames	The number of frames received by this port from its segment. A frame received on the interface corresponding to this port is counted by this object only if it is for a protocol being processed by the local bridging function, including bridge management frames.
		dot1dTpPortOutFrames	The number of frames transmitted by this port to its segment. A frame transmitted on the interface corresponding to this port is counted by this object if and only if it is for a protocol processed by

Category	Subcategory	Variable	Definition
			the local bridging function, including bridge management frames.
		dot1dTpLearnedEntryDiscards.0	The total number of Forwarding Database entries learned but discarded due to a lack of space to store them in the Forwarding Database. If this counter increases, it indicates that the forwarding database is regularly becoming full (a condition that has negative performance effects on the subnetwork). If this counter has a significant value but does not increase, it indicates that the problem occurred but is not persistent.
	Utilization	rcSysCpuUtil.0	Percentage of SF/CPU utilization.
		rcSysSwitchFabricUtil.0	Percentage of switching fabric utilization.
		rcSysBufferUtil.0	Buffer utilization as a percentage of the total amount of buffer space in the system. A high value indicates congestion.
		rcSysNVRamUsed.0	Nonvolatile RAM (NVRAM) in use in kilobytes.
		rcSysLastChange.0	Last management-initiated configuration change since sysUpTime.
		rcSysLastVlanChange.0	Last management-initiated VLAN configuration change since sysUpTime.
		rcSysLastSaveToNVRam.0	SysUpTime of the last time the NVRAM on the SF/CPU board was written to.
		rcSysLastSaveToStandbyNVRam.0	SysUpTime of the last time the standby NVRAM (on the backup SF/CPU board) was written to.
	RIP	rip2GlobalRoute Changes.0	The number of changes made to the IP Route database by RIP.
		rip2GlobalQueries.0	The number of responses sent to RIP queries from other systems.
		rip2IfStatSentUpdates	The number of triggered RIP updates actually sent on this interface.

Category	Subcategory	Variable	Definition
	OSPF	ospfExternLSACount.0	The number of external (LSA type 5) link-state advertisements in the link-state database.
		ospfOriginateNewLSAs.0	The number of new link-state advertisements that have originated. The number increments each time the router originates a new LSA.
		ospfrxNewLSAs.0	The number of link-state advertisements received determined to be new installations.
		ospfSpfRuns	Indicates the number of SPF calculations performed by OSPF.
		ospfAreaBdrRtrCount	The total number of area border routers reachable within this area.
		ospfASBdrRtrCount	The total number of autonomous system border routers reachable within this area.
		ospfAreaLSACount	The total number of link-state advertisements in this area's link state database.
		ospflfState	This signifies a change in the state of an OSPF virtual interface.
		ospflfEvents	The number of times this OSPF interface changed the state or an error occurred.
		ospfVirtLfState	The number of times this OSPF interface.
		ospfVirtLfEvents	The number of state changes or error events on this virtual link.
		ospfVirtNbrState	The state of the Virtual Neighbor Relationship.
		ospfVirtNbrEvents	The number of times this virtual link changed the state or an error occurred.
	Igmp	igmpInterfaceWrongVersions	The number of queries received whose IGMP version does not match. IGMP requires that all routers on the LAN are configured to run the same version of IGMP.
		igmpInterfaceJoins	The number of times a group membership was added on this interface.

Category	Subcategory	Variable	Definition
		igmpInterfaceLeaves	The number of times a group membership was deleted on this interface.
	MLT	rcStatMltIfExtnIfInMulticastPkts	The total number of multicast packets delivered to this MLT interface.
		rcStatMltIfExtnIfInBroadcastPkts	The total number of broadcast packets delivered to this MLT Interface.
		rcStatMltIfExtnIfOutMulticastPkts	The total number of MLT interface multicast packets delivered to this MLT interface.
		rcStatMltIfExtnIfOutBroadcastPkts	The total number of MLT interface broadcast packets delivered to this MLT interface.
		rcStatMltIfExtnIfHCInOctets	The total number of octets received on this MLT interface including framing characters detected by the high-count (64-bit) register.
		rcStatMltIfExtnIfHCInUcastPkts	The number of packets delivered by this MLT interface to a higher MLT that were not addressed to a multicast or broadcast address as detected by the high-count (64-bit) register.
		rcStatMltIfExtnIfHCInMulticastPkt	The total number of multicast packets delivered to this MLT interface detected by the high-count (64-bit) register.
		rcStatMltIfExtnIfHCInBroadcastPkt	The total number of broadcast packets delivered to this MLT interface detected by the high-count (64-bit) register.
		rcStatMltIfExtnIfHCOutOctets	The total number of octets transmitted from the MLT interface, including framing characters.
		rcStatMltIfExtnIfHCOutUcastPkts	The number of packets transmitted by this MLT interface to a higher MLT that were not addressed to a multicast or broadcast address as detected by the high-count (64-bit) register.

Category	Subcategory	Variable	Definition
		rcStatMltIfExtnIfHCOuMulticast	The total number of packets that higher-level protocols requested be transmitted, and that were addressed to a multicast address at this sublayer, including those that were discarded or not sent registered by the high-count (64-bit) register.
		rcStatMltIfExtnIfHCOuBroadcast	The total number of packets that higher-level protocols requested be transmitted, and that were addressed to a broadcast address at this sublayer, including those that were discarded or not sent registered by the high-count (64-bit) register.

RMON configuration using ACLI

Remote monitoring (RMON) is a management information base (MIB) or a group of management objects that you use to obtain or configure values using the Simple Network Management Protocol (SNMP).

Configuring RMON alarms and events

Configure RMON functions on Virtual Services Platform 9000 to set alarms and capture events.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enable RMON globally:

```
rmon
```

3. Configure an RMON alarm:

```
rmon alarm <1-65535> WORD <1-1536> <1-3600> {absolute|delta}
[falling-threshold <-2147483647-2147483647> event <1-65535>] [owner
WORD<1-127>] [rising-threshold <-2147483647-2147483647> event
<1-65535>]
```

4. Configure an RMON event:

```
rmon event <1-65535> [community WORD<1-127>] [description
WORD<1-127>] [log] [owner WORD<1-127>] [trap] [trap_dest
[{A.B.C.D}]] [trap_src [{A.B.C.D}]]
```

Before you delete an RMON event, remove all RMON alarms related to the RMON event.

Example

```
VSP-9012:1(config)#rmon
VSP-9012:1(config)#rmon alarm 4 rcCliNumAccessViolations.0 10 absolute rising-threshold 2
event 6000
VSP-9012:1(config)#rmon event 60534 community public description "Rising Event" log trap
```

Variable definitions

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

Variable	Value
alarm <1-65535> WORD <1-1536> <1-3600> {absolute delta} [falling-threshold <-2147483647-2147483647> event <1-65535>] [owner WORD<1-127>] [rising-threshold <-2147483647-2147483647> event <1-65535>]	<p>Create an alarm interface.</p> <ul style="list-style-type: none"> <1-65535> is the interface index number from 1–65535. WORD <1-1536> is the variable name or OID, case sensitive (string length 1–1536). {absolute delta} is the sample type. rising-threshold <-2147483648-2147483647> [<event: 1-65535>] is the rising threshold (–2147483648–2147483647) and the rising event number (1–65535). falling-threshold <-2147483648-2147483647> [<event: 1-65535>] is the falling threshold (–2147483648–2147483647) and the falling event number (1–65535). owner WORD<1-127> is the name of the owner (string length 1–127). The default value is CLI. <p>Use the default operator to reset the RMON alarms to their default configuration: <code>default rmon alarm <65535></code></p> <p>Use the no operator to disable RMON alarms: <code>no rmon alarm [<1-65535>]</code></p>
event <1-65535> [community WORD<1-127>] [description WORD<1-127>] [log] [owner WORD<1-127>] [trap] [trap_dest [{A.B.C.D}]] [trap_src [{A.B.C.D}]]	<p>Create an event.</p> <ul style="list-style-type: none"> <1-65535> is the event index number. [log] display information about configured traps. [trap] specify trap source and destination IP addresses. description WORD<1-127> is the event description (string length 0–127). owner WORD<1-127> is the name of the owner (string length 1–127). The default value is CLI. trap_src {A.B.C.D} is the trap source ip address. trap_dest {A.B.C.D} is the trap destination ip address.

Variable	Value
	<ul style="list-style-type: none"> community <i>WORD</i><1-127> is the event community (string length 1–127). <p>Use the no operator to delete a RMON event: <code>no rmon event [<1-65535>] [log]</code></p> <p>Note:</p> <p>Before you delete an RMON event, remove all RMON alarms related to the RMON event.</p>

Viewing RMON settings

View RMON settings to see information about alarms, statistics, events, or the status of RMON on Virtual Services Platform 9000.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View RMON settings:

```
show rmon {alarm|event|history|log|stats}
```

Example

```
CB-SWB:1(config)#show rmon event
```

```
=====
                                Rmon Event
=====
INDEX DESCRIPTION          TYPE          COMMUNITY OWNER          LAST_TIME_SENT
-----
60534 Rising Event        log-and-trap public          47.17.142.155 none
60535 Falling Event      log-and-trap public          47.17.142.155 8 day(s), 19:14:32
```

```
CB-SWB:1(config)#show rmon log
```

```
=====
                                Rmon Log
=====
INDEX    TIME                DESCRIPTION
-----
60535. 1 8 day(s), 19:14:45 1.3.6.1.4.1.2272.1.19.14.0 (absValue = 0, Falling
Threshold = 2, interval = 10)[alarmIndex.1][trap]
"Falling Event"
60535. 2 8 day(s), 19:14:45 1.3.6.1.4.1.2272.1.19.14.0 (absValue = 0, Falling
Threshold = 1, interval = 10)[alarmIndex.2][trap]
"Falling Event"
```

```
VSP-9012:1(config)#show rmon stats
```

```
=====
                                Rmon Ether Stats
=====
```

```

INDEX  PORT  OWNER
-----
1      cpp   monitor

```

Variable definitions

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

Variable	Value
alarm	Display the RMON Alarm table.
event	Display the RMON event table.
history	Display the RMON history table.
log	Display the RMON log table.
stats	Display the RMON statistics table.

RMON configuration using EDM

Remote monitoring (RMON) is a management information base (MIB) or a group of management objects that you use to obtain or configure values using the Simple Network Management Protocol (SNMP).

Enabling RMON globally

You must globally enable RMON before you can use an RMON function. If you attempt to enable an RMON function before the global flag is disabled, EDM informs you that the flag is disabled and prompts you to enable the flag.

About this task

If you want to use nondefault RMON parameter values, you can configure them before you enable RMON, or as you configure the RMON functions.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > Serviceability > RMON**.
2. Click **Options**.
3. Click the **Options** tab.
4. Select the **Enable** check box.
5. In the **UtilizationMethod** option, select a utilization method.
6. In the **TrapOption** option, select a trap option.
7. In the **MemSize** box, type a memory size.

- Click **Apply**.

Options field descriptions

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

Name	Description
Enable	Enables RMON. If you select the Enable check box, the RMON agent starts immediately if the amount of memory specified by MemSize is currently available in the device. To disable RMON, clear the Enable check box and click Apply to save the new setting to NVRAM, and then restart the device. The default is disabled.
UtilizationMethod	Controls whether RMON uses a half-duplex or full-duplex formula to calculate port usage. After you select halfDuplex, RMON uses InOctets and the speed of the port to calculate port usage (this is the standard RMON RFC1271 convention). After you select fullDuplex, RMON uses InOctets and OutOctets and 2X the speed of the port to calculate port usage. If you select fullDuplex, but the port operates in half-duplex mode, the calculation defaults to the RFC1271 convention. The default is halfDuplex.
TrapOption	Indicates whether the system sends RMON traps to the owner of the RMON alarm (the manager that created the alarm entry) or to all trap recipients in the system trap receiver table. The default value is toOwner.
MemSize	Specifies the RAM size, in bytes, available for RMON to use. The default value is 250 Kilobytes.

Enabling RMON history

Use RMON to establish a history for a port and configure the bucket interval. For example, to gather RMON statistics over the weekend, you must have enough buckets to cover two days. Configure the history to gather one bucket every hour, and cover a 48 hour period.

About this task

After you configure history characteristics, you cannot modify them; you must delete the history and create another one.

Procedure

- In the navigation pane, expand the following folders: **Configuration > Serviceability > RMON**.
- Click **Control**.
- In the **History** tab, click **Insert**.
- In the **Port** box, click the ellipsis (...) button.
- Select a port.
- Click **OK**.

7. In the **Buckets Requested** box, type the number of discrete time intervals to save data.
8. In the **Interval** box, type the interval in seconds.
9. In the **Owner** box, type the owner information.
10. Click **Insert**.

History field descriptions

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

Name	Description
Index	Specifies an index that uniquely identifies an entry in the historyControl table. Each entry defines a set of samples at a particular interval for an interface on the device. Index value ranges from 1–65535. The default value is 1.
Port	Identifies the source for which historical data is collected and placed in a media-specific table on behalf of this historyControlEntry. The source is an interface on this device. To identify a particular interface, the object identifies the instance of the ifIndex object, defined in (4,6), for the desired interface. For example, if an entry receives data from interface 1, the object is ifIndex 1. The statistics in this group reflect all packets on the local network segment attached to the identified interface. You cannot modify this object if the associated historyControlStatus object is equal to valid(1).
BucketsRequested	Specifies the requested number of discrete time intervals over which data is saved in the part of the media-specific table associated with this historyControlEntry. After this object is created or modified, the probe configures historyControlBucketsGranted as closely to this object as possible for the particular probe implementation and available resources. Values range from 1–65535. The default value is 50.
BucketsGranted	Specifies the number of discrete sampling intervals over which data is saved in the part of the media-specific table associated with this historyControlEntry. After the associated BucketsRequested object is created or modified, the probe sets this object as closely to the requested value as possible for the particular probe implementation and available resources. The probe must not lower this value except as a result of a modification to the associated BucketsRequested object. Occasionally, the actual number of buckets associated with this entry is less than the value of this object. In this case, at the end of each sampling interval, a new bucket is added to the media-specific table. After the number of buckets reaches the value of this object and a new bucket is to be added to the media-specific table, the oldest bucket associated with this entry is deleted by the agent so that the new bucket can be added. After the value of this object changes to a value less than the current value, entries are deleted from the media-specific table associated with this entry. The agent deletes the oldest of these entries so that their number remains less than or equal to the new value of this object. After the value of this object changes to a value greater than the

Name	Description
	current value, the number of associated media-specific entries is allowed to grow.
Interval	Specifies the interval in seconds over which the data is sampled for each bucket in the part of the media-specific table associated with this historyControlEntry. You can set this interval between 1–3600 seconds (1 hour). Because the counters in a bucket can overflow at their maximum value with no indication, you must take into account the possibility of overflow in all of the associated counters. Consider the minimum time in which a counter can overflow on a particular media type, and then set the historyControlInterval object to a value less than this interval, which is typically most important for the octets counter in a media-specific table. For example, on an Ethernet network, the etherHistoryOctets counter can overflow in approximately 1 hour at the maximum utilization. You cannot modify this object if the associated historyControlStatus object is equal to valid. The default value is 1800.
Owner	Specifies the entity that configured this entry and is using the assigned resources. The default value is SNMP.

Disabling RMON history

Disable RMON history on a port if you do not want to record a statistical sample from that port.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > Serviceability > RMON**.
2. Click **Control**.
3. In the **History** tab, select the row that contains the port ID to delete.
4. Click **Delete**.

Viewing RMON history statistics

View RMON history statistics when you want to see a statistical sample from the switch. You can create a graph of the statistics in a bar, pie, chart, or line format.

Procedure

1. On the Device Physical View, select a port.
2. In the navigation pane, expand the following folders: **Configuration > Graph**
3. Click **Port**.
4. Click the **RMON History** tab.
5. Select the statistics you want to graph.

6. Click the button for the type of graph you require (bar, pie, chart, or line).

RMON History field descriptions

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

Table 5: Variable definitions

Parameter	Description
SampleIndex	Identifies the particular sample this entry represents among all samples associated with the same history control entry. This index starts at one and increases by one as each new sample is taken.
Utilization	The best estimate of the mean physical layer network utilization on this interface during the sampling interval, in hundredths of a percent.
Octets	The total number of octets of data (including those in bad packets) received on the network (excluding framing bits but including FCS octets)
Pkts	The number of packets (including bad packets) received during this sampling interval.
BroadcastPkts	The number of good packets received during this sampling interval that were directed to the broadcast address.
MulticastPkts	The number of good packets received during this sampling interval that were directed to a multicast address. This number does not include packets addressed to the broadcast address.
DropEvents	The total number of events in which packets were dropped by the probe due to lack of resources during this sampling interval. This number is not necessarily the number of packets dropped; it is only the number of times this condition was detected.
CRCAAlignErrors	The number of packets received during this sampling interval that had a length (excluding framing bits but including FCS octets) from 64–1518 octets, inclusive, but had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).
UndersizePkts	The number of packets received during this sampling interval that were less than 64 octets (excluding framing bits but including FCS octets) and were otherwise well formed.
OversizePkts	The number of packets received during this sampling interval that were longer than 1518 octets (excluding framing bits but including FCS octets) but were otherwise well formed.
Fragments	The total number of packets received during this sampling interval that were less than 64 octets in length (excluding framing bits but including FCS octets) and had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). It is entirely normal for Fragments to increment because it counts both runts (which are normal occurrences due to collisions) and noise hits.

Parameter	Description
Collisions	<p>The best estimate of the total number of collisions on this Ethernet segment during this sampling interval. The value returned depends on the location of the RMON probe. Section 8.2.1.3 (10Base-5) and section 10.3.1.3 (10Base-2) of IEEE standard 802.3 states that a station must detect a collision in the receive mode if three or more stations transmit simultaneously. A repeater port must detect a collision when two or more stations transmit simultaneously. Thus, a probe placed on a repeater port can record more collisions than a probe connected to a station on the same segment.</p> <p>Probe location plays a small role when 10Base-T. 14.2.1.4 (10Base-T) of IEEE standard 802.3 defines a collision as the simultaneous presence of signals on the DO and RD circuits (transmitting and receiving at the same time). A 10Base-T station can detect only collisions when it transmits. Thus, probes placed on a station and a repeater can report the same number of collisions.</p> <p>A RMON probe inside a repeater can ideally report collisions between the repeater and one or more other hosts (transmit collisions as defined by IEEE 802.3k) plus receiver collisions observed on any coax segments to which the repeater is connected.</p>

Viewing the RMON log

About this task

View the trap log to see which activity occurred.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Serviceability > RMON**.
2. Click **Alarms**.
3. Click the **Log** tab.

Log field descriptions

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

Name	Description
Time	Specifies the creation time for this log entry.
Description	Specifies an implementation dependent description of the event that activated this log entry.

Creating an alarm

After you enable RMON globally, you also create a default rising and falling event. The default for the events is log-and-trap, which means that you receive notification through a trap as well as through a log file.

Before you begin

- You must globally enable RMON.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > Serviceability > RMON**.
2. Click **Alarms**.
3. Click the **Alarms** tab.
4. Click **Insert**.
5. In the **Variable** option, select a variable for the alarm.

If you select some variables, the system will prompt you for a port (or other object) on which you want to set an alarm.

6. In the **SampleType** option, select a sample type.
7. In the **Interval** box, type a sample interval in seconds.
8. In the **Index** box, type an index number.
9. In the **RisingThreshold** box, type a rising threshold value.
10. In the **RisingEventIndex** box, type a rising threshold event index.
11. In the **FallingThreshold** box, type a falling threshold value.
12. In the **FallingEventIndex** box, type a falling threshold event index.
13. In the **Owner** box, type the owner of the alarm.
14. Click **Insert**.

Alarms field descriptions

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

Name	Description
Index	Uniquely identifies an entry in the alarm table. Each entry defines a diagnostic sample at a particular interval for an object on the device. The default is 1.
Interval	Specifies the interval, in seconds, over which the data is sampled and compared with the rising and falling thresholds. deltaValue sampling—configure the interval short enough that the sampled variable is unlikely to increase or decrease by more than $2^{31}-1$ during a single sampling interval. The default is 10.

Name	Description
Variable	<p>Specifies the object identifier of the particular variable to be sampled. Only variables that resolve to an ASN.1 primitive type of INTEGER (INTEGER, Counter, Gauge, or TimeTicks) can be sampled.</p> <p>Alarm variables exist in three formats, depending on the type:</p> <ul style="list-style-type: none"> • A chassis, power supply, or fan-related alarm ends in x where the x index is hard-coded. No further information is required. • A card, spanning tree group (STG), Routing Information Protocol (RIP), Open Shortest Path First (OSPF), or EtherStat alarm ends with a dot (.). You must enter a card number, STG ID, IP address, or EtherStat information. • A port alarm ends with no dot or index and requires that you use the port shortcut menu. An example of a port alarm is ifInOctets (interface incoming octet count). <p>Because SNMP access control is articulated entirely in terms of the contents of MIB views, no access control mechanism exists to restrict the value of this object to identify only those objects that exist in a particular MIB view. Because there is no acceptable means of restricting the read access that is obtained through the alarm mechanism exists, the probe must grant only write access to this object in those views that have read access to all objects on the probe.</p> <p>After you configure a variable, if the supplied variable name is not available in the selected MIB view, a badValue error will be returned. After the variable name of an established alarmEntry is no longer available in the selected MIB view, the probe will change the status of this alarmEntry to invalid.</p> <p>You cannot modify this object if the associated alarmStatus object is equal to valid.</p>
SampleType	<p>Specifies the method of sampling the selected variable and calculating the value to be compared against the thresholds. If the value of this object is absoluteValue, the value of the selected variable is compared directly with the thresholds at the end of the sampling interval. If the value of this object is deltaValue, the value of the selected variable at the last sample is subtracted from the current value, and the difference compared with the thresholds. You cannot modify this object if the associated alarmStatus object is equal to valid. The default is deltaValue.</p>
Value	<p>Specifies the value of the statistic during the last sampling period. For example, if the sample type is deltaValue, this value is the difference between the samples at the beginning and end of the period. If the sample type is absoluteValue, this value is the sampled value at the end of the period. This value is compared with the rising and falling thresholds. The value during the current sampling period is not made available until the period is completed and remains available until the next period is complete.</p>
StartupAlarm	<p>Specifies the alarm that is sent after this entry is first set to valid. If the first sample after this entry becomes valid is greater than or equal to the risingThreshold and alarmStartupAlarm is equal to the risingAlarm or the risingOrFallingAlarm, then a single rising alarm is generated. If the first sample after this entry becomes valid is less than or equal to the fallingThreshold and alarmStartupAlarm is equal to the fallingAlarm or the risingOrFallingAlarm, then</p>

Name	Description
	a single falling alarm is generated. You cannot modify this object if the associated alarmStatus object is equal to valid.
RisingThreshold	Specifies a threshold for the sampled statistic. After the current sampled value is greater than or equal to this threshold, and the value at the last sampling interval was less than this threshold, a single event is generated. A single event is also generated if the first sample after this entry becomes valid is greater than or equal to this threshold and the associated alarmStartupAlarm is equal to risingAlarm or risingOrFallingAlarm. After a rising event is generated, another such event is not generated until the sampled value falls below this threshold and reaches the alarmFallingThreshold. You cannot modify this object if the associated alarmStatus object is equal to valid.
RisingEventIndex	Specifies the index of the eventEntry that is used after a rising threshold is crossed. The eventEntry identified by a particular value of this index is the same as identified by the same value of the eventIndex object. If there is no corresponding entry in the eventTable, no association exists. You cannot modify this object if the associated alarmStatus object is equal to valid. The default is 60534.
FallingThreshold	Specifies a threshold for the sampled statistic. If the current sampled value is less than or equal to this threshold, and the value at the last sampling interval was greater than this threshold, a single event is generated. A single event is also generated if the first sample after this entry becomes valid is less than or equal to this threshold and the associated alarmStartupAlarm is equal to fallingAlarm or risingOrFallingAlarm. After a falling event is generated, another such event is not generated until the sampled value rises above this threshold and reaches the alarmRisingThreshold. You cannot modify this object if the associated alarmStatus object is equal to valid.
FallingEventIndex	Specifies the index of the eventEntry that is used after a falling threshold is crossed. The eventEntry identified by a particular value of this index is the same as identified by the same value of the eventIndex object. If there is no corresponding entry in the eventTable, no association exists. You cannot modify this object if the associated alarmStatus object is equal to valid. The default is 60535.
Owner	Specifies the entity that configured this entry and is therefore using the resources assigned to it. The default value is CLI if the entry is configured using ACLI. The default is SNMP if the entry is configured using EDM or SNMP.
Status	Specifies the status of this alarm entry.

Creating a port history alarm

Create a port history alarm to track the number of alarms fired from a particular port.

Before you begin

- Ensure that you globally enable RMON.

Enabling RMON globally turns on logging and trapping.

Procedure

1. Select the port that has an alarm configured.
2. Right-click the port.
3. Choose **Enable Rmon Stats and Enable Rmon History**.

Viewing RMON alarms

View the RMON alarm information to see alarm activity.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > Serviceability > RMON**.
2. Click **Alarms**.
3. Click the **Alarm** tab.

Deleting an alarm

Delete an alarm if you no longer want it to appear in the log.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > Serviceability > RMON**.
2. Click **Alarms**.
3. Select the alarm you must delete.
4. Click **Delete**.

Creating a default RMON event

Create a default rising and falling event to specify if alarm information is sent to a trap, a log, or both.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > Serviceability > RMON**.
2. Click **Alarms**.
3. Click the **Events** tab.

4. Click **Insert**.
5. In the **Description** box, type a description for the event.
6. In the **Owner** box, type the owner of the event.
7. In the **Insert Events** dialog box, click **Insert**.

If Rmon is not globally enabled, the following message appears:

```
RMON is currently disabled. Do you want to enable it now?
```

8. Click **Yes**.

Events field descriptions

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

Name	Description
Index	Uniquely identifies an entry in the event table. Each such entry defines one event that is generated after the appropriate conditions occur. The default is 1.
Description	Specifies a comment describing this event entry.
Type	Specifies the type of notification that the probe makes about this event. In the case of a log, an entry is made in the log table for each event. In the case of SNMP traps, an SNMP trap is sent to one or more management stations.
Community	Specifies the SNMP community to where you can send SNMP traps.
LastTimeSent	Specifies the value of sysUpTime at the time this event entry last generated an event. If this entry has not generated events, this value is zero.
Owner	Specifies the entity that configured this entry and is therefore using the assigned resources. If this object contains a string starting with monitor and has associated entries in the log table, all connected management stations retrieve those log entries, as they have significance to all management stations connected to this device. The default value is CLI if the entry is configured using ACLI. The default is SNMP if the entry is configured using EDM or SNMP.

Creating a nondefault RMON event

Create a custom rising and falling event to specify if alarm information is sent to a trap, a log, or both.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > Serviceability > RMON**.
2. Click **Alarms**.
3. Click the **Events** tab.
4. Click **Insert**.

5. In the **Description** box, type an event name.
6. In the **Type** option, select an event type.

The default configuration is log-and-trap. To save memory, configure the event type to log. To reduce traffic from the system, configure the event type to snmp-log.

If you select snmp-trap or log, you must configure trap receivers.

7. In the **Community** box, type an SNMP community.
8. In the **Owner** box, type the owner of this event.
9. Click **Insert**.

Events field descriptions

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

Name	Description
Index	Uniquely identifies an entry in the event table. Each such entry defines one event that is generated after the appropriate conditions occur. The default is 1.
Description	Specifies a comment describing this event entry.
Type	Specifies the type of notification that the probe makes about this event. In the case of a log, an entry is made in the log table for each event. In the case of SNMP traps, an SNMP trap is sent to one or more management stations.
Community	Specifies the SNMP community to where you can send SNMP traps.
LastTimeSent	Specifies the value of sysUpTime at the time this event entry last generated an event. If this entry has not generated events, this value is zero.
Owner	Specifies the entity that configured this entry and is therefore using the assigned resources. If this object contains a string starting with monitor and has associated entries in the log table, all connected management stations retrieve those log entries, as they have significance to all management stations connected to this device. The default value is CLI if the entry is configured using ACLI. The default is SNMP if the entry is configured using EDM or SNMP.

Viewing RMON events

View RMON events to see how many events occurred.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > Serviceability > RMON**.
2. Click **Alarms**.
3. Click the **Events** tab.

Events field descriptions

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

Name	Description
Index	Uniquely identifies an entry in the event table. Each such entry defines one event that is generated after the appropriate conditions occur. The default is 1.
Description	Specifies a comment describing this event entry.
Type	Specifies the type of notification that the probe makes about this event. In the case of a log, an entry is made in the log table for each event. In the case of SNMP traps, an SNMP trap is sent to one or more management stations.
Community	Specifies the SNMP community to where you can send SNMP traps.
LastTimeSent	Specifies the value of sysUpTime at the time this event entry last generated an event. If this entry has not generated events, this value is zero.
Owner	Specifies the entity that configured this entry and is therefore using the assigned resources. If this object contains a string starting with monitor and has associated entries in the log table, all connected management stations retrieve those log entries, as they have significance to all management stations connected to this device. The default value is CLI if the entry is configured using ACLI. The default is SNMP if the entry is configured using EDM or SNMP.

Deleting an event

Delete an event after you no longer require the alarm information.

Before you delete an RMON event, remove all RMON alarms related to the RMON event.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > Serviceability > RMON**.
2. Click **Alarms**.
3. Click the **Events** tab.
4. Select the event you must delete.
5. Click **Delete**.

Chapter 10: Service Level Agreement Monitor

The switch supports the Service Level Agreement Monitor (SLA Mon™) agent as part of the Avaya SLA Mon solution.

SLA Mon uses a server and agent relationship to perform end-to-end network Quality of Service (QoS) validation and to distribute monitoring devices. You can use the test results to target under-performing areas of the network for deeper analysis.

SLA Mon server and agent

The switch supports the SLA Mon agent. You must have an Avaya Diagnostic Server with SLA Mon technology in your network to use the SLA Mon feature. Most of the SLA Mon configuration occurs on the server; configuration on the SLA Mon agent is minimal.

The SLA Mon server initiates the SLA Mon functions on one or more agents, and the agents run specific QoS tests at the request of the server. Agents can exchange packets between one another to conduct the QoS tests.

SLA Mon can monitor a number of key items, including the following:

- network paths
- Differentiated Services Code Point (DSCP) markings
- loss
- jitter
- delay

The following figure shows an SLA Mon implementation.

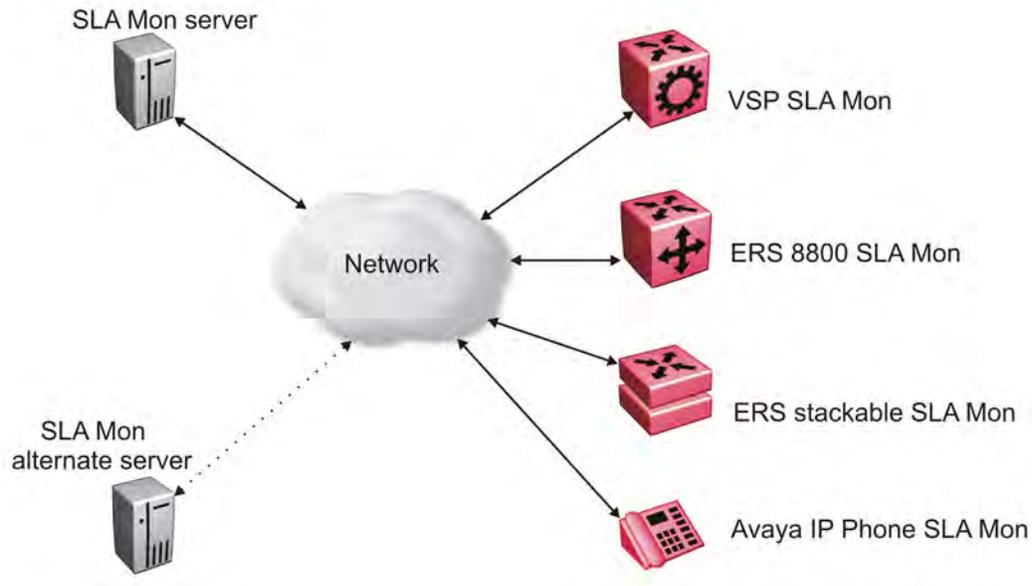


Figure 4: SLA Monitor network

An SLA Mon agent remains dormant until it receives a User Datagram Protocol (UDP) discovery packet from a server. The agent accepts the discovery packet to register with an SLA Mon server. If the registration process fails, the agent remains dormant until it receives another discovery packet.

An agent can attempt to register with an SLA Mon server once every 60 seconds. After a successful registration, the agent reregisters with the server every 6 hours to exchange a new encryption key.

An agent only accepts commands from the SLA Mon server to which it is registered. An agent can use alternate SLA Mon servers to provide backup for time-out and communication issues with the primary SLA Mon server.

Note:

If you configure the SLA Mon agent address under an IP address for a VLAN or brouter, you must remove the SLA Mon address before you can remove the IP address for the VLAN or brouter.

QoS tests

SLA Mon uses two types of tests to determine QoS benchmarks:

- Real Time Protocol (RTP)

This test measures network performance — for example, jitter, delay, and loss — by injecting a short stream of UDP packets from source to destination (an SLA Mon agent).

- New Trace Route (NTR)

This test is similar to traceroute but also includes DSCP values at each hop in the path from the source to the destination. The destination does not need to be an SLA Mon agent.

Limitations

SLA Mon agent communications are IPv4-based. Agent communications do not currently support IPv6.

SLA Mon configuration using ACLI

Configuring the SLA Mon agent

Configure the SLA Mon agent to communicate with an Avaya Diagnostic Server with SLA Mon technology to perform Quality of Service (QoS) tests of the network.

Before you begin

- To use the SLA Mon agent, you must have an Avaya Diagnostic Server with SLA Mon technology in your network.

About this task

To configure the SLA Mon agent, you must assign an IP address and enable it. Remaining agent parameters are optional and you can operate the agent using the default values.

In HA mode, the agent startup and initialization occurs only on the master CP module.

Note:

If you want to change SLA Mon parameters, you must first disable SLA Mon.

If you configure the SLA Mon agent address under an IP address for a VLAN or brouter, you must remove the SLA Mon address before you can remove the IP address for the VLAN or brouter. To remove the SLA Mon address, first use the command `no slamon oper-mode enable`, followed by `slamon agent ip address 0.0.0.0`.

Procedure

1. Enter Application Configuration mode:

```
enable
configure terminal
application
```
2. Configure the SLA Mon agent IP address:

Note:

The SLA Mon agent IP address must not use the IP address of an IP interface on the switch.

```
slamon agent ip address {A.B.C.D} [vrf WORD<1-16>]
```

3. **(Optional)** Configure the UDP port for agent-server communication:

```
slamon agent port <0-65535>
```

4. **(Optional)** Restrict which servers an agent can use:

```
slamon server ip address {A.B.C.D} [{A.B.C.D}]
```

```
slamon server port <0-65535>
```

5. **(Optional)** Control the port used for Real Time Protocol (RTP) and New Trace Route (NTR) testing:

```
slamon agent-comm-port <0-65535>
```

6. **(Optional)** Install a Secure Socket Layer (SSL) certificate for the agent:

```
slamon install-cert-file WORD<0-128>
```

7. Enable the agent:

```
slamon oper-mode enable
```

8. Verify the agent configuration:

```
show application slamon agent
```

Example

Configure the SLA Mon agent IP address. Configure the agent so that it only accepts registration packets from a specific server communicating on a specific port. Finally, enable the SLA Mon agent, 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)#application
Switch:1(config-app)#slamon agent ip address 192.0.2.1
Switch:1(config-app)#slamon server ip address 192.0.2.25
Switch:1(config-app)#slamon server port 50011
Switch:1(config-app)#slamon oper-mode enable
Switch:1(config-app)#show application slamon agent
```

```
=====
                          SLA Monitor Agent Info
=====
```

```
SLAMon Operational Mode: Enabled
SLAMon Agent Address: 192.0.2.1
SLAMon Agent Port: 50011
SLAMon Agent Registration Status: Registered
SLAMon Registered Server Address: 192.0.2.25
SLAMon Registered Server Port: 50011
SLAMon Server Registration Time: 130
SLAMon Encryption Mode: Supported
SLAMon Configured Agent Address: 192.0.2.1
SLAMon Configured Agent Port: 0
SLAMon Configured Server Address: 192.0.2.25 0.0.0.0
```

```
SLAMon Configured Server Port: 50011 0
SLAMon Agent-To-Agent Communication Port: 50012
SLAMon Configured Agent-To-Agent Communication Port: 0
SLAMon Configured Agent Address Vrf Name:
```

Next steps

If you have configured SLA Mon, but the agent does not function as expected, use the **show khi performance pthread** [{slot[-slot] [, ...]}] command to verify that the slamon task is running.

If the SLA Mon agent is not running, use the commands **no slamon oper-mode enable** and **slamon oper-mode enable** to start the agent.

If the agent task is running, perform typical troubleshooting steps to verify agent accessibility:

- Verify IP address assignment and port use.
- Ping the server IP address.
- Verify the server configuration.
- Use the **trace level 192 <0-4>** command to observe the status of the SLA Mon software module.

Variable definitions

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

Variable	Value
agent-comm-port <0-65535>	Configures the port used for RTP and NTR testing in agent-to-agent communication. The default port is 50012. If you configure this value to zero (0), the default port is used.
agent ip address {A.B.C.D}	Configures the SLA Mon agent IP address. You must configure the IP address before the agent can process received discovery packets from the server. The agent ip address is a mandatory parameter. The default value is 0.0.0.0.
agent port <0-65535>	Configures the UDP port for agent-server communication. The SLA Mon agent receives discovery packets on this port. The default is port 50011. The server must use the same port.
install-cert-file	Installs an SSL certificate. <i>WORD</i> <0-128>specifies the file name and path of the certificate to install. If you install a certificate on the SLA Mon agent, you must ensure a matching configuration on the server. By default, the agent uses an Avaya SIP certificate to secure communications with the server.
oper-mode enable	Enables the SLA Mon agent. The default is disabled.

Variable	Value
	<p>If you disable the agent, it does not respond to discovery packets from a server.</p> <p>If you disable the agent because of resource concerns, consider changing the server configuration instead, to alter the test frequency or duration, or the number of targets.</p>
server ip address {A.B.C.D} [{A.B.C.D}]	<p>Restricts the SLA Mon agent to use of this server IP address only. The default is 0.0.0.0, which means the agent can register with any server.</p> <p>You can specify a secondary server as well.</p>
server port <0–65535>	<p>Restricts the SLA Mon agent to use of this registration port only. The default is 0, which means the agent disregards the source port information in server traffic.</p> <p>The server must use the same port.</p>
vrf WORD<1-16>	Specifies the name of a VRF.

SLA Mon configuration using EDM

Configuring the SLA Mon agent

Configure the SLA Mon agent to communicate with an Avaya Diagnostic Server with SLA Mon technology to perform Quality of Service (QoS) tests of the network.

Before you begin

- To use the SLA Mon agent, you must have an Avaya Diagnostic Server with SLA Mon technology in your network.

About this task

To configure the SLA Mon agent, you must assign an IP address and enable it. Remaining agent parameters are optional and you can operate the agent using the default values.

In HA mode, the agent startup and initialization occurs only on the master CP module.

Note:

If you want to change SLA Mon parameters, you must first disable SLA Mon.

If you configure the SLA Mon agent address under an IP address for a VLAN or router, you must remove the SLA Mon address, before you can remove the IP address for the VLAN or router. To remove the SLA Mon address, first select disabled from the **Status** field, then configure the IP address in the **ConfiguredAgentAddr** field to 0.0.0.0.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > Serviceability**.
2. Click **SLA Monitor**.
3. Click the **SLA Monitor** tab.
4. For the status, select **enabled**.
5. In the **ConfiguredAgentAddr** field, enter the SLA Mon agent IP address
6. Configure optional parameters as required.
7. Click **Apply**.

SLA Monitor field descriptions

Use the data in the following table to use the **SLA Monitor** tab.

Name	Description
Status	Enables or disables the SLA Mon agent. The default is disabled. If you disable the agent, it does not respond to discovery packets from a server. If you disable the agent because of resource concerns, consider changing the server configuration instead, to alter the test frequency or duration, or the number of targets.
CertFileInstallAction	Installs or uninstalls a Secure Sockets Layer (SSL) certificate file. The default is noAction.
CertFile	Specifies the file name and path of the SSL certificate. If you install a certificate on the SLA Mon agent, you must ensure a matching configuration on the server. By default, the agent uses an Avaya SIP certificate to secure communications with the server.
ConfiguredAgentAddrType	Specifies the address type of the agent: IPv4.
ConfiguredAgentAddr	Configures the agent IP address. You must configure the IP address before the agent can process received discovery packets from the server. The agent IP address is a mandatory parameter. The default value is 0.0.0.0.
ConfiguredAgentPort	Configures the UDP port for agent-server communication. The SLA Mon agent receives discovery packets on this port. The default is port 50011. The server must use the same port.
ConfiguredAgentVrfName	Specifies the name of a VRF.
ConfiguredServerAddrType	Specifies the address type of the server: IPv4.

Name	Description
ConfiguredServerAddr	Restricts the SLA Mon agent to use of this server IP address only. If the default of 0.0.0.0 is used, then the SLA Mon agent can register with any server.
ConfiguredServerPort	Restricts the SLA Mon agent to use of this registration port only. The default is 0, which means the agent disregards the source port information in server traffic. The server must use the same port.
ConfiguredAltServerAddrType	Specifies the address type of the secondary server: IPv4.
ConfiguredAltServerAddr	Configures a secondary server in the event that the primary server is unreachable.
ConfiguredAltServerPort	Restricts the SLA Mon agent to use of this registration port on the secondary server only. The default is 0, which means the agent disregards the source port information in server traffic. The server must use the same port.
SupportedApps	Shows the type of testing supported: Real Time Protocol (RTP) and New Trace Route (NTR).
AgentAddressType	Shows the SLA Mon agent address type.
AgentAddress	Shows the configured SLA Mon agent IP address.
AgentPort	Shows the configured SLA Mon agent port.
RegisteredWithServer	Indicates if the SLA Mon agent has registered with a server.
RegisteredServerAddrType	Shows the address type for the registered server.
RegisteredServerAddr	Shows the IP address for the registered server.
RegisteredServerPort	Shows the port number for the registered server.
RegistrationTime	Shows the amount of time, in seconds, since the SLA Mon agent registered with the server.
AgentToAgentPort	Shows the port for SLA Mon agent-to-agent communication.
ConfiguredAgentToAgentPort	Configures the port used for RTP and NTR testing in SLA Mon agent-to-agent communication. The default port is 50012. If you configure this value as zero (0), the default port is used.

Chapter 11: Statistics

This chapter provides the procedures for using statistics to help monitor the performance of the Avaya Virtual Services Platform 9000 using Enterprise Device Manager (EDM) and Avaya command line interface (ACLI).

SPBM packet drop statistics

This enhanced packet drop feature keeps statistics on a per-port basis. The SPBM Drop Stats by Port tables display information on the source (SA) or destination (DA) MAC address of the last packet dropped.

This feature collects and displays frame drops on ingress at SPBM NNI interfaces in the following categories:

- Last drop
Displays information for the last packets dropped.
- Reverse Path Forwarding Check (RPFC) multicast SA drops
Displays the total number of SPBM RPFC multicast packets dropped. These drops occur when packets with a specific source MAC address ingress on a port different than what IS-IS expected.
- RPFC unicast SA drops
Displays the total number of SPBM RPFC unicast packets dropped. These drops occur when packets with a specific source MAC address ingress on a port different than what IS-IS expected.
- Unknown unicast DA drops
Displays the total number of SPBM unknown unicast DA packets dropped. These drops occur when the unicast destination MAC address of the packet is not known.
- Unknown unicast SA drops
Displays the total number of SPBM unknown unicast SA packets dropped. These drops occur when the unicast source MAC address of the packet is not known. VSP 9000 does not support the unknown-unicast-sa drop count parameter for second generation modules in this release. The device always displays the ACLI output for second generation modules as 0 for this counter.

- Unknown multicast DA drops

Displays the total number of SPBM unknown multicast DA packets dropped. These drops occur when the multicast destination MAC address of the packet is not known.

Viewing statistics using ACLI

This section contains procedures to view statistics in the ACLI.

Viewing TCP statistics

View TCP statistics to manage network performance.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. View TCP statistics:

```
show ip tcp statistics
```

Example

```
VSP-9012:1#show ip tcp statistics
show ip tcp global statistics:
-----
ActiveOpens:      0
PassiveOpens:    37
AttemptFails:    0
EstabResets:     34
CurrEstab:       1
InSegs:          6726
OutSegs:         7267
RetransSegs:     10
InErrs:          0
OutRsts:         10
```

Job aid

The following table describes the output for the `show ip tcp statistics` command.

Table 6: show ip tcp statistics command output

Field	Description
ActiveOpens	The count of transitions by TCP connections to the SYN-SENT state from the CLOSED state.
PassiveOpens	The count of transitions by TCP connections to the SYN-RCVD state from the LISTEN state.
AttemptFails	The count of transitions by TCP connections to the CLOSED state from either the SYN-SENT state or

Field	Description
	the SYN-RCVD state, plus the count of transitions to the LISTEN state from the SYN-RCVD state.
EstabResets	The count of transitions by TCP connections to the CLOSED state from the ESTABLISHED or CLOSE-WAIT state.
CurrEstab	The count of TCP connections for which the current state is either ESTABLISHED or CLOSE-WAIT.
InSegs	The total count of segments received, including those received in error. This count includes segments received on currently established connections.
OutSegs	The total number of segments sent, including those on current connections but excluding those containing only retransmitted octets.
RetransSegs	The total count of TCP segments transmitted containing one or more previously transmitted octets.
InErrs	The count of segments received in error.
OutRsts	The count of TCP segments sent containing the RST flag.

Viewing port routing statistics

View port routing statistics to manage network performance.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. View port routing statistics:

```
show routing statistics interface [gigabitethernet] [{slot/port[-slot/port] [, ...]}
```

Example

```
VSP-9012:1#show routing statistics interface gigabitethernet 4/7-4/9
```

```
=====
                          Port Stats Routing
=====
```

PORT NUM	IN_FRAME UNICAST	IN_FRAME MULTICAST	IN DISCARD	OUT_FRAME UNICAST	OUT_FRAME MULTICAST
4/7	1386	0	0	1344	0
4/8	1302	0	0	1344	0
4/9	0	0	0	0	0

```
=====
```

Variable definitions

Use the data in the following table to use the `show routing statistics interface` command.

Variable	Value
gigabitethernet	Specifies the interface type.
{slot/port[-slot/port][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).

Job aid

The following table describes the output for the `show routing statistics interface` command.

Table 7: show routing statistics interface field descriptions

Parameter	Description
PORT NUM	Indicates the port number.
IN_FRAME UNICAST	The count of inbound unicast frames.
IN_FRAME MULTICAST	The count of inbound multicast frames.
IN DISCARD	The count of inbound discarded frames.
OUT_FRAME UNICAST	The count of outbound unicast frames.
OUT_FRAME MULTICAST	The count of outbound multicast frames.

Displaying bridging statistics for specific ports

Display individual bridging statistics for specific ports to manage network performance.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View bridging statistics for a specific port:

```
show interfaces GigabitEthernet statistics bridging [{slot/port[-slot/port][,...]}]
```

Example

```
VSP-9012:1#show interfaces gigabitEthernet statistics bridging
```

```
=====
                          Port Stats Bridge
=====
```

PORT NUM	IN_FRAME UNICAST	IN_FRAME MULTICAST	IN_FRAME BROADCAST	OUT_FRAME	IN_FRAME xSTP BPDU	OUT_FRAME xSTP BPDU	IN_DISCARD
4/1	179325	0	0	119310	179325	0	0
4/2	187951	26078	42	689486	179324	0	25617
4/3	0	0	0	0	0	0	0
4/4	0	0	0	0	0	0	0

```
=====
```

Statistics

```
4/5 0 0 0 0 0 0 0
4/6 394 0 0 948942 360 0 0
4/7 4689 0 0 863403 360 0 0
4/8 4369 3206 116 958752 360 0 3995
4/9 0 0 0 0 0 0 0
4/10 0 0 0 0 0 0 0
4/11 0 0 0 0 0 0 0
4/12 0 0 0 0 0 0 0
4/13 179325 0 0 42040 179325 0 0
4/14 187864 0 0 50437 179324 0 0
4/15 0 0 0 0 0 0 0
4/16 0 0 0 0 0 0 0
--More-- (q = quit)
```

Variable definitions

Use the data in the following table to use the **show interfaces GigabitEthernet statistics bridging** command.

Variable	Value
{slot/port[-slot/port][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).

Job aid

The following table describes parameters for the **show interfaces GigabitEthernet statistics bridging** command.

Table 8: show interfaces gigabitEthernet statistic bridging field descriptions

Parameter	Description
PORT NUMB	Port index of the statistics table.
IN_FRAME UNICAST	The count of inbound Unicast frames.
IN_FRAME MULTICAST	The count of inbound Multicast frames.
IN_FRAME BROADCAST	The count of inbound Broadcast frames.
OUT_FRAME	The count of outbound frames.

Displaying DHCP-relay statistics for specific ports

Display individual DHCP-relay statistics for specific ports to manage network performance.

Procedure

1. Enter Privileged EXEC mode:
enable
2. View DHCP-relay statistics for a specific port or VRF.

```
show interfaces GigabitEthernet statistics dhcp-relay [vrf
WORD<0-16>] [vrfids WORD<0-255>] [{slot/port[-slot/port][, ...]]}
```

Variable definitions

Use the data in the following table to use the **show interfaces GigabitEthernet statistics dhcp-relay** command.

Variable	Value
vrf WORD<0-16>	Specifies a VRF instance by VRF name.
vrfids WORD<0-255>	Specifies the ID of the VRF.
{slot/port[-slot/port][, ...]}	Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2–3/4), or a series of slots and ports (3/2, 5/3, 6/2).

Job aid

The following table describes parameters for the **show interfaces GigabitEthernet statistics dhcp-relay** command output.

Table 9: show interfaces gigabitethernet statistics dhcp-relay field descriptions

Variable	Value
PORT_NUM	Indicates the port number.
VRF NAME	Identifies the VRF
NUMREQUEST	Indicates the total number of DHCP requests on this interface
NUMREPLY	Indicates the total number of DHCP replies on this interface.

Displaying DHCP-relay statistics for all interfaces

Display DHCP-relay statistics for all interfaces to manage network performance.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Show the number of requests and replies for each interface:

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

3. Show counters for Option 82:

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

Example

```
VSP-9012:1>show ip dhcp-relay counters option82
=====
                        DHCP Counters Option82 - GlobalRouter
=====
      FOUND DROP  CIRCUIT ADD  REMOVE REMOTE      ADD  REMOVE
INTERFACE OPT82 PKT   ID    CIRC  CIRC  ID              REMOTE REMOTE
-----
Port6/12  0    0   395    0    0    00:24:7f:9d:0a:00  0    0
Vlan40    0    0   2088   0    0    00:24:7f:9d:0a:01  0    0
```

Variable definitions

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

Variable	Value
vrf <i>WORD</i> <1-16>	Specifies a VRF instance by the VRF name.
vrfids <i>WORD</i> <0-512>	Specifies the ID of the VRF.

Job aid

The following table explains the output from the **show ip dhcp-relay counters option82** command.

Table 10: show ip dhcp-relay counters option82 command

Heading	Description
INTERFACE	Shows the name of the interface on which you enabled option 82. Shows the port number if the interface is a brouter port or the VLAN number if the interface is a VLAN.
FOUND OPT82	Shows the number of packets that the interface received that already had option82 in them.
DROP PKT	Shows the number of packets the interface dropped because of option 82–related issues. These reasons could be that the packet was received from an untrusted source or spoofing was detected. To determine the cause of the drop, you must enable trace on level 170.
CIRCUIT ID	Show the value inserted in the packets as the circuit ID. The value is the index of the interface.
ADD CIRC	Shows on how many packets (requests from client to server) the circuit ID was inserted for that interface. If you expect this value to increase but it does not, and the interface does not drop a packet, it is possible the packet does not have enough space to insert the option. You must enable trace on level 170 to determine the cause.

Heading	Description
REMOVE CIRC	Shows on how many packets (replies from server to client) the circuit id was removed for that interface.
REMOTE ID	Shows the value inserted in the packets as the remote ID. The value is the MAC address of the interface.
ADD REMOTE	Shows on how many packets (requests from client to server) the remote ID was inserted for that interface. If you expect this value to increase but it does not, and the interface does not drop a packet, it is possible the packet does not have enough space to insert the option. You must enable trace on level 170 to determine the cause.
REMOVE REMOTE	Shows on how many packets (replies from server to client) the remote ID was removed for that interface.

Viewing IPv6 DHCP Relay statistics

Display individual IPv6 DHCP Relay statistics for specific interfaces to manage network performance.

Procedure

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

```
show ipv6 dhcp-relay counters
```

Note:

Use the `sys action reset counters` command to clear DHCP Relay statistics.

Example

```
Switch:1#show ipv6 dhcp-relay counters
```

```
=====
                        DHCPv6 Counters
=====
INTERFACE                REQUESTS    REPLIES
-----
1111:0:0:0:0:0:1111      1          1
```

Job aid

The following table explains the output of the `show ipv6 dhcp-relay counters` command.

Table 11: show ipv6 dhcp-relay counters command output

Heading	Description
REQUESTS	Shows the number of DHCP and BootP requests on this interface.
REPLIES	Shows the number of DHCP and BootP replies on this interface.

Displaying LACP statistics for specific ports

Display individual LACP statistics for specific ports to manage network performance.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View statistics for specific ports:

```
show interfaces GigabitEthernet statistics lacp [{slot/port[-slot/port] [, ...]]
```

Variable definitions

Use the data in the following table to use the **show interfaces GigabitEthernet statistics lacp** command.

Variable	Value
{slot/port[-slot/port][, ...]}	Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).

Job aid

The following table describes parameters for the **show interfaces GigabitEthernet statistics lacp** command.

Table 12: show interfaces GigabitEthernet statistics lacp field descriptions

Parameter	Description
PORT_NUM	Indicates the port number.
TX LACPDU	The count of transmitted LACP data units.
RX LACPDU	The count of received LACP data units.
TX MARKERPDU	The count of transmitted marker protocol data units.
RX MARKERPDU	The count of received marker protocol data units.

Parameter	Description
TX MARKERRESPDU	The count of transmitted marker protocol response data units.
RX MARKERRESPDU	The count of received marker protocol response data units.
RX UNKNOWN	The count of received unknown frames.
RX ILLEGAL	The count of received illegal frames.

Displaying RMON statistics for specific ports

Display individual RMON statistics for specific ports to manage network performance.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View statistics for specific ports:

```
show interfaces GigabitEthernet statistics rmon {slot/port[-slot/
port][,...]}
```

Variable definitions

Use the data in the following table to use the **show interfaces GigabitEthernet statistics rmon** command.

Variable	Value
{slot/port[-slot/port][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).

Job aid

The following table describes parameters for the **show interfaces GigabitEthernet statistics rmon** command output.

Table 13: show interfaces GigabitEthernet statistics rmon field descriptions

Parameter	Description
PORT NUM	Indicates the port number.
OCTETS	The total number of octets of data (including those in bad packets) received on the network (excluding framing bits but including FCS octets).
PKTS	The total number of packets (including bad packets, broadcast packets, and multicast packets) received.

Parameter	Description
MULTICAST	The total number of packets received that were directed to a multicast address. This number does not include packets directed to the broadcast address.
BROADCAST	The total number of packets received that were directed to the broadcast address. This number does not include multicast packets.
CRC ALLIGN	The total number of packets received that had a length (excluding framing bits, but including FCS octets) between 64 and 1518 octets, inclusive, but had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error), or a bad FCS with a nonintegral number of octets (Alignment Error).
UNDERSIZE	The total number of packets received that were less than 64 octets long (excluding framing bits, but including FCS octets) and were otherwise well formed.
OVERSIZE	The total number of packets received that were longer than 1518 octets (excluding framing bits, but including FCS octets) and were otherwise well formed.
FRAGMENT	The total number of packets received that were less than 64 octets in length (excluding framing bits but including FCS octets) and had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).
COLLISION	An estimated value for the total number of collisions on this Ethernet segment.

Displaying detailed statistics for ports

Display detailed statistics for specific ports to manage network performance.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View statistics for specific ports:

```
show interfaces GigabitEthernet statistics verbose {slot/port[-slot/
port][,...]}
```

Example

```
VSP-9012:1#show interface gigabitEthernet statistics verbose
=====
Port Stats Interface Extended
=====
PORT_NUM IN_UNICST OUT_UNICST IN_MULTICST OUT_MULTICST IN_BRDCST OUT_BRDCST IN_LSM OUT_LSM
-----
4/1      0      0      422479      221764      0      1      0      0
4/2      400     1      564619      1431955     4      68     0      0
4/3      0      0      0           0           0      0      0      0
4/4      0      0      0           0           0      0      0      0
4/5      0      0      0           0           0      0      0      0
4/6      0      0      0           0           0      0      0      0
4/7      0      0      0           0           0      0      0      0
4/8      0      0      0           0           0      0      0      0
4/9      0      0      0           0           0      0      0      0
4/10     0      0      0           0           0      0      0      0
4/11     0      0      0           0           0      0      0      0
4/12     0      0      0           0           0      0      0      0
4/13     0      0      422479      99046       0      0      0      0
4/14     0      0      442596      118859      0      0      0      0
4/15     0      0      0           0           0      0      0      0

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

Variable definitions

Use the data in the following table to use the **show interfaces GigabitEthernet statistics verbose** command.

Variable	Value
{slot/port[-slot/port][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).

Job aid

The following table describes parameters for the **show interfaces GigabitEthernet statistics verbose** command.

Table 14: how interfaces GigabitEthernet statistics verbose field descriptions

Parameter	Description
PORT_NUM	Indicates the port number.
IN_UNICAST	The count of inbound Unicast packets.
OUT_UNICAST	The count of outbound Unicast packets.
IN_MULTICAST	The count of inbound Multicast packets.
OUT_MULTICAST	The count of outbound Multicast packets.
IN_BRDCST	The count of inbound broadcast packets.
OUT_BRDCST	The count of outbound broadcast packets.

Displaying IS-IS statistics and counters

Use the following procedure to display the IS-IS statistics and counters.

Procedure

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

2. Display IS-IS system statistics:

```
show isis statistics
```

3. Display IS-IS interface counters:

```
show isis int-counters
```

4. Display IS-IS level 1 control packet counters:

```
show isis int-l1-cntl-pkts
```

Note:

The current release uses level 1 IS-IS. The current release does not support level 2 IS-IS. The ACLI command `show isis int-l2-cntl-pkts` is not supported in the current release because the IEEE 802.1aq standard currently only defines the use of one hierarchy, Level 1.

5. Enter Privileged EXEC mode:

```
enable
```

6. Clear IS-IS statistics:

```
clear isis stats [error-counters] [packet-counters]
```

Example

```
VSP-9012:1# show isis statistics
```

```
=====
                        ISIS System Stats
=====
LEVEL      CORR   AUTH   AREA   MAX SEQ   SEQ NUM   OWN   LSP   BAD ID   PART   LSP   DB
          LSPs  FAILS  DROP   EXCEEDED SKIPS    PURGE  LEN   CHANGES OLOAD
-----
Level-1    0      0      0      0         1         0     0     0       0     0
```

```
VSP-9012:1#show isis int-counters
```

```
=====
                        ISIS Interface Counters
=====
IFIDX     LEVEL   AUTH   ADJ      INIT     REJ     ID LEN   MAX AREA LAN   DIS
          LEVEL  FAILS  CHANGES  FAILS    FAILS   ADJ      CHANGES  CHANGES
-----
Mlt2     Level 1-2  0      1         0       0       0       0       0       0
```

```
Port3/21 Level 1-2 0 1 0 0 0 0 0
```

```
VSP-9012:1#show isis int-l1-ctrl-pkts
```

```
=====
                        ISIS L1 Control Packet counters
=====
IFIDX          DIRECTION    HELLO        LSP          CSNP         PSNP
-----
Mlt2           Transmitted    13346        231          2            229
Mlt2           Received      13329        230          1            230
Port3/21       Transmitted    13340        227          2            226
Port3/21       Received      13335        226          1            227
=====
```

Variable definitions

Use the data in the following table to use the `clear isis stats` command.

Variable	Value
error-counters	Clears IS-IS stats error-counters.
packet-counters	Clears IS-IS stats packet-counters.

Job aid

show isis statistics

The following table describes the fields in the output for the `show isis statistics` command.

Parameter	Description
LEVEL	Shows the level of the IS-IS interface (Level 1 in the current release).
CORR LSPs	Shows the number of corrupted LSPs detected.
AUTH FAILS	Shows the number of times authentication has failed on the global level.
AREA DROP	Shows the number of manual addresses dropped from the area.
MAX SEQ EXCEEDED	Shows the number of attempts to exceed the maximum sequence number.
SEQ NUM SKIPS	Shows the number of times the sequence number was skipped.
OWN LSP PURGE	Shows how many times the local LSP was purged.
BAD ID LEN	Shows the number of ID field length mismatches.
PART CHANGES	Shows the number of partition link changes.
LSP DB OLOAD	Show the number of times the Virtual Services Platform 9000 was in the overload state.

show isis int-counters

The following table describes the fields in the output for the `show isis int-counters` command.

Parameter	Description
IFIDX	Shows the interface index for the Ethernet or MLT interface.

Parameter	Description
LEVEL	Shows the level of the IS-IS interface (Level 1 in the current release).
AUTH FAILS	Shows the number of times authentication has failed per interface.
ADJ CHANGES	Shows the number of times the adjacencies have changed.
INIT FAILS	Shows the number of times the adjacency has failed to establish.
REJ ADJ	Shows the number of times the adjacency was rejected by another router.
ID LEN	Shows the ID field length mismatches.
MAX AREA	Shows the maximum area address mismatches.
LAN DIS CHANGES	Shows the number of times the DIS has changed.

show isis int-l1-ctrl-pkts

The following table describes the fields in the output for the `show isis int-l1-ctrl-pkts` command.

Parameter	Description
IFIDX	Shows the interface index for the Ethernet or MLT interface.
DIRECTION	Shows the packet flow (Transmitted or Received).
HELLO	Shows the amount of interface-level Hello packets.
LSP	Shows the amount of LSP packets.
CSNP	Shows the amount of CSNPs.
PSNP	Shows the amount of PSNPs.

Displaying SPBM packet drop statistics by port

Use this procedure to display the SPBM drop statistics.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display the last dropped statistics:

```
show isis spbm drop-stats port last-drop [{slot/port[-slot/port]
[,...]]}
```

3. Display the RPFC multicast source MAC drop statistics:

```
show isis spbm drop-stats port rpfc-multicast-sa [{slot/port[-slot/
port][,...]]}
```

4. Display the RPFC unicast source MAC drop statistics:

```
show isis spbm drop-stats port rpfc-unicast-sa [{slot/port[-slot/
port][,...]]}
```

5. Display the unknown multicast destination MAC drop statistics:

```
show isis spbm drop-stats port unknown-multicast-da [{slot/port[-slot/port][, ...]]}
```

6. Display the unknown unicast destination MAC drop statistics:

```
show isis spbm drop-stats port unknown-unicast-da [{slot/port[-slot/port][, ...]]}
```

7. Display the unknown unicast source MAC drop statistics:

```
show isis spbm drop-stats port unknown-unicast-sa [{slot/port[-slot/port][, ...]]}
```

Note:

Virtual Services Platform 9000 does not support the **show isis spbm drop-stats port unknown-unicast-sa** drop count parameter for second generation modules in this release. The device always displays the ACLI output for second generation modules as 0 for this counter.

Example

The following output shows examples for the **show isis spbm drop-stats port** command.

```
VSP-9012:1#show isis spbm drop-stats port last-drop
```

```
=====
                        SPBM Drop Stats By Port
                        Last Drop
=====
PORT  PRIMARY B-VID      DA  B-MAC      SECONDARY B-VID      DA  B-MAC
NUM  HOST NAME
-----
3/11  evp                  N  00:13:0a:e6:73:df  evp                  N  00:13:0a:e6:73:df
3/13  evs                  N  00:13:0a:e6:43:df  evs                  N  00:13:0a:e6:43:df
```

```
VSP-9012:1#show isis spbm drop-stats port unknown-unicast-sa 3/11
```

```
=====
                        SPBM Drop Stats By Port
                        Unknown Unicast Source Address
=====
PORT  PRIMARY B-VID      SECONDARY B-VID
NUM  PKT DROP           PKT DROP
-----
3/11  4                  4
```

```
VSP-9012:1#show isis spbm drop-stats port rpf-check-unicast-sa
```

```
=====
                        SPBM Drop Stats By Port
                        Reverse Path Forwarding Check Unicast Source Address
=====
PORT  PRIMARY B-VID      SECONDARY B-VID
NUM  PKT DROP           PKT DROP
-----
3/11  0                  0
3/13  0                  0
```

Variable definitions

Use the data in the following table to use the `show isis spbm drop-stats port` commands.

Variable	Value
{slot/port[-slot/port][,....]}	Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).

Job aid

The following table describes the fields in the output for the `show isis spbm drop-stats port last-drop` command.

Table 15: show isis spbm drop-stats port last-drop field descriptions

Field	Description
PORT NUM	Shows the slot/port number that identifies the ingress port of the dropped packet.
PRIMARY B-VID HOST NAME	Shows the name of the primary SPBM B-VID.
DA	Shows whether there were dropped packets. <ul style="list-style-type: none"> • Y indicates that the last drop B-MAC is a destination MAC. • N indicates that the last drop B-MAC is a source MAC.
B-MAC	Shows the backbone MAC address of the primary SPBM B-VLAN.
SECONDARY B-VID HOST NAME	Shows the name of the secondary SPBM B-VID, if known.
DA	Shows whether there were dropped packets: Yes (Y) or No (N).
B-MAC	Shows the backbone MAC address of the secondary SPBM B-VLAN.

The following table describes the fields in the output for all of the other `show isis spbm drop-stats port` commands.

Table 16: show isis spbm drop-stats port field descriptions

Field	Description
PORT NO	Shows the slot/port number that identifies the ingress port of the dropped packet.
PRIMARY B-VID PKT DROP	Shows the total number of RPFC multicast drops for each primary SPBM B-VLAN.
SECONDARY B-VID PKT DROP	Shows the total number of RPFC multicast drops for each secondary SPBM B-VLAN.

Clearing SPBM packet drop statistics

Clear drop statistics to reset the counters.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Clear SPBM port-based drop statistics:

```
clear isis spbm drop-stats [{slot/port[-slot/port][,...]} ]
```

Variable definitions

Use the data in the following table to use the `clear isis spbm drop-stats` command.

Variable	Value
{slot/port[-slot/port][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).

Displaying policing statistics

View statistics to ensure proper QoS performance.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. View policing statistics:

```
show qos statistics policy [<1-16000>] [slot {slot[-slot][,...]}]
```

Variable definitions

Use the data in the following table to use the `show qos statistics policy` command.

Variable	Value
1-16000	Specifies an optional policy ID,
slot {slot[-slot][,...]}	Identifies the slot in one of the following formats: a single slot (3), a range of slots (3-4), or a series of slots (3,5,6).

Job aid

The following table describes the output for the `show qos statistics policy` command.

Table 17: show qos statistics policy field descriptions

Parameter	Description
Policer Name	Specifies the packet policer name.
Id	Identifies a global policer (GP) ID value that corresponds to the local policer. Valid values range from 1 to 16 383.
lane ports	Specifies a port number for a set of lanes.
Total pkts	Specifies the total packets.
Total Bytes	Specifies the total bytes.
BytesOvr SvcRate	Specifies the bytes over the local policer service rate.
BytesOvr PeakRate	Specifies the bytes over the local policer peak rate.
DropOth pkts	Specifies other dropped packets.

Clearing ACL statistics

Clear default ACL statistics if you no longer require previous statistics.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Enter the following command to clear default ACL statistics:

```
clear filter acl statistics default [<1-2048>]
```

3. Enter the following command to clear global ACL statistics:

```
clear filter acl statistics global [<1-2048>]
```

4. Enter the following command to clear all ACL statistics:

```
clear filter acl statistics all
```

5. Enter the following command to clear statistics associated with a particular ACL, ACE, or ACE type:

```
clear filter acl statistics [<1-2048>] [<1-2000>][qos] [security]
```

Variable definitions

Use the information in the following table to use the `clear filter acl statistics` command.

Variable	Value
1-2048	Specifies the ACL ID.
1-2000	Specifies the ACE ID.

Viewing ACE statistics

View ACE statistics to ensure that the filter operates correctly.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. View ACE statistics for a specific ACL, ACE, or ACE type:
`show filter acl statistics <1-2048> [<1-2000>] [qos] [security]`
3. View all ACE statistics:
`show filter acl statistics all`
4. View default ACE statistics:
`show filter acl statistics default [<1-2048>]`
5. View global statistics for ACEs:
`show filter acl statistics global [<1-2048>]`

Variable definitions

Use the data in the following table to use the `show filter acl statistics` command.

Variable	Value
<code>1-2048</code>	Specifies the ACL ID.
<code>1-2000</code>	Specifies the ACE ID.

Job aid

The following table describes output for the `show filter acl statistics port` command.

Table 18: show filter acl statistics port field descriptions

Parameter	Description
Acl ID	Specifies the identifier for the ACL.
Acl Name	Specifies the name for the ACL.
Acl Type	Specifies the ACL type.
Ace Id	Specifies the ACE identifier.
Port Num	Specifies the port number.
Packets	Specifies the number of packets on the port.
Bytes	Specifies the number of bytes on the port.

Viewing MSTP statistics

Display MSTP statistics to see MSTP related bridge-level statistics.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display the MSTP related bridge-level statistics:

```
show spanning-tree mstp statistics
```

Example

```
VSP-9012:1#show spanning-tree mstp statistics
```

```
=====
MSTP Bridge Statistics
=====
Mstp UP Count           : 1
Mstp Down Count         : 0
Region Config Change Count : 12
Time since topology change : 8 day(s), 02H:54M:33S
Topology change count   : 10
New Root Bridge Count   : 25
```

Job aid

The following table describes the output for the **show spanning-tree mstp statistics** command.

Table 19: show spanning-tree mstp statistics field descriptions

Parameter	Description
MSTP Up Count	The number of times the MSTP Module has been enabled. A Trap is generated on the occurrence of this event.
MSTP Down Count	The number of times the MSTP Module has been disabled. A Trap is generated on the occurrence of this event.
Region Config Change Count	The number of times the switch detects a Region Configuration Identifier Change. The switch generates a trap on the occurrence of this event.
Time since topology change	The time (in hundredths of a second) since the TcWhile Timer for any port in this Bridge was non-zero for Common Spanning Tree context.
Topology change count	The count of at least one non zero TcWhile timers on this Bridge for Common Spanning Tree context.
New Root Bridge Count	The number of times this Bridge has detected a Root Bridge change for Common Spanning Tree context. A Trap is generated on the occurrence of this event.

Viewing RSTP statistics

View Rapid Spanning Tree Protocol statistics to manage network performance.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. View RSTP stats with the following command:

```
show spanning-tree rstp statistics
```

Job aid

The following table describes output for the `show spanning-tree rstp statistics` command.

Table 20: show spanning-tree rstp statistics field descriptions

Parameter	Description
RSTP Up Count	The number of times RSTP Module has been enabled. A Trap is generated on the occurrence of this event.
RSTP Down Count	The number of times RSTP Module has been disabled. A Trap is generated on the occurrence of this event.
Count of Root Bridge Changes	The number of times this Bridge has detected a Root Bridge change for Common Spanning Tree context.
STP Time since Topology change	The time (in hundredths of a second) since the "TcWhile" Timer for any port in this Bridge was non zero for this spanning tree instance.
Total number of topology changes	The number of times that there have been atleast one non zero "TcWhile" Timer on this Bridge for this spanning tree instance.

Viewing RSTP port statistics

View RSTP statistics on ports to manage network performance.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. View RSTP statistics on a port:

```
show spanning-tree rstp port statistics [{slot/port[-slot/port]}
[,...]]
```

Variable definitions

Use the data in the following table to use the `show spanning-tree rstp port statistics` command.

Variable	Value
{slot/port[-slot/port][,....]}	Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).

Job aid

The following table describes output for the `show spanning-tree rstp port statistics` command.

Table 21: show spanning-tree rstp port statistics field descriptions

Parameter	Description
RxRstBpduCount	The number of RSTP BPDUs received on this port.
RxConfigBpduCount	The number of configuration BPDUs received on this port.
RxTcnBpduCount	The number of TCN BPDUs received on this port.
TxRstBpduCount	The number of RSTP BPDUs transmitted by this port.
TxConfigBpduCount	The number of Config BPDUs transmitted by this port.
TxTcnBpduCount	The number of TCN BPDUs transmitted by this port.
InvalidRstBpduRxCount	The number of invalid RSTP BPDUs received on this port. A trap is generated on the occurrence of this event.
InvalidConfigBpduRx Count	The number of invalid configuration BPDUs received on this port. A trap is generated on the occurrence of this event.
InvalidTcnBpduRxCount	The number of invalid TCN BPDUs received on this port. A trap is generated on the occurrence of this event.
ProtocolMigrationCount	The number of times this port migrated from one STP protocol version to another. The relevant protocols are STP-Compatible and RSTP. A trap is generated on the occurrence of this event.

Viewing MLT statistics

View MLT statistics to display MultiLinkTrunking statistics for the switch or for the specified MLT ID.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. View MLT statistics:

```
show mlt stats [<1-512>]
```

Example

```
VSP-9012:1#show mlt stats
```

```
=====
                        Mlt Interface
=====
ID IN-OCTETS           OUT-OCTETS           IN-UNICST           OUT-UNICST
-----
1  256676904           183670662           1397                456
2  61737348498         61584347982         1450182             1490619
4  229256124           47472778            0                   0
100 251678170           32332107            0                   0

ID IN-MULTICST         OUT-MULTICST         IN-BROADCAST         OUT-BROADCAST        MT
-----
1  2419514              2295274              41                   268194               E
2  962303832           960067410           765                  237                   E
4  2159884              666153               0                     90                    E
100 2095269             504965               13                    0                      E

ID IN-LSM              OUT-LSM
-----
1  0                    0
2  957925732           957929399
4  0                    0

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

Variable definitions

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

Variable	Value
<1-512>	Specifies the MLT ID.

Job aid

The following table describes the output for the `show mlt stats` command.

Table 22: show mlt stats field descriptions

Parameter	Description
ID IN-OCTETS	The total number of inbound octets of data (including those in bad packets).
OUT-OCTETS	The total number of outbound octets of data.
IN-UNICAST	The count of inbound Unicast packets.
OUT-UNICAST	The count of outbound unicast packets.
ID IN-MULTICAST	The count of inbound multicast packets.

Parameter	Description
OUT-MULTICAST	The count of outbound multicast packets.
IN-BROADCAST	The count of inbound broadcast packets.
OUT-BROADCAST	The count of outbound broadcast packets.
MT	The MLT type: P for POS, E for Ethernet, A for ATM.

Showing OSPF error statistics on a port

Use statistics to help you monitor Open Shortest Path First (OSPF) performance. You can also use statistics in troubleshooting procedures.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display extended information about OSPF errors for the specified port or for all ports:

```
show interfaces GigabitEthernet error ospf [{slot/port[-slot/port]}
[,...]]
```

Variable definitions

Use the following table to help you use the `show interfaces GigabitEthernet error ospf` command.

Variable	Value
{slot/port[-slot/port][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).

Job aid

The following table describes the output for the `show interfaces GigabitEthernet error ospf` command.

Table 23: show interfaces GigabitEthernet error ospf field descriptions

Parameters	Description
PORT NUM	Indicates the port number.
VERSION MISMATCH	Indicates the number of version mismatches this interface receives.
AREA MISMATCH	Indicates the number of area mismatches this interface receives.
AUHTYPEMISMATCH	Indicates the number of AuthType mismatches this interface receives.
AUTH FAILURES	Indicates the number of authentication failures.

Parameters	Description
NET_MASK MISMATCH	Indicates the number of net mask mismatches this interface receives.
HELLOINT MISMATCH	Indicates the number of hello interval mismatches this interface receives.
DEADINT MISMATCH	Indicates the number of dead interval mismatches this interface receives.
OPTION MISMATCH	Indicates the number of options mismatches this interface receives.

Viewing OSPF interface statistics

Use statistics to help you monitor Open Shortest Path First (OSPF) performance. You can also use statistics in troubleshooting procedures.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display OSPF interface statistics:

```
show ip ospf ifstats [detail] [mismatch] [vlan <1-4084>] [vrf
WORD<1-16>] [vrfids WORD<0-512>]
```

Example

```
VSP-9012:1#show ip ospf ifstats
```

```
=====
                        OSPF Interface Statistics - GlobalRouter
=====
INTERFACE          ---HELLOS---  ---DBS---  -LS REQ--  --LS UPD---  --LS ACK---
                   RX      TX      RX  TX  RX  TX  RX  TX  RX  TX
-----
2.2.2.32            76035  76355  33   32   4   9   2483 2551 2525 1247
30.30.30.32        76038  76349   0    0   0   0    0    0    0    0
40.1.1.32          153207 76355  38   44   6  11  2899 3797 4203 1601
=====
```

Variable definitions

Use this table to help you use the `show ip ospf ifstats` command.

Variable	Value
detail	Shows detailed information.
mismatch	Shows the number of times the area ID is not matched.
vlan <1-4084>	Specifies a VLAN ID.
vrf WORD<1-16>	Specifies a VRF instance by VRF name.
vrfids WORD<0-512>	Specifies a VRF or range of VRFs by ID.

Job aid

The following table describes the output for the `show ip ospf ifstats` command.

Table 24: show ip ospf ifstats field descriptions

Field	Description
INTERFACE	Indicates the IP address of the host.
HELLOS RX	Indicates the number of hello packets received by this interface.
HELLOS TX	Indicates the number of hello packets transmitted by this interface.
DBS RX	Indicates the number of database descriptor packets received by this interface.
DBS TX	Indicates the number of database descriptor packets transmitted by this interface.
LS REQ	Indicates the number of link state request packets received by this interface.
LS TX	Indicates the number of link state request packets transmitted by this interface.
LS UDP RX	Indicates the number of link state update packets received by this interface.
LS UDP TX	Indicates the number of link state update packets transmitted by this interface.
LS ACK RX	Indicates the number of link state acknowledge packets received by this interface.
LS ACK TX	Indicates the number of link state acknowledge packets transmitted by this interface.
VERSION	Indicates the OSPF version.
AREA	Indicates the OSPF area.
AUTHTYPE	Indicates the OSPF authentication type.
AUTHFAIL	The count of authentication fail messages.
NETMASK	Indicates the net mask.
HELLO	The count of Hello messages.
DEADTRR OPTION	The dead TRR option.

Viewing OSPF range statistics

Use statistics to help you monitor Open Shortest Path First (OSPF) performance. You can also use statistics in troubleshooting procedures. OSPF range statistics include area ID, range network address, range subnet mask, range flag, and LSDB type.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display the OSPF range statistics:

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

Example

```
VSP-9012:1#show ip ospf stats
```

```
=====
                        OSPF Statistics - GlobalRouter
=====
      NumBufAlloc: 239603
      NumBufFree: 239603
NumBufAllocFail: 0
NumBufFreeFail: 0
      NumTxPkt: 239655
      NumRxPkt: 317562
NumTxDropPkt: 0
NumRxDropPkt: 0
      NumRxBadPkt: 0
      NumSpfRun: 47
      LastSpfRun: 2 day(s), 04:18:58
      LsdbTblSize: 16
NumAllocBdDDP: 24
NumFreeBdDDP: 24
      NumBadLsReq: 0
NumSeqMismatch: 3
      NumOspfRoutes: 4
      NumOspfAreas: 1
NumOspfAdjacencies: 3
--More-- (q = quit)
```

Variable definitions

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

Variable	Value
vrf <i>WORD</i> <1-16>	Specifies a VRF instance by VRF name.
vrfids <i>WORD</i> <0-512>	Specifies a VRF or range of VRFs by ID.

Job aid

The following table describes the show command output.

Table 25: show ip ospf stats command parameters

Parameter	Description
NumBufAlloc	Indicates the number of buffers allocated for OSPF.
NumBufFree	Indicates the number of buffers that are freed by the OSPF.
NumBufAllocFail	Indicates the number of times that OSPF failed to allocate buffers.
NumBufFreeFail	Indicates the number of times that OSPF failed to free buffers.
NumTxPkt	Indicates the number of packets transmitted by OSPF.
NumRxPkt	Indicates the number of packets received by OSPF.
NumTxDropPkt	Indicates the number of packets dropped before transmission by OSPF.
NumRxDropPkt	Indicates the number of packets dropped before reception by OSPF.
NumRxBadPkt	Indicates the number of packets received by OSPF that are bad.

Parameter	Description
NumSpfRun	Indicates the total number of SPF calculations performed by OSPF, which also includes the number of partial route table calculation for incremental updates.
LastSpfRun	Indicates the time (SysUpTime) since the last SPF calculated by OSPF.
LsdbTblSize	Indicates the number of entries in the link state database table.
NumAllocBdDDP	Indicates the number of times buffer descriptors were allocated for OSPF database description packets.
NumFreeBdDDP	Indicates the number of times buffer descriptors were freed after use as OSPF database description packets.
NumBadLsReq	Indicates the number of bad LSDB requests.
NumSeqMismatch	Indicates the number of mismatches for sequence numbers.
NumOspfRoutes	The count of OSPF routes.
NumOspfAreas	The count of OSPF areas.
NumOspfAdjacencies	The count of Adjacencies.

Viewing basic OSPF statistics for a port

Use statistics to help you monitor Open Shortest Path First (OSPF) performance. You can also use statistics in troubleshooting procedures.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. View basic OSPF statistics:

```
show ports statistics ospf main [{slot/port[-slot/port][, ...]}
```

Variable definitions

Use the data in the following table to use the `show ports statistics ospf main` command.

Variable	Value
{slot/port[-slot/port][, ...]}	Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).
Variable	Value
{slot/port[-slot/port][, ...]}	Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).

Job aid

The following table describes the output for the `show ports statistics ospf main` command.

Table 26: show ports statistics ospf main output description

Field	Description
PORT NUM	Indicates the port number.
RX_HELLO	Indicates the number of hello packets this interface receives.
TX_HELLO	Indicates the number of hello packets this interface transmitted.
RXDB_DESCR	Indicates the number of database descriptor packets this interface receives.
TXDB_DESCR	Indicates the number of database descriptor packets this interface transmitted.
RXLS_UPDATE	Indicates the number of link state update packets this interface receives.
TXLS_UPDATE	Indicates the number of link state update packets this interface transmitted.

Showing extended OSPF statistics

Use statistics to help you monitor Open Shortest Path First (OSPF) performance. You can also use statistics in troubleshooting procedures.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display extended OSPF information about the specified port or for all ports:

```
show ports statistics ospf extended [{slot/port[-slot/port][,...]]
```

Variable definitions

Use the data in the following table to use the `show ports statistics ospf extended` command.

Variable	Value
{slot/port[-slot/port][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).

Job aid

The following table describes the output for the `show ports statistics ospf extended` command.

Table 27: show ports statistics ospf extended output description

Parameters	Description
PORT_NUM	Indicates the port number.
RXLS_REQS	Indicates the number of link state update request packets received by this interface.
TXLS_REQS	Indicates the number of link state request packets transmitted by this interface.
RXLS_ACKS	Indicates the number of link state acknowledge packets received by this interface.
TXLS_ACKS	Indicates the number of link state acknowledge packets transmitted by this interface.

Viewing IPv6 OSPF statistics

View OSPF statistics to analyze trends.

Procedure

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

```
show ipv6 ospf statistics
```

Example

Job aid

The following table explains the output of the `show ipv6 ospf statistics` command.

Field	Description
NumTxPkt	Shows the count of sent packets.
NumRxPkt	Shows the count of received packets.
NumTxDropPkt	Shows the count of sent, dropped packets.
NumRxDropPkt	Shows the count of received, dropped packets.
NumRxBadPkt	Shows the count of received, bad packets.
NumSpfRun	Shows the count of intra-area route table updates with calculations using this area linkstate database.
LastSpfRun	Shows the count of the most recent SPF run.
LsdbTblSize	Shows the size of the link state database table.
NumBadLsReq	Shows the count of bad link requests.
NumSeqMismatch	Shows the count of sequence mismatched packets.

Showing the EAPoL status of the device

Display the current device configuration.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display the current device configuration by using the following command:

```
show eapol system
```

Example

```
Switch:1#show eapol system
                        eap : enabled
                        sess-manage : false
```

Showing EAPoL authenticator statistics

Display the authenticator statistics to manage network performance.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display the authenticator statistics:

```
show eapol auth-stats interface [gigabitEthernet [{slot/port[-slot/
port][, ...]}]] [vlan <1-4084>]
```

Example

```
Switch:1#show eapol auth-stats interface
```

```
=====
                        Eap Authenticator Statistics
=====
```

PORT	TOTAL RX	TOTAL TX	START RCVD	LOGOFF RCVD	RESP_ID RCVD	RESP RCVD	REQ-ID TX	REQ TX	INVALID FRAMES	LENGTH ERROR	FRAME VER	LAST-SRC MAC
4/1	0	1	0	0	0	0	0	1	0	0	0	00:00:00:00:00:00
4/2	0	1	0	0	0	0	0	1	0	0	0	00:00:00:00:00:00
4/3	0	0	0	0	0	0	0	0	0	0	0	00:00:00:00:00:00
4/4	0	0	0	0	0	0	0	0	0	0	0	00:00:00:00:00:00
4/5	0	0	0	0	0	0	0	0	0	0	0	00:00:00:00:00:00
4/6	0	9	0	0	0	0	0	9	0	0	0	00:00:00:00:00:00
4/7	0	9	0	0	0	0	0	9	0	0	0	00:00:00:00:00:00
4/8	0	9	0	0	0	0	0	9	0	0	0	00:00:00:00:00:00
4/9	0	0	0	0	0	0	0	0	0	0	0	00:00:00:00:00:00
4/10	0	0	0	0	0	0	0	0	0	0	0	00:00:00:00:00:00
4/11	0	0	0	0	0	0	0	0	0	0	0	00:00:00:00:00:00
4/12	0	0	0	0	0	0	0	0	0	0	0	00:00:00:00:00:00
4/13	0	1	0	0	0	0	0	1	0	0	0	00:00:00:00:00:00
4/14	0	1	0	0	0	0	0	1	0	0	0	00:00:00:00:00:00
4/15	0	0	0	0	0	0	0	0	0	0	0	00:00:00:00:00:00
4/16	0	0	0	0	0	0	0	0	0	0	0	00:00:00:00:00:00

Variable definitions

Use the data in the following table to use the `show eapol auth-stats interface` command.

Variable	Value
gigabitEthernet [{slot/port[-slot/port][,...]]	Specifies the slot and port. If you do not specify a slot and port, the show command returns results for all GigabitEthernet interfaces.
<1-4084>	Specifies the VLAN ID for which to show the statistics.

Showing EAPoL session statistics

View EAPoL session statistics to manage network performance.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display the session statistics:

```
show eapol session-stats interface [gigabitEthernet [{slot/port[-slot/port][,...]]] [vlan <1-4084>]
```

Example

```
Switch:1#show eapol session-stats interface
```

```
=====
Eap Authenticator Session Statistics
=====
```

PORT	TOTAL OCTETS RCVD	TOTAL OCTETS TXMT	TOTAL FRAMES RCVD	TOTAL FRAMES TXMT	SESSION ID	AUTHENTIC METHOD	SESSION TIME	TERMINATE CAUSE	USER NAME
4/1	0	0	0	0		local-server	0 day(s),	00:00:00	supp-logoff
4/2	0	0	0	0		local-server	0 day(s),	00:00:00	supp-logoff
4/3	0	0	0	0		local-server	0 day(s),	00:00:00	supp-logoff
4/4	0	0	0	0		local-server	0 day(s),	00:00:00	supp-logoff
4/5	0	0	0	0		local-server	0 day(s),	00:00:00	supp-logoff
4/6	0	0	0	0		local-server	0 day(s),	00:00:00	supp-logoff
4/7	0	0	0	0		local-server	0 day(s),	00:00:00	supp-logoff
4/8	0	0	0	0		local-server	0 day(s),	00:00:00	supp-logoff
4/9	0	0	0	0		local-server	0 day(s),	00:00:00	supp-logoff
4/10	0	0	0	0		local-server	0 day(s),	00:00:00	supp-logoff
4/11	0	0	0	0		local-server	0 day(s),	00:00:00	supp-logoff
4/12	0	0	0	0		local-server	0 day(s),	00:00:00	supp-logoff
4/13	0	0	0	0		local-server	0 day(s),	00:00:00	supp-logoff
4/14	0	0	0	0		local-server	0 day(s),	00:00:00	supp-logoff
4/15	0	0	0	0		local-server	0 day(s),	00:00:00	supp-logoff

Variable definitions

Use the data in the following table to use the `show eapol session-stats interface` command.

Variable	Value
gigabitEthernet [{slot/port[-slot/port][,...]]	Specifies the slot and port. If you do not specify a slot and port, the show command returns results for all GigabitEthernet interfaces.

Variable	Value
<1-4084>	Specifies the VLAN ID for which to show the statistics.

Job aid

The following table describes the output for the `show eapol session-stats interface` command.

Table 28: show eapol session interface field descriptions

Parameter	Description
TOTAL OCTETS RCVD	Displays the number of octets received in user data frames on this port during the session.
TOTAL OCTETS TXMT	Displays the number of octets transmitted in user data frames on this port during the session.
TOTAL FRAMES RCVD	Displays the number of user data frames received on this port during the session.
TOTAL FRAMES TXMT	Displays the number of user data frames transmitted on this port during the session.
SESSION ID	Displays a unique identifier for the session that is at least three characters.
AUTHENTIC METHOD	Displays the authentication method (remote or local RADIUS server) used to establish the session.
SESSION TIME	Displays the duration of the session (in seconds).
TERMINATE CAUSE	Displays the reason the session terminated.
USER NAME	Displays the user name of the Supplicant PAE.

Showing RADIUS server statistics

You cannot collect the following network statistics from a console port: the number of input and output packets, and the number of input and output bytes. All other statistics from console ports are available to assist with debugging.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Display RADIUS server statistics:
`show radius-server statistics`
3. Clear server statistics:
`clear radius statistics`

Example

```
VSP-9012:1#show radius-server statistics
Responses with invalid server address: 0
  Radius Server (UsedBy) : 47.17.143.58 (cli)
-----
  Access Requests : 52
  Access Accepts : 0
  Access Rejects : 0
  Bad Responses : 52
  Client Retries : 52
  Pending Requests : 0
  Acct On Requests : 1
  Acct Off Requests : 0
  Acct Start Requests : 47
  Acct Stop Requests : 46
  Acct Interim Requests : 0
  Acct Bad Responses : 94
  Acct Pending Requests : 0
  Acct Client Retries : 94
  Access Challenges : 0
  Round-trip Time :
  Nas Ip Address : 47.17.10.32
  Radius Server (UsedBy) : 47.17.143.58 (snmp)
-----
  Access Requests : 0
  Access Accepts : 0
  Access Rejects : 0
  Bad Responses : 0
  Client Retries : 0
  Pending Requests : 0
  Acct On Requests : 0
  Acct Off Requests : 0
  Acct Start Requests : 0
  Acct Stop Requests : 0
  Acct Interim Requests : 0
  Acct Bad Responses : 0
  Acct Pending Requests : 0
  Acct Client Retries : 0
  Access Challenges : 0
  Round-trip Time :
  Nas Ip Address : 47.17.10.32
--More-- (q = quit)
```

Job aid

The following table shows the field descriptions for the **show radius-server statistics** command output.

Table 29: show radius-server statistics command fields

Parameter	Description
RADIUS Server	The IP address of the RADIUS server.
AccessRequests	Number of access-response packets sent to the server; does not include retransmissions.

Parameter	Description
AccessAccepts	Number of access-accept packets, valid or invalid, received from the server.
AccessRejects	Number of access-reject packets, valid or invalid, received from the server.
BadResponses	Number of invalid access-response packets received from the server.
PendingRequests	Access-request packets sent to the server that have not yet received a response, or have timed out.
ClientRetries	Number of authentication retransmissions to the server.
AcctOnRequests	Number of accounting On requests sent to the server.
AcctOffRequests	Number of accounting Off requests sent to the server.
AcctStartRequests	Number of accounting Start requests sent to the server.
AcctStopRequests	Number of accounting Stop requests sent to the server.
AcctInterimRequests	Number of accounting Interim Requests sent to the server. The AcctInterimRequests counter increments only if the parameter acct-include-cli-commands is set to true.
AcctBadResponses	Number of Invalid Responses from the server that are discarded.
AcctPendingRequests	Number of requests waiting to be sent to the server.
AcctClientRetries	Number of retries made to this server.

Viewing RMON statistics

View RMON statistics to manage network performance.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. View RMON statistics:

```
show rmon stats
```

Example

```
VSP-9012:1(config)#show rmon stats
```

```
=====
                                Rmon Ether Stats
=====
INDEX  PORT   OWNER
-----
1      cpp    monitor
```

Job aid

The following table describes parameters in the output for the `show rmon stats` command.

Table 30: show rmon stats field descriptions

Parameter	Description
Index	An index that uniquely identifies an entry in the Ethernet statistics table.
Port	Identifies the source of the data that this entry analyzes.
Owner	The entity that configured this entry and is therefore using the assign resources. The default value is CLI if the entry was configured using ACLI. The default is SNMP if the entry was configured using EDM or SNMP.

Viewing PCAP statistics

View PCAP statistics to manage network performance.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. View PCAP statistics:

```
show pcap stats
```

Example

View PCAP statistics:

```
VSP-9012:1#show pcap stats
```

```
Stat Information for PCAP
=====
Packet Capacity Count : 381300
Number of packets received in PCAP engine : 0
Number of packets accumulated in PCAP engine : 0
Number of packets dropped in PCAP engine by filters : 0
Number of packets dropped in Hardware : 0
```

Job aid

The following table describes parameters for the `show pcap stats` command output.

Table 31: show pcap stats field descriptions

Parameter	Description
Packet Capacity Count	The maximum number of packets that currently can be stored in the PCAP engine buffer. Reset-stat does not reset this value.
Number of packets received in PCAP engine	The number of packets currently in the PCAP engine buffer. When buffer-wrap occurs, the value is set to 0 and the count starts again.

Parameter	Description
	<p>Important:</p> <p>When buffer-wrap occurs, the second field is set to 0 and the third field is not set to zero. From the capture log, the user can determine how many times buffer-wrap occurred.</p>
Number of packets accumulated in PCAP engine	<p>This is the number of packets accumulated in the PCAP engine.</p> <p>Important:</p> <p>When buffer-wrap occurs, the second field is set to 0 and the third field is not set to zero. From the capture log, the user can determine how many times buffer-wrap occurred.</p>
Number of packets dropped in PCAP engine by filters	The number of packets dropped when ingress packets match the filter criteria and the PCAP action is set to drop.
Number of packets dropped in Hardware	The number of packets dropped by the PCAP engine hardware when the amount of packets being forwarded cannot be processed.

Viewing IPFIX statistics

View the exporter statistics for each slot to see the following information:

- collector IP address
- packets sent since you enabled IPFIX
- bytes sent since you enabled IPFIX
- packets lost within the device
- IPFIX protocol status

View the hashing statistics to view total hash overflows.

About this task

If you do not specify a slot, all slots appear in the command output.

Procedure

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

2. View exporter statistics:

```
show ip ipfix export [{slot[-slot]][,...]]
```

3. View hashing statistics:

```
show ip ipfix hash-statistics [{slot[-slot]][,...]]
```

Virtual Services Platform 9000 does not support the `show ip ipfix hash-statistics` command for second generation modules in this release.

Example

```
VSP-9012:1#show ip ipfix export 4
=====
IPFIX Exporter-Statistics
=====
SlotNum      Collector-IP      Number of      Number of      Number of
Address      packets sent     bytes sent     packets lost
-----
4            47.17.143.146   20             3280           0

VSP-9012:1#show ip ipfix hash-statistics 4
=====
IPFIX Hash-Statistics
=====
SlotNum      Hash Overflows   Hash Drops
(resource contention)
-----
4            0                0
```

Variable definitions

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

Variable	Value
{slot[-slot][,...]}	Identifies the slot in one of the following formats: a single slot (3), a range of slots (3-4), or a series of slots (3,5,6).

Clearing IPFIX statistics

Clear IPFIX statistics to remove the exporter and hash statistics.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Clear exporter statistics:
clear ip ipfix stats [{slot[-slot][,...]}
3. Clear hash statistics:
clear ip ipfix hash-stats [{slot[-slot][,...]}

Virtual Services Platform 9000 does not support the `clear ip ipfix hash-stats` command for second generation modules in this release.

Example

```
VSP-9012:1#clear ip ipfix stats
VSP-9012:1#clear ip ipfix hash-stats 4
```

Variable definitions

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

Variable	Value
{slot[-slot][,...]}	Identifies the slot in one of the following formats: a single slot (3), a range of slots (3-4), or a series of slots (3,5,6). If you do not specify a slot, you clear the statistics for all slots.

Clearing IPv6 statistics

Clear all IPv6 statistics if you do not require previous statistics.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Enter the following command to clear all the IPv6 statistics:

```
clear ipv6 statistics all
```

3. Enter the following command to clear interface statistics:

```
clear ipv6 statistics interface [general|icmp] [gigabitethernet
<slot/port>|mgmtethernet <slot/port>|vlan <1-4084>]
```

4. Enter the following command to clear TCP statistics:

```
clear ipv6 statistics tcp
```

5. Enter the following command to clear UDP statistics:

```
clear ipv6 statistics udp
```

Variable definitions

Use the information in the following table to use the `clear ipv6 statistics` command.

Variable	Value
<1-4084>	Specifies the VLAN ID in the range of 1-4084. VLAN IDs 1 to 4084 are configurable; VLAN IDs 4085 to 4094 are reserved for internal use. VLAN ID 1 is the default VLAN and cannot be created nor deleted.
slot/port	Specifies the slot and the port number to clear router interface statistics.
slot/port	Specifies the slot and the port number to clear statistics on a management port.

Viewing multicast routing process statistics

Enable the collection and display of multicast routing process statistics. The statistics display based on multicast group classification. The statistics are not related to the interface (port) statistics. By default, mroute statistics collection is disabled.

Before you begin

- To use the monitor command, you must log on to Privileged EXEC mode.
- To enable the collection of statistics, you must log on to Global Configuration mode.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display multicast routing process statistics at regular intervals:

```
monitor ip mroute stats WORD<7-160>
```

You can change the duration or interval for monitoring in Global Configuration mode.

3. Enter Global Configuration mode:

```
configure terminal
```

4. Enable multicast routing process statistics collection:

```
ip mroute stats enable
```

5. Disable multicast routing process statistics collection:

```
no ip mroute stats enable
```

6. View multicast routing process statistics:

```
show ip mroute stats [WORD<7-160>]
```

7. Clear the multicast routing process statistics:

```
clear ip mroute stats
```

Example

View multicast routing process statistics:

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1#ip mroute stats enable
VSP-9012:1(config)#show ip mroute stats 233.252.0.1
=====
                          Multicast Stats - GlobalRouter
=====
GroupAddress SourceCounter IngressPackets IngressBytes AverageSize Packets/Second
DropPackets DropBytes
-----
233.252.0.1 1                1090179126      89394689233 82                1225653
0            0
```

Variable definitions

Use the information in the following table to use the `monitor ip mroute stats` and `show ip mroute stats` commands.

Variable	Value
<code>WORD<7-160></code>	Specifies the multicast group IDs for which to display statistics. You can specify a maximum of 10 groups.

Job aid

The following table shows the field descriptions for the `show ip mroute stats` command.

Table 32: show ip mroute stats

Field	Description
GroupAddress	Specifies the multicast group IP address for which to show statistics.
SourceCounter	Specifies the number of sources associated with the multicast route record.
IngressPackets	Specifies the number of packets received for the associated IP address.
IngressBytes	Specifies the number of bytes received for the associated IP address.
AverageSize	Specifies the average packet length for the associated group IP address. This information indicates only the ingress packet length and is calculated using the following formula: ingress packet/ingress byte.
Packets/Second	Specifies the average speed. This field is only valid in the monitor output. The value is calculated using the following formula: (current ingress packet – last ingress packet)/ monitor interval. With the first monitor multicast statistics output, this field is null. Subsequent outputs provide valid values.
DropPackets	Specifies the number of dropped packets for the associated VRF and group IP address.
DropBytes	Specifies the number of dropped bytes for the associated VRF and group IP address.

Viewing IPv6 VRRP statistics

View IPv6 VRRP statistics to monitor network performance

Procedure

1. Log on to the switch to enter User EXEC mode.
2. View statistics for the device and for all interfaces:

```
show ipv6 vrrp statistics [link-local WORD<0-127>]] [vrid <1-255>]
```

Example

```
Switch:1(config)#show ipv6 vrrp statistics vrid 1
```

```

=====
                        VRRP Global Stats - GlobalRouter
=====
CHK_SUM_ERR   VERSION_ERR   VRID_ERR
0              0              0

=====
                        VRRP Interface Stats - GlobalRouter
=====
VRID  P/V   BECOME_MASTER  ADVERTISE_RCV
-----
1     20    1              0

VRID  P/V   ADVERTISE_INT_ERR  TTL_ERR      PRIO_0_RCV
-----
1     20    0                  0            0

VRID  P/V   PRIO_0_SENT   INVALID_TYPE_ERR  ADDRESS_LIST_ERR  UNKNOWN_AUTHTYPE
-----
1     20    0              0                  0                  0

VRID  P/V   PACKLEN_ERR
-----
1     20    0

```

Variable definitions

Use the data in the following table to use the `show ipv6 vrrp statistics` command.

Variable	Value
link-local <i>WORD<0-127></i>	Shows statistics for a specific link-local address.
vrid <i><1-255></i>	Shows statistics for a specific VRID.

Job aid

The following table describes the output for the `show ipv6 vrrp statistics` command.

Table 33: show ipv6 vrrp statistics command output

Heading	Description
CHK_SUM_ERR	Shows the number of VRRP packets received with an invalid VRRP checksum value.
VERSION_ERR	Shows the number of VRRP packets received with an unknown or unsupported version number.
VRID_ERR	Shows the number of VRRP packets received with an invalid VrID for this virtual router.

Heading	Description
BECOME_MASTER	Shows the total number of times that the state of this virtual router has transitioned to master. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of DiscontinuityTime.
ADVERTISE_RCV	Shows the total number of VRRP advertisements received by this virtual router. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of DiscontinuityTime.
ADVERTISE_INT_ERR	Shows the total number of VRRP advertisement packets received for which the advertisement interval is different than the one configured for the local virtual router. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of DiscontinuityTime.
TTL_ERR	Shows the total number of VRRP packets received by the virtual router with IPv4 TTL (for VRRP over IPv4) or IPv6 Hop Limit (for VRRP over IPv6) not equal to 255. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of DiscontinuityTime.
PRIO_0_RCV	Shows the total number of VRRP packets received by the virtual router with a priority of 0. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of DiscontinuityTime.
PRIO_0_SENT	Shows the total number of VRRP packets sent by the virtual router with a priority of 0. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of DiscontinuityTime.
INVALID_TYPE_ERR	Shows the number of VRRP packets received by the virtual router with an invalid value in the 'type' field. Discontinuities in the value of this counter can occur at reinitialization of the management system, and at other times as indicated by the value of DiscontinuityTime.
ADDRESS_LIST_ERR	Shows the total number of packets received for which the address list does not match the locally configured list for the virtual router. Discontinuities in the value of this counter can occur at reinitialization

Heading	Description
	of the management system, and at other times as indicated by the value of DiscontinuityTime.
UNKNOWN_AUTHTYPE	Shows the total number of packets received with an unknown authentication type.
PACKLEN_ERR	Shows the total number of packets received with a packet length less than the length of the VRRP header. Discontinuities in the value of this counter can occur at reinitialization of the management system, and at other times as indicated by the value of DiscontinuityTime.

Viewing ICMP statistics

View IPv6 ICMP statistics on an interface for ICMP messages sent over a particular interface.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. View IPv6 ICMP statistics

```
show ipv6 interface icmpstatistics
```

Example

View ICMP statistics:

```
Switch:1>show ipv6 interface icmpstatistics
=====
                                Icmp Stats
=====
Icmp stats for IfIndex = 192

IcmpInMsgs: 0
IcmpInErrors: 0
IcmpInDestUnreachs : 0
IcmpInAdminProhibs : 0
IcmpInTimeExcds : 0
IcmpInParmProblems : 0
IcmpInPktTooBigs : 0
IcmpInEchos : 0
IcmpInEchoReplies : 0
IcmpInRouterSolicits : 0
IcmpInRouterAdverts : 0
InNeighborSolicits : 0
InNbrAdverts : 0
IcmpInRedirects : 0
IcmpInGroupMembQueries : 0
IcmpInGroupMembResponses : 0
```

Variable definitions

Use the data in the following table to use the `show ipv6 interface icmpstatistics` command

Variable	Value
<1-4084>	Shows ICMP statistics for the specific interface index. If you do not specify an interface index, the command output includes all IPv6 ICMP interfaces. Specifies the VLAN ID in the range of 1-4084. VLAN IDs 1 to 4084 are configurable; VLAN IDs 4085 to 4094 are reserved for internal use. VLAN ID 1 is the default VLAN and cannot be created nor deleted.

Viewing IPv6 statistics on an interface

View IPv6 statistics to view information about the IPv6 datagrams on an interface.

Procedure

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

```
show ipv6 interface statistics [<1-4084>]
```

Variable definitions

Use the data in the following table to use the `show ipv6 interface statistics` command

Variable	Value
<1-4084>	Shows statistics for the specific interface index. If you do not specify an interface index, the command output includes all IPv6 interfaces. Specifies the VLAN ID in the range of 1-4084. VLAN IDs 1 to 4084 are configurable; VLAN IDs 4085 to 4094 are reserved for internal use. VLAN ID 1 is the default VLAN and cannot be created nor deleted.

Viewing statistics using EDM

This section contains the procedures to view statistics using EDM.

Enabling RMON statistics

Enable Ethernet statistics collection for RMON.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > Serviceability > RMON**.
2. Click **Control**.
3. Click the **Ethernet Statistics** tab.
4. Click **Insert**.
5. Next to the **Port** box, click the ellipsis (...) button.
6. Select a port.
7. Click **OK**.
8. In the **Owner** box, type the name of the owner entity.
9. Click **OK**.
10. Click **Insert**.

Ethernet Statistics field descriptions

Use the data in the following table to use the **Ethernet Statistics** tab.

Name	Description
Index	Uniquely identifies an entry in the Ethernet Statistics table. The default is 1.
Port	Identifies the source of the data that this etherStats entry is configured to analyze.
Owner	Specifies the entity that configured this entry and therefore uses the assigned resources. The default value is SNMP.

Disabling RMON statistics

Disable RMON statistics on a port after you do not need to gather statistics on that port.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Serviceability > RMON**.
2. Click **Control**.
3. Click the **Ethernet Statistics** tab.
4. Select the row that contains the port ID for which you must disable statistics.
5. Click **Delete**.

Viewing RMON statistics

Use the following procedure to view RMON statistics for each port.

Before you begin

- You must enable RMON statistics collection.

Procedure

1. On the Device Physical View, select a port.
2. In the navigation pane, expand the following folders: **Configuration > Graph**
3. Click **Port**.
4. Click the **RMON** tab.
5. Select the statistics you want to graph.
6. Select a graph type:
 - bar
 - pie
 - chart
 - line

RMON field descriptions

The following table describes fields on the **RMON** tab with second generation modules.

Name	Description
InOctets	<p>Specifies the number of subnetwork unicast packets delivery to the protocol.</p> <p>Specifies the number of octets of data (including those in bad packets) received on the network (excluding framing bits but including FCS octets).</p> <p>You can use this object as a reasonable estimate of Ethernet utilization. If additional precision is desired, sample the Pkts and Octets objects before and after a common interval. The differences in the sampled values are Pkts and Octets, and the number of seconds in the interval is Interval. These values are used to calculate the Utilization as follows:</p> $\text{Pkts} * (9.6+6.4) + (\text{Octets} * .8)$ <p>Utilization =</p> <p>Interval * 10,000</p> <p>The result of this equation is the value Utilization, which is the percent utilization of the Ethernet segment on a scale of 0 to 100 percent.</p>
OutOctets	Specifies the total number of octets transmitted out of the interface, including framing characters.
InUcastPkts	Specifies the number of subnetwork unicast packets delivered to the protocol.
OutUcastPkts	Specifies the total number of packets transmitted to a subnetwork unicast address, including those that were discarded or not sent.

Name	Description
InMulticastPkts	Specifies the number of packets delivered by this sublayer to a higher sublayer, which were addressed to a multicast address at this sub-layer. For a MAC layer protocol this includes both use and functional addresses.
OutMulticastPkts	Specifies the total number of packets that higher-level protocols requested be transmitted, and which were addressed to a multicast address at this sublayer, including those that were discarded or not sent. For a MAC layer protocol, this includes both group and functional addresses.
InBroadcastPkts	Specifies the number of packets delivered by this sublayer to a higher sublayer, which were addressed to a broadcast address at this sublayer.
OutBroadcastPkts	Specifies the total number of packets that higher-level protocols requested be transmitted, and which were addressed to a broadcast address at this sublayer, including those that were discarded or not sent. For a MAC layer protocol, this includes both group and functional addresses.
InDiscards	Specifies the number of inbound packets which were chosen to be discarded even though no errors had been detected to prevent their being deliverable to a higher layer protocol. One possible reason for discarding such a packet could be to free up buffer space.
InErrors	Specifies the number of inbound packets that contained errors preventing them from being deliverable to a higher level protocol.
InUnknownProtos	Specifies the number of packets received through the interface which were discarded because of an unknown or unsupported protocol.
InFlowCtrlPkts	Specifies the total number of flow control packets received by this interface.
OutFlowCtrlPkts	Specifies the total number of low control packets transmitted by this interface.
InPfcPkts	Specifies the total number of Priority Flow Control (PFC) packets received by this interface.
OutPfcPkts	Specifies the total number of Priority Flow Control (PFC) packets transmitted by this interface.

The following table describes fields on the **RMON** tab with first generation modules.

Name	Description
InOctets	<p>Specifies the number of subnetwork unicast packets delivery to the protocol.</p> <p>Specifies the number of octets of data (including those in bad packets) received on the network (excluding framing bits but including FCS octets).</p> <p>You can use this object as a reasonable estimate of Ethernet utilization. If additional precision is desired, sample the Pkts and Octets objects before and after a common interval. The differences in the sampled values are Pkts and Octets, and the number of seconds in the interval is Interval. These values are used to calculate the Utilization as follows:</p> $Pkts * (9.6+6.4) + (Octets * .8)$ <p>Utilization =</p> <p>Interval * 10,000</p>

Name	Description
	The result of this equation is the value Utilization, which is the percent utilization of the Ethernet segment on a scale of 0 to 100 percent.
OutOctets	Specifies the total number of octets transmitted out of the interface, including framing characters.
InUcastPkts	Specifies the number of subnetwork unicast packets delivered to the protocol.
OutUcastPkts	Specifies the total number of packets transmitted to a subnetwork unicast address, including those that were discarded or not sent.
InMulticastPkts	Specifies the number of packets delivered by this sublayer to a higher sublayer, which were addressed to a multicast address at this sub-layer. For a MAC layer protocol this includes both use and functional addresses.
OutMulticastPkts	Specifies the total number of packets that higher-level protocols requested be transmitted, and which were addressed to a multicast address at this sublayer, including those that were discarded or not sent. For a MAC layer protocol, this includes both group and functional addresses.
InBroadcastPkts	Specifies the number of packets delivered by this sublayer to a higher sublayer, which were addressed to a broadcast address at this sublayer.
OutBroadcastPkts	Specifies the total number of packets that higher-level protocols requested be transmitted, and which were addressed to a broadcast address at this sublayer, including those that were discarded or not sent. For a MAC layer protocol, this includes both group and functional addresses.
InDiscards	Specifies the number of inbound packets which were chosen to be discarded even though no errors had been detected to prevent their being deliverable to a higher layer protocol. One possible reason for discarding such a packet could be to free up buffer space.
InErrors	Specifies the number of inbound packets that contained errors preventing them from being deliverable to a higher level protocol.
InUnknownProtos	Specifies the number of packets received through the interface which were discarded because of an unknown or unsupported protocol.
HCInPfcPkts	Specifies the total number of Priority Flow Control (PFC) packets received by this interface.
HCOutPfcPkts	Specifies the total number of Priority Flow Control (PFC) packets transmitted by this interface.
HCInFlowCtrlPkts	Specifies the total number of flow control packets received by this interface.
HCOutFlowCtrlPkts	Specifies the total number of low control packets transmitted by this interface.

Graphing chassis statistics

Create graphs of chassis statistics to generate a visual representation of your data.

Procedure

1. On the Device Physical View, select the chassis.

2. In the navigation pane, expand the following folders: **Configuration** > **Graph**.
3. Click **Chassis**.
4. On the Graph Chassis tab, select the tab with the data you want to graph:
 - System
 - SNMP
 - IP
 - ICMP In
 - ICMP Out
 - TCP
 - UDP
 - Protocol Drop
5. Select the statistic you want to graph.
6. Select the graph type:
 - line chart
 - area chart
 - bar chart
 - pie chart

Graphing port statistics

You can create graphs for many port statistics to generate a visual representation of your data.

Procedure

1. On the Device Physical View, select the port or ports for which you want to create a graph.
2. Perform the following steps:
 - Right-click a port or multiple ports. On the shortcut menu, choose **Graph**.
 - In the navigation tree, expand the following folders: **Configuration** > **Graph**, and then click **Port**.
3. When the graph port dialog box appears, click the tab for which you want to graph the statistics.
4. Select the item for which you want to graph the statistics.
5. Select a graph type:
 - bar
 - pie
 - chart

- line

Viewing chassis system statistics

Use the following procedure to create graphs for chassis statistics.

Procedure

1. On the Device Physical View, select the chassis.
2. In the navigation pane, expand the following folders: **Configuration > Graph**.
3. Click **Chassis**.
4. Click the **System** tab.

System field descriptions

The following table describes the fields on the **System** tab.

Name	Description
MemUsed	The percentage of DRAM space used. Only the AbsoluteValue column is valid in the System tab. All other columns display as N/A because they are percentages and not actual memory counters.
MemFree	The amount in kilobytes of free DRAM.
CpuCurrentUtil	Percentage of CPU utilization.

Viewing chassis SNMP statistics

View chassis SNMP statistics to monitor network performance.

Procedure

1. On the Device Physical View, select the chassis.
2. In the navigation pane, expand the following folders: **Configuration > Graph**.
3. Click **Chassis**.
4. Click the **SNMP** tab.

SNMP field descriptions

The following table describes parameters on the **SNMP** tab.

Name	Description
InPkts	The number of messages delivered to the SNMP entity from the transport service.

Name	Description
OutPkts	The number of SNMP messages passed from the SNMP protocol entity to the transport service.
InTotalReqVars	The number of MIB objects retrieved successfully by the SNMP protocol entity as the result of receiving valid SNMP Get-Request and Get-Next PDUs.
InTotalSetVars	The number of MIB objects altered successfully by the SNMP protocol entity as the result of receiving valid SNMP Set-Request PDUs.
InGetRequests	The number of SNMP Get-Request PDUs the SNMP protocol accepts and processes.
InGetNexts	The number of SNMP Get-Next PDUs the SNMP protocol accepts and processes.
InSetRequests	The number of SNMP Set-Request PDUs the SNMP protocol accepts and processes.
InGetResponses	The number of SNMP Get-Response PDUs the SNMP protocol accepts and processes.
OutTraps	The number of SNMP Trap PDUs the SNMP protocol generates.
OutTooBig	The number of SNMP PDUs the SNMP protocol generates for which the value of the error-status field is tooBig.
OutNoSuchNames	The number of SNMP PDUs the SNMP protocol generates for which the value of the error-status field is noSuchName.
OutBadValues	The number of SNMP PDUs the SNMP protocol generates for which the value of the error-status field is badValue.
OutGenErrs	The number of SNMP PDUs the SNMP protocol generates for which the value of the error-status field is genErr.
InBadVersions	The number of SNMP messages delivered to the SNMP protocol entity for an unsupported SNMP version.
InBadCommunityNames	The number of SNMP messages delivered to the SNMP protocol entity that used an SNMP community name not known to said entity.
InBadCommunityUses	The number of SNMP messages delivered to the SNMP protocol entity that represented an SNMP operation not allowed by the SNMP community named in the message.
InASNParseErrs	The number of ASN.1 or BER errors the SNMP protocol encountered when decoding received SNMP messages.
InTooBig	The number of SNMP PDUs delivered to the SNMP protocol entity for which the value of the error-status field is tooBig.
InNoSuchNames	The number of SNMP PDUs delivered to the SNMP protocol entity for which the value of the error-status field is noSuchName.
InBadValues	The number of SNMP PDUs delivered to the SNMP protocol entity for which the value of the error-status field is badValue.
InReadOnly	The number of SNMP PDUs delivered to the SNMP protocol entity for which the value of the error-status field is readOnly. It is a protocol error to generate an SNMP PDU containing the value "readOnly" in the error-status

Name	Description
	field. This object is provided to detect incorrect implementations of the SNMP.
InGenErrs	The number of SNMP PDUs delivered to the SNMP protocol entity for which the value of the error-status field is genErr.

Viewing chassis IP statistics

View chassis IP statistics to monitor network performance.

Procedure

1. On the Device Physical View, select the chassis.
2. In the navigation pane, expand the following folders: **Configuration > Graph**.
3. Click **Chassis**.
4. Click the **IP** tab.

IP field descriptions

The following table describes parameters on the **IP** tab.

Name	Description
InReceives	The number of input datagrams received from interfaces, including those received in error.
InHdrErrors	The number of input datagrams discarded due to errors in the IP headers, including bad checksums, version number mismatch, other format errors, time-to-live exceeded, errors discovered in processing their IP options.
InAddrErrors	The number of input datagrams discarded because the IP address in the IP header destination field was not a valid address to be received at this entity. This count includes invalid addresses (for example, 0.0.0.0) and addresses of unsupported Classes (for example, Class E). For entities that are not IP Gateways and therefore do not forward datagrams, this counter includes datagrams discarded because the destination address was not a local address.
ForwDatagrams	The number of input datagrams for which this entity was not their final IP destination, as a result of which an attempt was made to find a route to forward them to that final destination. In entities that do not act as IP Gateways, this counter includes only those packets that were Source-Routed by way of this entity and had successful Source-Route option processing.
InUnknownProtos	The number of locally addressed datagrams received successfully but discarded because of an unknown or unsupported protocol.
InDiscards	The number of input IP datagrams for which no problems were encountered to prevent their continued processing but that were discarded

Name	Description
	(for example, for lack of buffer space). This counter does not include any datagrams discarded while awaiting reassembly.
InDelivers	The number of input datagrams successfully delivered to IP user-protocols (including ICMP).
OutRequests	The number of IP datagrams that local IP user-protocols (including ICMP) supplied to IP in requests for transmission. This counter does not include any datagrams counted in ipForwDatagrams.
OutDiscards	The number of output IP datagrams for which no problem was encountered to prevent their transmission to their destination, but that were discarded (for example, for lack of buffer space). This counter includes datagrams counted in ipForwDatagrams if any such packets met this (discretionary) discard criterion.
OutNoRoutes	The number of IP datagrams discarded because no route was found to transmit them to their destination. This counter includes any packets counted in ipForwDatagrams that meet this no-route criterion. This counter includes any datagrams a host cannot route because all default gateways are down.
FragOKs	The number of IP datagrams that were successfully fragmented at this entity.
FragFails	The number of IP datagrams that were discarded because they needed to be fragmented at this entity but can not be, for example, because the Don't Fragment flags were set.
FragCreates	The number of IP datagram fragments that were generated as a result of fragmentation at this entity.
ReasmReqds	The number of IP fragments received that needed to be reassembled at this entity.
ReasmOKs	The number of IP datagrams successfully reassembled.
ReasmFails	The number of failures detected by the IP reassembly algorithm (for whatever reason: timed out, errors, and so on). This number is not necessarily a count of discarded IP fragments because some algorithms (notably the algorithm in RFC 815) can lose track of the number of fragments by combining them as they are received.

Viewing chassis ICMP In statistics

View chassis ICMP In statistics to monitor network performance.

Procedure

1. On the Device Physical View, select the chassis.
2. In the navigation pane, expand the following folders: **Configuration > Graph**.
3. Click **Chassis**.
4. Click the **ICMP In** tab.

ICMP In field descriptions

The following table describes parameters on the **ICMP In** tab.

Name	Description
SrcQuenches	The number of ICMP Source Quench messages received.
Redirects	The number of ICMP Redirect messages received.
Echos	The number of ICMP Echo (request) messages received.
EchoReps	The number of ICMP Echo Reply messages received.
Timestamps	The number of ICMP Timestamp (request) messages received.
TimestampReps	The number of ICMP Timestamp Reply messages received.
AddrMasks	The number of ICMP Address Mask Request messages received.
AddrMaskReps	The number of ICMP Address Mask Reply messages received.
ParmProbs	The number of ICMP Parameter Problem messages received.
DestUnreachs	The number of ICMP Destination Unreachable messages received.
TimeExcds	The number of ICMP Time Exceeded messages received.

Viewing chassis ICMP Out statistics

View chassis ICMP Out statistics to monitor network performance.

Procedure

1. On the Device Physical View, select the chassis.
2. In the navigation pane, expand the following folders: **Configuration > Graph**.
3. Click **Chassis**.
4. Click the **ICMP Out** tab.

ICMP Out field descriptions

The following table describes parameters on the **ICMP Out** tab.

Name	Description
SrcQuenches	The number of ICMP Source Quench messages sent.
Redirects	The number of ICMP Redirect messages received. For a host, this object is always zero, because hosts do not send redirects.
Echos	The number of ICMP Echo (request) messages sent.
EchoReps	The number of ICMP Echo Reply messages sent.
Timestamps	The number of ICMP Timestamp (request) messages sent.
TimestampReps	The number of ICMP Timestamp Reply messages sent.
AddrMasks	The number of ICMP Address Mask Request messages sent.

Name	Description
AddrMaskReps	The number of ICMP Address Mask Reply messages sent.
ParmProbs	The number of ICMP Parameter Problem messages sent.
DestUnreachs	The number of ICMP Destination Unreachable messages sent.
TimeExcds	The number of ICMP Time Exceeded messages sent.

Viewing ICMP statistics

View ICMP statistics for ICMP configuration information.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > IPv6**.
2. Click **IPv6**.
3. Click **Interfaces** tab.
4. Select the interface on which you want to view the ICMP statistics.
5. Click **ICMPstats** option from the menu.

ICMP stats field descriptions

Use the data in the following table to use the ICMP **Statistics** tab.

Name	Description
InMsgs	Specifies the total number of ICMP messages which the entity received. Note: This counter includes all those counted by icmpInErrors.
InErrors	Specifies the number of ICMP messages which the entity received but determined as having ICMP-specific errors (bad ICMP checksums, bad length, etc.).
InDestUnreachs	Specifies the number of ICMP Destination Unreachable messages received by the interface.
InAdminProhibs	Specifies the number of ICMP destination unreachable/communication administratively prohibited messages received by the interface.
InTimeExcds	Specifies the number of ICMP Time Exceeded messages by the interface.
InParmProblems	Specifies the number of ICMP Parameter Problem messages received by the interface.

Name	Description
InPktTooBig	Specifies the number of ICMP Packet Too Big messages received by the interface.
InEchos	Specifies the number of ICMP Echo (request) messages received by the interface.
InEchoReplies	Specifies the number of ICMP Echo Reply messages received by the interface.
InRouterSolicits	Specifies the number of ICMP Router Solicit messages received by the interface.
InRouterAdvertisements	Specifies the number of ICMP Router Advertisement messages received by the interface
InNeighborSolicits	Specifies the number of ICMP Neighbor Solicit messages received by the interface.
InNeighborAdvertisements	Specifies the number of ICMP Neighbor Advertisement messages received by the interface.
InRedirects	Specifies the number of ICMP Redirect messages received by the interface.
InGroupMembQueries	Specifies the number of ICMPv6 Group Membership Query messages received by the interface
InGroupMembResponses	Specifies the number of ICPv6 Group Membership Response messages received by the interface.
InGroupMembReductions	Specifies the number of ICMPv6 Group Membership Reduction messages received by the interface.
OutMsgs	Specifies the total number of ICMP messages which this interface attempted to send. Note that this counter includes all those counted by icmpOutErrors.
OutErrors	Specifies the number of ICMP messages which this interface did not send due to problems discovered within ICMP such as a lack of buffers. This value should not include errors discovered outside the ICMP layer such as the inability of IPv6 to route the resultant datagram. In some implementations there may be no types of error which contribute to this counter's value.
OutDestUnreachs	Specifies the number of ICMP Destination Unreachable messages sent by the interface.
OutAdminProhibs	Specifies the number of ICMP dest unreachable/ communication administratively prohibited messages sent.
OutTimeExclds	Specifies the number of ICMP Time Exceeded messages sent by the interface.
OutParmProblems	Specifies the number of ICMP Parameter Problem messages sent by the interface.

Name	Description
OutPktTooBigs	Specifies the number of ICMP Packet Too Big messages sent by the interface.
OutEchos	Specifies the number of ICMP Echo (request) messages sent by the interface.
OutEchoReplies	Specifies the number of ICMP Echo Reply messages sent by the interface.
OutRouterSolicits	Specifies the number of ICMP Router Solicitation messages sent by the interface.
OutRouterAdvertisements	Specifies the number of ICMP Router Advertisement messages sent by the interface.
OutNeighborSolicits	Specifies the number of ICMP Neighbor Solicitation messages sent by the interface.
OutNeighborAdvertisements	Specifies the number of ICMP Neighbor Advertisement messages sent by the interface.
OutRedirects	Specifies the number of Redirect messages sent. For a host, this object will always be zero, since hosts do not send redirects.
OutGroupMembQueries	Specifies the number of ICMPv6 Group Membership Query messages sent.
OutGroupMembResponses	Specifies the number of ICMPv6 Group Membership Response messages sent.
OutGroupMembReductions	Specifies the number of ICMPv6 Group Membership Reduction messages sent.

Viewing chassis TCP statistics

View TCP statistics to monitor network performance.

Procedure

1. On the Device Physical View, select the chassis.
2. In the navigation pane, expand the following folders: **Configuration > Graph**.
3. Click **Chassis**.
4. Click the **TCP** tab.

TCP field descriptions

The following table describes parameters on the **TCP** tab.

Name	Description
ActiveOpens	The number of times TCP connections made a direct transition to the SYN-SENT state from the CLOSED state.

Name	Description
PassiveOpens	The number of times TCP connections made a direct transition to the SYN-RCVD state from the LISTEN state.
AttemptFails	The number of times TCP connections made a direct transition to the CLOSED state from either the SYN-SENT state or the SYN-RCVD state, plus the number of times TCP connections made a direct transition to the LISTEN state from the SYN-RCVD state.
EstabResets	The number of times TCP connections made a direct transition to the CLOSED state from either the ESTABLISHED state or the CLOSE-WAIT state.
CurrEstab	The number of TCP connections for which the current state is either ESTABLISHED or CLOSE-WAIT.
InSegs	The number of segments received, including those received in error. This count includes segments received on currently established connections.
OutSegs	The number of segments sent, including those on current connections, but excluding those containing only retransmitted octets.
RetransSegs	The number of segments retransmitted that is, the number of TCP segments transmitted containing one or more previously transmitted octets.
InErrs	The number of segments received in error (for example, bad TCP checksums).
OutRsts	The number of TCP segments sent containing the RST flag.
HCInSegs	The number of segments received, including those received in error. This count includes segments received on currently established connections. This object is the 64-bit equivalent of InSegs.
HCOutSegs	The number of segments sent, including those on current connections, but excluding those containing only retransmitted octets. This object is the 64-bit equivalent of OutSegs.

Viewing chassis UDP statistics

Display User Datagram Protocol (UDP) statistics to see information about the UDP datagrams.

Procedure

1. On the Device Physical View, select the chassis.
2. In the navigation pane, expand the following folders: **Configuration > Graph**.
3. Click **Chassis**.
4. Click the **UDP** tab.
5. Select the information you want to graph.
6. Select the type of graph you want:
 - line
 - area

- bar
- pie

7. To clear counters, click **Clear Counters**. Discontinuities in the value of these counters can occur when the management system reinitializes, and at other times as indicated by discontinuities in the value of sysUpTime.

UDP field descriptions

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

Name	Description
NoPorts	<p>The number of received UDP datagrams with no application at the destination port.</p> <p>Discontinuities in the value of this counter can occur at reinitialization of the management system, and at other times as indicated by discontinuities in the value of sysUpTime.</p>
InErrors	<p>The number of received UDP datagrams that were not delivered for reasons other than the lack of an application at the destination port.</p> <p>Discontinuities in the value of this counter can occur at reinitialization of the management system and at other times as indicated by discontinuities in the value of sysUpTime.</p>
InDatagrams	<p>The number of UDP datagrams delivered to UDP users, for devices that can receive more than 1 000 000 UDP datagrams for each second.</p> <p>Discontinuities in the value of this counter can occur at reinitialization of the management system, and at other times as indicated by discontinuities in the value of sysUpTime.</p>
OutDatagrams	<p>The number of UDP datagrams sent from this entity.</p> <p>Discontinuities in the value of this counter can occur at reinitialization of the management system, and at other times as indicated by discontinuities in the value of sysUpTime.</p>
HCInDatagrams	<p>The number of TCP connections for which the current state is either ESTABLISHED or CLOSE-WAIT.</p>
HCOutDatagrams	<p>The number of UDP datagrams sent from this entity, for devices that can transmit more than 1 million UDP datagrams for each second.</p> <p>Discontinuities in the value of this counter can occur at reinitialization of the management system, and at other times as indicated by discontinuities in the value of sysUpTime.</p>

Configuring Switch Fabric statistics capture

Configure the statistic settings for the Switch Fabric modules to determine what to capture.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > Edit**.
2. Click **Switch Fabric**.
3. Click the **Switch Fabric Stats Settings** tab.
4. Specify the Class of Service in the **StatsCosId** box.
5. Select the port in the **StatsPortId** box.
6. Select **StatsCapture**.
7. Click **Apply**.

Switch Fabric Stats Settings field descriptions

The following table describes the variable on the **Switch Fabric Stats Settings** tab.

Name	Description
StatsCosId	Specifies the Class of Service on which to collect statistics.
StatsPortId	Specifies the port on which to collect statistics. You must select a data port.
StatsCapture	Turns statistics collection on or off. The default is off.
DeviceRead	Collects statistics for Switch Fabric counters directly from devices. If you clear this variable, a cached copy is returned. The default is selected (enabled).

Viewing Switch Fabric statistics

View statistics for the Switch Fabric modules to manage network performance.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > Edit**
2. Click **Switch Fabric**.
3. Click the **Switch Fabric Stats** tab.

Switch Fabric Stats field descriptions

The following table describes the statistics captured on the **Switch Fabric Stats** tab.

Name	Description
StatsPortId	Shows the data port on which statistics are collected.
DropPrec1AcceptedPackets	Shows the accepted number of packets from drop precedence 1. The drop precedence is a function of the front-end policer.

Name	Description
DropPrec1AcceptedBytes	Shows the accepted number of bytes from drop precedence 1.
DropPrec1CongestionMarkedPackets	Shows the congestion marked number of packets from drop precedence 1.
DropPrec1CongestionMarkedBytes	Shows the congestion marked number of bytes from drop precedence 1.
DropPrec1DiscardDroppedPackets	Shows the number of WRED dropped packets from drop precedence 1.
DropPrec1DiscardDroppedBytes	Shows the number of WRED dropped bytes from drop precedence 1.
DropPrec2AcceptedPackets	Shows the accepted number of packets from drop precedence 2. The drop precedence is a function of the front-end policer.
DropPrec2AcceptedBytes	Shows the accepted number of bytes from drop precedence 2.
DropPrec2CongestionMarkedPackets	Shows the congestion marked number of packets from drop precedence 2.
DropPrec2CongestionMarkedBytes	Shows the congestion marked number of bytes from drop precedence 2.
DropPrec2DiscardDroppedPackets	Shows the number of WRED dropped packets from drop precedence 2.
DropPrec2DiscardDroppedBytes	Shows the number of WRED dropped bytes from drop precedence 2.
DropPrec3AcceptedPackets	Shows the accepted number of packets from drop precedence 3. The drop precedence is a function of the front-end policer.
DropPrec3AcceptedBytes	Shows the accepted number of bytes from drop precedence 3.
DropPrec3CongestionMarkedPackets	Shows the congestion marked number of packets from drop precedence 3.
DropPrec3CongestionMarkedBytes	Shows the congestion marked number of bytes from drop precedence 3.
DropPrec3DiscardDroppedPackets	Shows the number of WRED dropped packets from drop precedence 3.
DropPrec3DiscardDroppedBytes	Shows the number of WRED dropped bytes from drop precedence 3.
DropPrec4AcceptedPackets	Shows the accepted number of packets from drop precedence 4. The drop precedence is a function of the front-end policer.
DropPrec4AcceptedBytes	Shows the accepted number of bytes from drop precedence 4.

Name	Description
NonWredDroppedPackets	Shows the number of dropped packets due to non Random Early Detection.
NonWredDroppedBytes	Shows the number of dropped bytes due to non Random Early Detection.
DequeuedPackets	Shows the number of packets dequeued once inside the Switch Fabric.
DequeuedBytes	Shows the number of bytes dequeued once inside the Switch Fabric.
DropPrec1DroppedPackets	Shows the numbers of non-WRED dropped packets from drop precedence 1.
DropPrec1DroppedBytes	Shows the number of non-WRED dropped bytes from drop precedence 1.
DropPrec2DroppedPackets	Shows the numbers of non-WRED dropped packets from drop precedence 2.
DropPrec2DroppedBytes	Shows the number of non-WRED dropped bytes from drop precedence 2.
DropPrec3DroppedPackets	Shows the numbers of non-WRED dropped packets from drop precedence 3.
DropPrec3DroppedBytes	Shows the number of non-WRED dropped bytes from drop precedence 3.
DropPrec4CongestionMarkedPackets	Shows the congestion marked number of packets from drop precedence 4.
DropPrec4CongestionMarkedBytes	Shows the congestion marked number of bytes from drop precedence 4.
DropPrec4DiscardDroppedPackets	Shows the number of WRED dropped packets from drop precedence 4.
DropPrec4DiscardDroppedBytes	Shows the number of WRED dropped bytes from drop precedence 4.
DropPrec4DroppedPackets	Shows the numbers of non-WRED dropped packets from drop precedence 4.
DropPrec4DroppedBytes	Shows the number of non-WRED dropped bytes from drop precedence 4.
OverSubscribeTotalDroppedPkts	Shows the number of dropped packets due to free list underflow.
OverSubscribeTotalDroppedBytes	Shows the number of dropped bytes due to free list underflow.
OverSubscribeGuaranteeDroppedPkts	Shows the number of dropped packets due to the global buffer threshold guarantee.
OverSubscribeGuaranteeDroppedBytes	Shows the number of dropped bytes due to the global buffer threshold guarantee.
OutPkts	Shows the number of packets out from the egress interface.

Name	Description
OutBytes	Shows the number of bytes out from the egress interface.

Viewing port interface statistics

View port interface statistics to manage network performance.

Procedure

1. On the Device Physical View, select a port.
2. In the navigation pane, expand the following folders: **Configuration > Graph**.
3. Click **Port**.
4. Click the **Interface** tab.

Interface field descriptions

The following table describes parameters for second generation modules on the **Interface** tab.

Name	Description
InOctets	Specifies the number of octets received on the interface, including framing characters.
OutOctets	Specifies the number of octets transmitted from the interface, including framing characters.
InUcastPkts	Specifies the number of packets delivered by this sublayer to a higher sublayer that were not addressed to a multicast or broadcast address at this sublayer.
OutUcastPkts	Specifies the number of packets that higher-level protocols requested be transmitted that were not addressed to a multicast address at this sublayer. The total number includes those packets discarded or not sent.
InMulticastPkts	Specifies the number of packets delivered by this sublayer to a higher sublayer that were addressed to a multicast address at this sublayer. For a MAC layer protocol, this number includes both group and functional addresses.
OutMulticastPkts	Specifies the number of packets that higher-level protocols requested be transmitted, and that are addressed to a multicast address at this sublayer, including those that were discarded or not sent. For a MAC layer protocol, this number includes both group and functional addresses.
InBroadcastPkts	Specifies the number of packets delivered by this sublayer to a higher sublayer that are addressed to a broadcast address at this sublayer.
OutBroadcastPkts	Specifies the number of packets that higher-level protocols requested be transmitted, and that were addressed to a broadcast address at this sublayer, including those that were discarded or not sent.

Name	Description
InDiscards	Specifies the number of inbound packets that are discarded because of frames with errors or invalid frames or, in some cases, to fill up buffer space.
InErrors	For packet-oriented interfaces, specifies the number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol. For character-oriented or fixed-length interfaces, the number of inbound transmission units that contained errors preventing them from being deliverable to a higher-layer protocol.
InUnknownProtos	For packet-oriented interfaces, specifies the number of packets received through the interface that are discarded because of an unknown or unsupported protocol. For character-oriented or fixed-length interfaces that support protocol multiplexing, the number of transmission units received through the interface that were discarded because of an unknown or unsupported protocol. For any interface that does not support protocol multiplexing, this counter is always 0.
InFlowCtrlPkts	Specifies the number of flow control packets received by this interface.
OutFlowCtrlPkts	Specifies the number of flow control packets transmitted by this interface.
InPfcPkts	Specifies the total number of Priority Flow Control (PFC) packets received by this interface.
OutPfcPkts	Specifies the total number of Priority Flow Control (PFC) packets transmitted by this interface.

The following table describes parameters for first generation modules on the **Interface** tab.

Name	Description
InOctets	Specifies the number of octets received on the interface, including framing characters.
OutOctets	Specifies the number of octets transmitted from the interface, including framing characters.
InUcastPkts	Specifies the number of packets delivered by this sublayer to a higher sublayer that were not addressed to a multicast or broadcast address at this sublayer.
OutUcastPkts	Specifies the number of packets that higher-level protocols requested be transmitted that were not addressed to a multicast address at this sublayer. The total number includes those packets discarded or not sent.
InMulticastPkts	Specifies the number of packets delivered by this sublayer to a higher sublayer that were addressed to a multicast address at this sublayer. For a MAC layer protocol, this number includes both group and functional addresses.
OutMulticastPkts	Specifies the number of packets that higher-level protocols requested be transmitted, and that are addressed to a multicast address at this sublayer, including those that were discarded or not sent. For a MAC

Name	Description
	layer protocol, this number includes both group and functional addresses.
InBroadcastPkts	Specifies the number of packets delivered by this sublayer to a higher sublayer that are addressed to a broadcast address at this sublayer.
OutBroadcastPkts	Specifies the number of packets that higher-level protocols requested be transmitted, and that were addressed to a broadcast address at this sublayer, including those that were discarded or not sent.
InDiscards	Specifies the number of inbound packets that are discarded because of frames with errors or invalid frames or, in some cases, to fill up buffer space.
InErrors	For packet-oriented interfaces, specifies the number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol. For character-oriented or fixed-length interfaces, the number of inbound transmission units that contained errors preventing them from being deliverable to a higher-layer protocol.
InUnknownProtos	For packet-oriented interfaces, specifies the number of packets received through the interface that are discarded because of an unknown or unsupported protocol. For character-oriented or fixed-length interfaces that support protocol multiplexing, the number of transmission units received through the interface that were discarded because of an unknown or unsupported protocol. For any interface that does not support protocol multiplexing, this counter is always 0.
HCInPfcPkts	Specifies the total number of Priority Flow Control (PFC) packets received by this interface.
HCOutPfcPkts	Specifies the total number of Priority Flow Control (PFC) packets transmitted by this interface.
HCInFlowCtrlPkts	Specifies the number of flow control packets received by this interface.
HCOutFlowCtrlPkts	Specifies the number of flow control packets transmitted by this interface.
NumStateTransition	Specifies the number of times the port went in and out of service. The number of state transitions from up to down.

Viewing port Ethernet errors statistics

View port Ethernet errors statistics to manage network performance.

Procedure

1. On the Device Physical View, select a port.
2. In the navigation pane, expand the following folders: **Configuration > Graph**.
3. Click **Port**.
4. Click the **Ethernet Errors** tab.

Ethernet Errors field descriptions

The following table describes parameters on the **Ethernet Errors** tab.

Name	Description
AlignmentErrors	Specifies account of frames received on a particular interface that are not an integral number of octets in length and do not pass the FCS check. The count represented by an instance of this object increments when the alignmentError status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions obtain are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.
FCSErrors	Specifies a count of frames received on a particular interface that are an integral number of octets in length but do not pass the FCS check. The count represented by an instance of this object increments when the frameCheckError status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions obtained are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.
InternalMacTransmitErrors	Specifies a count of frames for which transmission on a particular interface fails due to an internal MAC sublayer transmit error. A frame is only counted by an instance of this object if it is not counted by the corresponding instance of either the LateCollisions object, the ExcessiveCollisions object, or the CarrierSenseErrors object. The precise meaning of the count represented by an instance of this object is implementation-specific. In particular, an instance of this object can represent a count of transmission errors on a particular interface that are not otherwise counted.
InternalMacReceiveErrors	Specifies a count of frames for which reception on a particular interface fails due to an internal MAC sublayer receive error. A frame is only counted by an instance of this object if it is not counted by the corresponding instance of either the FrameTooLongs object, the AlignmentErrors object, or the FCSErrors object. The precise meaning of the count represented by an instance of this object is implementation-specific. In particular, an instance of this object can represent a count of receive errors on a particular interface that are not otherwise counted.
CarrierSenseErrors	Specifies the number of times that the carrier sense condition is lost or not asserted when the switch attempts to transmit a frame on a particular interface. The count represented by an instance of this object increments at most once for each transmission attempt, even if the carrier sense condition fluctuates during a transmission attempt.
FrameTooLongs	Specifies a count of frames received on a particular interface that exceed the maximum permitted frame size. The count represented by an instance of this object increments when the frameTooLong status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions obtained are,

Name	Description
	according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.
SQETestErrors	Specifies a count of times that the SQE TEST ERROR message is generated by the PLS sublayer for a particular interface. The SQE TEST ERROR message is defined in section 7.2.2.2.4 of ANSI/IEEE 802.3-1985 and its generation described in section 7.2.4.6 of the same document.
DeferredTransmissions	Specifies a count of frames for which the first transmission attempt on a particular interface is delayed because the medium is busy. The count represented by an instance of this object does not include frames involved in collisions.
SingleCollisionFrames	Specifies a count of successfully transmitted frames on a particular interface for which transmission is inhibited by exactly one collision. A frame that is counted by an instance of this object is also counted by the corresponding instance of either the UcastPkts, MulticastPkts, or BroadcastPkts objects and is not counted by the corresponding instance of the MultipleCollisionFrames object.
MultipleCollisionFrames	Specifies a count of successfully transmitted frames on a particular interface for which transmission is inhibited by more than one collision. A frame that is counted by an instance of this object is also counted by the corresponding instance of either the UcastPkts, MulticastPkts, or BroadcastPkts objects and is not counted by the corresponding instance of the SingleCollisionFrames object.
LateCollisions	Specifies the number of times that a collision is detected on a particular interface later than 512 bit-times into the transmission of a packet; 512 corresponds to 51.2 microseconds on a 10 Mb/s system. A (late) collision included in a count represented by an instance of this object is also considered as a (generic) collision for purposes of other collision-related statistics.
ExcessiveCollisions	Specifies a count of frames for which transmission on a particular interface fails due to excessive collisions.
FrameTooShorts	Specifies the number of frames, encountered on this interface, that are too short.
LinkFailures	Specifies the number of link failures encountered on this interface.
PacketErrors	Specifies the number of packet errors encountered on this interface.
CarrierErrors	Specifies the number of carrier errors encountered on this interface.
LinkInactiveErrors	Specifies the number of link inactive errors encountered on this interface.

Viewing port bridging statistics

View port bridging errors statistics to manage network performance.

Procedure

1. On the Device Physical View, select a port.
2. In the navigation pane, expand the following folders: **Configuration > Graph**.
3. Click **Port**.
4. Click the **Bridging** tab.

Bridging field descriptions

The following table describes parameters on the **Bridging** tab.

Name	Description
InUnicastFrames	The number of incoming unicast frames bridged.
InMulticastFrames	The number of incoming multicast frames bridged.
InBroadcastFrames	The number of incoming broadcast frames bridged.
InDiscards	The number of frames discarded by the bridging entity.
OutFrames	The number of outgoing frames bridged.

Viewing port spanning tree statistics

View port spanning tree statistics to manage network performance.

Procedure

1. On the Device Physical View, select a port.
2. In the navigation pane, expand the following folders: **Configuration > Graph**.
3. Click **Port**.
4. Click the **Spanning Tree** tab.

Spanning Tree field descriptions

The following table describes parameters on the **Spanning Tree** tab.

Name	Description
InConfigBpdus	The number of Config BPDUs received.
InTcnBpdus	The number of Topology Change Notifications BPDUs received.
InBadBpdus	The number of unknown or malformed BPDUs received.
OutConfigBpdus	The number of Config BPDUs transmitted.
OutTcnBpdus	The number of Topology Change Notifications BPDUs transmitted.

Viewing port routing statistics

View port routing statistics to manage network performance.

Procedure

1. On the Device Physical View, select a port.
2. In the navigation pane, expand the following folders: **Configuration > Graph**.
3. Click **Port**.
4. Click the **Routing** tab.

Routing field descriptions

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

Name	Description
InUnicastFrames	The number of incoming unicast frames routed.
InMulticastFrames	The number of incoming multicast frames routed.
InDiscards	The number of frames discarded by the routing entity.
OutUnicastFrames	The number of outgoing unicast frames routed.
OutMulticastFrames	The number of outgoing multicast frames routed.

Viewing IPv6 statistics for an interface

View IPv6 statistics to view information about the IPv6 datagrams on an interface.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > IPv6**.
2. Click **IPv6**.
3. Click the **Interfaces** tab.
4. Select an interface.
5. Click **IfStats**.
6. **(Optional)** Select one or more values, and then click on the type of graph to graph the data.

Statistics field descriptions

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

Name	Description
InReceives	Shows the total number of input datagrams received by the interface, including those received in error.

Name	Description
InHdrErrors	Shows the number of input datagrams discarded due to errors in their IPv6 headers, including version number mismatch, other format errors, hop count exceeded, and errors discovered in processing the IPv6 options.
InTooBigErrors	Shows the number of input datagrams that could not be forwarded because their size exceeded the link MTU of the outgoing interface.
InNoRoutes	Shows the number of input datagrams discarded because no route could be found to transmit them to their destination.
InAddrErrors	Shows the number of input datagrams discarded because the IPv6 address in the IPv6 header destination field was not a valid address to be received at this entity. This count includes invalid addresses, for example, ::0, and unsupported addresses, for example, addresses with unallocated prefixes. For entities which are not IPv6 routers and do not forward datagrams, this counter includes datagrams discarded because the destination address was not a local address.
InUnknownProtos	Shows the number of locally-addressed datagrams received successfully but discarded because of an unknown or unsupported protocol. This counter is incremented at the interface to which these datagrams were addressed, which is not always the input interface for some of the datagrams.
InTruncatedPkts	Shows the number of input datagrams discarded because the datagram frame did not carry enough data.
InDiscards	Shows the number of input IPv6 datagrams for which no problems were encountered to prevent their continued processing, but which were discarded, for example, for lack of buffer space. This counter does not include datagrams discarded while awaiting re-assembly.
InDelivers	Shows the total number of datagrams successfully delivered to IPv6 user-protocols (including ICMP). This counter is incremented at the interface to which these datagrams were addressed which is not always the input interface for some of the datagrams.
OutForwDatagrams	Shows the number of output datagrams which this entity received and forwarded to their final destinations. In entities which do not act as IPv6 routers, this counter includes only those packets which were Source-Routed using this entity, and the

Name	Description
	Source-Route processing was successful. For a successfully forwarded datagram the counter of the outgoing interface is incremented.
OutRequests	Shows the total number of IPv6 datagrams which local IPv6 user-protocols (including ICMP) supplied to IPv6 in requests for transmission. This counter does not include datagrams counted in OutForwDatagrams .
OutDiscards	Shows the number of output IPv6 datagrams for which no problem was encountered to prevent their transmission to their destination, but which were discarded, for example , for lack of buffer space. This counter includes datagrams counted in OutForwDatagrams if such packets met this (discretionary) discard criterion.
OutFragOKs	Shows the number of IPv6 datagrams that have been successfully fragmented at this output interface.
OutFragFails	Shows the number of IPv6 datagrams that have been discarded because they needed to be fragmented at this output interface but could not be.
OutFragCreates	Shows the number of output datagram fragments that have been generated as a result of fragmentation at this output interface.
ReasmReqds	Shows the number of IPv6 fragments received which needed to be reassembled at this interface. This counter is incremented at the interface to which these fragments were addressed, which is not always the input interface for some of the fragments.
ReasmOKs	Shows the number of IPv6 datagrams successfully reassembled. This counter is incremented at the interface to which these datagrams were addressed, which is not always the input interface for some of the fragments.
ReasmFails	Shows the number of failures detected by the IPv6 re- assembly algorithm). This value is not necessarily a count of discarded IPv6 fragments because some algorithms can lose track of the number of fragments by combining them as they are received. This counter is incremented at the interface to which these fragments were addressed, which is not always the input interface for some of the fragments.
InMcastPkts	Shows the number of multicast packets received by the interface.
OutMcastPkts	Shows the number of multicast packets transmitted by the interface.

Viewing DHCP statistics for an interface

View DHCP statistics to manage network performance.

Procedure

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

Interfaces Stats field descriptions

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

Name	Description
IfIndex	Identifies the physical interface.
AgentAddr	Shows the IP address configured as the relay on this interface. This address is either the IP of the physical interface or the IP of the VRRP address.
NumRequests	Shows the number of DHCP and BootP requests on this interface.
NumReplies	Shows the number of DHCP and BootP replies on this interface.

Viewing IPv6 DHCP Relay statistics for a port

Display individual IPv6 DHCP Relay statistics for specific ports to manage network performance. You can also create a graph of selected statistical values.

Procedure

1. On the Device Physical view, select a port.
2. In the navigation pane, expand the following folders: **Configuration > IPv6**
3. Click the **DHCP Relay** tab.
4. Click the **Interface** tab.
5. Select the interface on which you want to view the IPv6 DHCP Relay statistics.
6. Click **Statistics**.
7. Select one or more values.
8. Click the type of graph.

Statistics field descriptions

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

Name	Description
NumRequests	Shows the number of DHCP and BootP requests on this interface.
NumReplies	Shows the number of DHCP and BootP replies on this interface.

Graphing DHCP statistics for a port

View DHCP statistics to manage network performance.

Procedure

1. On the Device Physical View, select a port.
2. In the navigation pane, expand the following folders: **Configuration > Graph**.
3. Click **Port**.
4. Click the **DHCP** tab.
5. Select one or more values.
6. Click the type of graph to create.

DHCP field descriptions

The following table describes parameters on the **DHCP** tab.

Name	Description
NumRequests	The number of DHCP and/or BootP requests on this interface.
NumReplies	The number of DHCP and/or BootP replies on this interface.

Viewing DHCP statistics for a port

View DHCP statistics to manage network performance.

Procedure

1. On the Device Physical view, select a port.
2. In the navigation pane, open the following folders: **Configuration > Edit > Port**
3. Click **IP**.
4. Click the **DHCP Relay** tab.
5. Click **Graph**.
6. Select one or more values.
7. Click the type of graph.

DHCP Stats field descriptions

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

Name	Description
NumRequests	The number of DHCP and BootP requests on this interface.
NumReplies	The number of DHCP and BootP replies on this interface.

Graphing DHCP statistics for a VLAN

View DHCP statistics to manage network performance.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > VLAN**
2. Click **VLANs**.
3. On the **Basic** tab, select a VLAN.
4. Click **IP**.
5. Click the **DHCP Relay** tab.
6. Click **Graph**.
7. Select one or more values.
8. Click the type of graph.

DHCP Stats field descriptions

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

Name	Description
NumRequests	The number of DHCP and BootP requests on this interface.
NumReplies	The number of DHCP and BootP replies on this interface.

Displaying DHCP-relay statistics for Option 82

Display DHCP-relay statistics for all interfaces to manage network performance.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > IP**.
2. Click **DHCP-Relay**.
3. Click the **Option 82 Stats** tab.

Option 82 Stats field descriptions

Use the data in the following table to use the **Option 82 Stats** tab.

Name	Description
IfIndex	Shows the name of the interface on which you enabled option 82. Shows the port number if the interface is a brouter port or the VLAN number if the interface is a VLAN.
AgentAddr	Shows the IP address configured as the relay on this interface. This address is either the IP of the physical interface or the IP of the VRRP address.
FoundOp82	Shows the number of packets that the interface received that already had option82 in them.
Dropped	Shows the number of packets the interface dropped because of option 82–related issues. These reasons could be that the packet was received from an untrusted source or spoofing was detected. To determine the cause of the drop, you must enable trace on level 170.
CircuitId	Shows the value inserted in the packets as the circuit ID. The value is the index of the interface.
AddedCircuitId	Shows how many packets (requests from client to server) the circuit ID was inserted for that interface. If you expect this value to increase but it does not, and the interface does not drop a packet, it is possible the packet does not have enough space to insert the option. You must enable trace on level 170 to determine the cause.
RemovedCircuitId	Shows how many packets (replies from server to client) the circuit id was removed for that interface.
Remoteld	Shows the value inserted in the packets as the remote ID. The value is the MAC address of the interface.
AddedRemoteld	Shows how many packets (requests from client to server) the remote ID was inserted for that interface. If you expect this value to increase but it does not, and the interface does not drop a packet, it is possible the packet does not have enough space to insert the option. You must enable trace on level 170 to determine the cause.
RemovedRemoteld	Shows how many packets (replies from server to client) the remote ID was removed for that interface.

Viewing port OSPF statistics

View port OSPF statistics to manage network performance.

Procedure

1. On the Device Physical View, select a port.
2. In the navigation pane, expand the following folders: **Configuration > Graph**.
3. Click **Port**.
4. Click the **OSPF** tab.

OSPF field descriptions

The following table describes parameters on the **OSPF** tab.

Name	Description
VersionMismatches	Specifies the number of version mismatches received by this interface.
AreaMismatches	Specifies the number of area mismatches received by this interface.
AuthTypeMismatches	Specifies the number of authentication type mismatches received by this interface.
AuthFailures	Specifies the number of authentication failures.
NetmaskMismatches	Specifies the number of net mask mismatches received by this interface.
HelloIntervalMismatches	Specifies the number of hello interval mismatches received by this interface.
DeadIntervalMismatches	Specifies the number of dead interval mismatches received by this interface.
OptionMismatches	Specifies the number of option mismatches in the hello interval or dead interval fields received by this interface.
RxHellos	Specifies the number of hello packets received by this interface.
RxDBDescrs	Specifies the number of database descriptor packets received by this interface.
RxLSUpdates	Specifies the number of link state update packets received by this interface.
RxLSReqs	Specifies the number of link state request packets received by this interface.
RxLSAcks	Specifies the number of link state acknowledge packets received by this interface.
TxHellos	Specifies the number of hello packets transmitted by this interface.

Name	Description
TxDBDescrs	Specifies the number of database descriptor packets transmitted by this interface.
TxLSUpdates	Specifies the number of link state update packets transmitted by this interface.
TxLSReqs	Specifies the number of link state request packets transmitted by this interface.
TxLSAcks	Specifies the number of link state acknowledge packets transmitted by this interface.

Viewing LACP port statistics

View LACP port statistics to monitor the performance of the port.

Procedure

1. On the Device Physical View, select a port.
2. In the navigation pane, expand the following folders: **Configuration > Graph**.
3. Click **Port**.
4. Click the **LACP** tab.
5. To change the poll interval, in the toolbar click the **Poll Interval** box, and then select a new interval.

LACP field descriptions

Use the data in the following table to view the LACP statistics.

Name	Description
LACPDUsRx	The number of valid LACPDU received on this aggregation port.
MarkerPDUsRx	The number of valid marker PDUs received on this aggregation port.
MarkerResponsePDUsRx	The number of valid marker response PDUs received on this aggregation port.
UnknownRx	The number of frames received that either: <ul style="list-style-type: none"> • carry Slow Protocols Ethernet type values, but contain an unknown PDU. • are addressed to the Slow Protocols group MAC Address, but do not carry the Slow Protocols Ethernet Type.
IllegalRx	The number of frames received that carry the Slow Protocols Ethernet Type value (43B.4), but contain a badly formed PDU or an illegal value of Protocol Subtype (43B.4).
LACPDUsTx	The number of LACPDU transmitted on this aggregation port.
MarkerPDUsTx	The number of marker PDUs transmitted on this aggregation port.

Name	Description
MarkerResponsePDUsTx	The number of marker response PDUs transmitted on this aggregation port.

Viewing port policer statistics

View port policer statistics to manage network performance.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > Graph**.
2. Click **Port**.
3. Click the **Policer** tab.

Policer field descriptions

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

Name	Description
TotalPkts	Shows the total number of packets received on the port.
TotalBytes	Shows the total number of bytes received on the port.
YellowBytes	Shows the total number of bytes received on the port that were above the committed rate but below the peak rate.
RedBytes	Shows the total number of bytes received on the port that were above the peak rate.

Displaying file statistics

Display the amount of memory used and available for both onboard flash memory and installed external storage devices, as well as the number of files in each location.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > Edit**.
2. Click **File System**.
3. Click the **Device Info** tab.

Device Info field descriptions

Use the data in the following table to use the **Device Info** tab.

Name	Description
Slot	Specifies the slot number of the CP module.
FlashBytesUsed	Specifies the number of bytes used in the internal flash memory.
FlashBytesFree	Specifies the number of bytes available for use in the internal flash memory.
FlashNumFiles	Specifies the number of files in the internal flash memory.
ExtflashBytesUsed	Specifies the number of bytes used in the external flash memory.
ExtflashBytesFree	Specifies the number of bytes available for use in the external flash memory.
ExtflashNumFiles	Specifies the number of files in the external flash memory.
UsbBytesUsed	Specifies the number of bytes used on the USB device.
UsbBytesFree	Specifies the number of bytes available on the USB device.
UsbNumFiles	Specifies the number of files on the USB device.

Viewing QoS policy statistics

Use policy statistics to better tailor policy parameters to suit customer needs.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > QoS**.
2. Click **Statistics**.
3. Click the **Policy Stats** tab.

Policy Stats field descriptions

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

Name	Description
PolicyId	Shows the global policer ID that corresponds to this local policer.
Slot	Shows the slot number of the chassis on which statistics are collected. Valid slots are IO slots.
TotalPkts	Shows the global policer total packet count.
TotalBytes	Shows the global policer total byte count.
GreenPackets	Shows the total number of packets received that were below the committed rate.
GreenBytes	Shows the total number of bytes received that were below the committed rate.
YellowPackets	Shows the total number of packets received that were above the committed rate but below the peak rate.

Name	Description
YellowBytes	Shows the total number of bytes received that were above the committed rate but below the peak rate.
RedPackets	Shows the total number of packets received that were above the peak rate.
RedBytes	Shows the total number of bytes received that were above the peak rate.

Graphing QoS policy statistics

Graph QoS policy statistics to create a visual comparison between data values.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > QoS**.
2. Click **Statistics**.
3. Click the **Policy Stats** tab.
4. Select a policy.
5. Click **Graph**.
6. Select one or more values.
7. Click the button for the type of graph you wish to view.

Policy field descriptions

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

Name	Description
TotalPkts	Shows the total packet count.
TotalBytes	Shows the total byte count.
GreenPackets	Shows the total number of packets received that were below the committed rate.
GreenBytes	Shows the total number of bytes received that were below the committed rate.
YellowPackets	Shows the total number of packets received that were above the committed rate but below the peak rate.
YellowBytes	Shows the total number of bytes received that were above the committed rate but below the peak rate.
RedPackets	Shows the total number of packets received that were above the peak rate.
RedBytes	Shows the total number of bytes received that were above the peak rate.

Viewing statistics for a specific QoS policy

View statistics for a specific QoS policy.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > QOS**.
2. Click **Policy**.
3. Select a policy.
4. Click **Stats**.

Policy Stats field descriptions

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

Name	Description
PolicyId	Shows the global policer ID that corresponds to this local policer.
Slot	Shows the slot number of the chassis on which statistics are collected. Valid slots are IO slots.
TotalPkts	Shows the global policer total packet count.
TotalBytes	Shows the global policer total byte count.
GreenPackets	Shows the total number of packets received that were below the committed rate.
GreenBytes	Shows the total number of bytes received that were below the committed rate.
YellowPackets	Shows the total number of packets received that were above the committed rate but below the peak rate.
YellowBytes	Shows the total number of bytes received that were above the committed rate but below the peak rate.
RedPackets	Shows the total number of packets received that were above the peak rate.
RedBytes	Shows the total number of bytes received that were above the peak rate.

Viewing ACE port statistics

Use port statistics to ensure that the ACE is operating correctly.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > Security > Data Path**.

2. Click **Advanced Filters (ACE/ACLs)**.
3. Click the **ACL** tab.
4. Select a field on the **ACL** tab.
5. Click **ACE**.
6. Click the **Statistics** tab.

Statistics field descriptions

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

Name	Description
AcId	Specifies the associated ACL index.
AcId	Specifies the ACE index.
MatchCountPkts	Specifies a packet count of the matching packets.
MatchCountOctets	Specifies the number of octets of the matching packets.

Viewing ACL statistics

Graph statistics for a specific ACL ID to view default statistics.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > Security > Data Path**.
2. Click **Advanced Filters (ACE/ACLs)**.
3. Click the **ACL** tab.
4. Select an ACL.
5. Click **Graph**.
6. You can click **Clear Counters** to clear the **Statistics** fields.

Statistics field descriptions

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

Name	Description
AcId	Specifies the ACL ID.
MatchDefaultSecurityPkts	Shows a security packet count of traffic that does not match an ACE rule or hits the default action if count is configured in the ACL global action.
MatchDefaultSecurityOctets	Shows a security byte count of traffic that does not match an ACE rule or hits the default action if count is configured in the ACL global action.

Name	Description
MatchDefaultQosPkts	Shows a QoS packet count of traffic that does not match an ACE rule or hits the default action if count is configured in the ACL global action.
MatchDefaultQosOctets	Shows a QoS byte count of traffic that does not match an ACE rule or hits the default action if count is configured in the ACL global action.
MatchGlobalSecurityPkts	Shows a security packet count of traffic that matches an ACE rule or hits the default action if count is configured in the ACL global action.
MatchGlobalSecurityOctets	Shows a security byte count of traffic that matches an ACE rule or hits the default action if count is configured in the ACL global action.
MatchGlobalQosPkts	Shows a QoS packet count of traffic that matches an ACE rule or hits the default action if count is configured in the ACL global action.
MatchGlobalQosOctets	Shows a QoS byte count of traffic that matches an ACE rule or hits the default action if count is configured in the ACL global action.

Clearing ACL statistics

Clear ACL statistics when you want to gather a new set of statistics.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > Security > Data Path**.
2. Click **Advanced Filters (ACE/ACLs)**.
3. Click the **ACL** tab.
4. Select a field.
5. Click **ClearStats**.

Viewing VLAN and Spanning Tree CIST statistics

View CIST port statistics to manage network performance.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > VLAN > Spanning Tree**.
2. Click **MSTP**.
3. Click the **CIST Port** tab.

4. Select a port, and then click **Graph**.

CIST field descriptions

The following table describes parameters on the **CIST** tab.

Name	Descriptions
ForwardTransitions	Specifies the number of times this port has transitioned to the forwarding state.
RxMstBpduCount	Specifies the number of MSTP BPDUs received on this port.
RxRstBpduCount	Specifies the number of RSTP BPDUs received on this port.
RxConfigBpduCount	Specifies the number of configuration BPDUs received on this port.
RxTcnBpduCount	Specifies the number of TCN BPDUs received on this port.
TxMstBpduCount	Specifies the number of MSTP BPDUs transmitted from this port.
TxRstBpduCount	Specifies the number of RSTP BPDUs transmitted from this port.
TxConfigBpduCount	Specifies the number of configuration BPDUs transmitted from this port.
TxTcnBpduCount	Specifies the number of TCN BPDUs transmitted from this port.
InvalidMstBpduRxCount	Specifies the number of Invalid MSTP BPDUs received on this port.
InvalidRstBpduRxCount	Specifies the number of Invalid RSTP BPDUs received on this port.
InvalidConfigBpduRxCount	Specifies the number of invalid configuration BPDUs received on this port.
InvalidTcnBpduRxCount	Specifies the number of invalid TCN BPDUs received on this port. The number of times this port has migrated from one STP protocol version to another. The relevant protocols are STP-compatible and RSTP/MSTP. A trap is generated on the occurrence of this event.
ProtocolMigrationCount	Specifies the number of times this port has migrated from one STP protocol version to another. The relevant protocols are STP-compatible and RSTP. A trap is generated on the occurrence of this event.

Viewing VLAN and Spanning Tree MSTI statistics

View multiple spanning tree instance (MSTI) port statistics to manage network performance.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > VLAN > Spanning Tree**.
2. Click **MSTP**.
3. Click the **MSTI Port** tab.
4. Select a port, and then click **Graph**.

MSTI field descriptions

The following table describes parameters on the **MSTI** tab.

Name	Description
ForwardTransitions	Specifies the number of times this port has transitioned to the forwarding state for this specific instance.
ReceivedBPDUs	Specifies the number of BPDUs received by this port for this spanning tree instance.
TransmittedBPDUs	Specifies the number of BPDUs transmitted on this port for this spanning tree instance.
InvalidBPDUsRcvd	Specifies the number of invalid BPDUs received on this port for this spanning tree instance.

Viewing VRRP interface stats

View VRRP statistics to manage network performance.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > IP**.
2. Click **VRRP**.
3. Select the **Interface** tab.
4. Select an interface.
5. Click **Graph**.

Interface field descriptions

The following table describes parameters on the **Interface** tab.

Name	Description
BecomeMaster	Specifies the number of times that this virtual router state transitions from BACKUP to MASTER.
AdvertiseRcvd	Specifies the number of VRRP advertisements received by this virtual router.
AdvertiseIntervalErrors	Specifies the number of received VRRP advertisement packets with a different interval is than configured for the local virtual router.
IPtTtlErrors	Specifies the number of VRRP packets received by the virtual router with IP TTL (Time-To-Live) not equal to 255.
PriorityZeroPktsRcvd	Specifies the number of VRRP packets received by the virtual router with a priority of 0.

Name	Description
PriorityZeroPktsSent	Specifies the number of VRRP packets sent by the virtual router with a priority of 0'.
InvalidTypePktsRcvd	Specifies the number of VRRP packets received by the virtual router with an invalid value in the 'type' field.
AddressListErrors	Specifies the packets received address list the address list does not match the locally configured list for the virtual router.
AuthTypeMismatch	Specifies the count of authentication type mismatch messages.
PacketLengthErrors	Specifies the count of packet length errors.
AuthFailures	Specifies the count of authentication failure messages.

Viewing VRRP statistics

View VRRP statistics to monitor network performance.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > IP**.
2. Click **VRRP**.
3. Select the **Stats** tab.

Stats field descriptions

The following table describes parameters on the VRRP statistics tab.

Name	Description
ChecksumErrors	Specifies the number of VRRP packets received with an invalid VRRP checksum value.
VersionErrors	Specifies the number of VRRP packets received with an unknown or unsupported version number.
VrIDErrors	Specifies the number of VRRP packets received with an invalid VrID for this virtual router.

Viewing IPv6 VRRP statistics for an interface

View IPv6 VRRP statistics for a VLAN or port.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > IPv6**.

2. Click **VRRP**.
3. Click the **Interface** tab.
4. Select an interface.
5. Click **Statistics**.

Statistics field descriptions

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

Name	Description
MasterTransitions	Shows the total number of times that the state of this virtual router has transitioned to master. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of DiscontinuityTime.
RcdAdvertisements	Shows the total number of VRRP advertisements received by this virtual router. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of DiscontinuityTime.
AdvIntervalErrors	Shows the total number of VRRP advertisement packets received for which the advertisement interval is different than the one configured for the local virtual router. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of DiscontinuityTime.
IpTtlErrors	Shows the total number of VRRP packets received by the virtual router with IPv4 TTL (for VRRP over IPv4) or IPv6 Hop Limit (for VRRP over IPv6) not equal to 255. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of DiscontinuityTime.
RcvdPriZeroPackets	Shows the total number of VRRP packets received by the virtual router with a priority of 0. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of DiscontinuityTime.
SentPriZeroPackets	Shows the total number of VRRP packets sent by the virtual router with a priority of 0. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of DiscontinuityTime.

Name	Description
RcvdInvalidTypePkts	Shows the number of VRRP packets received by the virtual router with an invalid value in the 'type' field. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of DiscontinuityTime.
AddressListErrors	Shows the total number of packets received for which the address list does not match the locally configured list for the virtual router. Discontinuities in the value of this counter can occur at reinitialization of the management system, and at other times as indicated by the value of DiscontinuityTime.
PacketLengthErrors	Shows the total number of packets received with a packet length less than the length of the VRRP header. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of DiscontinuityTime.
RcvdInvalidAuthentications	Shows the total number of packets received with an unknown authentication type.

Viewing IPv6 VRRP statistics

View IPv6 VRRP statistics to monitor network performance.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > IPv6**.
2. Click **VRRP**.
3. Click the **Stats** tab.

Stats field descriptions

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

Name	Description
ChecksumErrors	Shows the number of VRRP packets received with an invalid VRRP checksum value.
VersionErrors	Shows the number of VRRP packets received with an unknown or unsupported version number.
VrldErrors	Shows the number of VRRP packets received with an invalid VrID for this virtual router.

Viewing SMLT statistics

View SMLT statistics to manage network performance.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > VLAN**.
2. Click **MLT/LACP**.
3. Select the **Ist/SMLT Stats** tab.

IST/SMLT Stats field descriptions

The following table describes parameters on the **IST/SMLT Stats** tab.

Name	Description
SmltIstDownCnt	The number of times the session between the two peering switches has gone down since last boot.
SmltHelloTxMsgCnt	The count of transmitted hello messages.
SmltHelloRxMsgCnt	The count of received hello messages.
SmltLearnMacAddrTxMsgCnt	The count of transmitted learned MAC address messages.
SmltLearnMacAddrRxMsgCnt	The count of received learned MAC address messages.
SmltMacAddrAgeOutTxMsgCnt	The count of transmitted aging out MAC address messages.
SmltMacAddrAgeOutRxMsgCnt	The count of received aging out MAC address messages.
SmltMacAddrAgeExpTxMsgCnt	The count of transmitted MAC address age expired messages.
SmltMacAddrAgeExpRxMsgCnt	The count of received MAC address age expired messages.
SmltStgInfoTxMsgCnt	The count of transmitted STG information messages.
SmltStgInfoRxMsgCnt	The count of received STG information messages.
SmltDelMacAddrTxMsgCnt	The count of transmitted MAC address deleted messages.
SmltDelMacAddrRxMsgCnt	The count of received MAC address received messages.
SmltSmltDownTxMsgCnt	The count of transmitted SMLT down messages.
SmltSmltDownRxMsgCnt	The count of received SMLT down messages.
SmltUpTxMsgCnt	The count of transmitted SMLT up messages.
SmltUpRxMsgCnt	The count of received SMLT up messages.
SmltSendMacTblTxMsgCnt	The count of sent send MAC table messages.

Name	Description
SmltSendMacTblRxMsgCnt	The count of received send MAC table messages.
SmltIcmpTxMsgCnt	The count of sent IGMP messages.
SmltIcmpRxMsgCnt	The count of received IGMP messages.
SmltPortDownTxMsgCnt	The count of sent port down messages.
SmltPortDownRxMsgCnt	The count of received port down messages.
SmltReqMacTblTxMsgCnt	The count or sent MAC table request messages.
SmltReqMacTblRxMsgCnt	The count of received MAC table request messages.
SmltRxUnknownMsgTypeCnt	The count of received unknown message type messages.

Viewing RSTP status statistics

You can view status statistics for Rapid Spanning Tree Protocol (RSTP).

Procedure

1. In the navigation pane, expand the following folders: **Configuration > VLAN > Spanning Tree**.
2. Click **RSTP**.
3. In the **RSTP Status** tab, select a port, and then click **Graph**.

RSTP Status field descriptions

The following table describes the **RSTP Status** fields.

Name	Description
RxRstBpduCount	Specifies the number of RSTP BPDUs this port received.
RxConfigBpduCount	Specifies the number of configuration BPDUs this port received.
RxTcnBpduCount	Specifies the number of TCN BPDUs this port received.
TxRstBpduCount	Specifies the number of RSTP BPDUs this port transmitted.
TxConfigBpduCount	Specifies the number of Config BPDUs this port transmitted.
TxTcnBpduCount	Specifies the number of TCN BPDUs this port transmitted.
InvalidRstBpduRxCount	Specifies the number of invalid RSTP BPDUs this port received. A trap is generated on the occurrence of this event.
InvalidConfigBpduRx Count	Specifies the number of invalid configuration BPDUs this port received. A trap is generated on the occurrence of this event.
InvalidTcnBpduRxCount	Specifies the number of invalid TCN BPDUs this port received. A trap is generated on the occurrence of this event.
ProtocolMigrationCount	Specifies the number of times this port migrated from one STP protocol version to another. The relevant protocols are STP-

Name	Description
	Compatible and RSTP. A trap is generated on the occurrence of this event.
ForwardTransitions	Specifies the number of times this port has transitioned from the learning state to the forwarding state.

Viewing MLT interface statistics

Use MLT interface statistics tab to view interface statistics for the selected MLT.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > VLAN**.
2. Click **MLT/LACP**.
3. Click the **MultiLink/LACP Trunks** tab.
4. Select an MLT.
5. Click **Graph**.

MultiLink/LACP Trunks field descriptions

Use the data in the following table to use the **MultiLink/LACP Trunks** tab.

Name	Description
InOctets	Specifies the total number of octets received on the MLT interface, including framing characters.
OutOctets	Specifies the total number of octets transmitted out of the MLT interface, including framing characters.
InUcastPkts	Specifies the number of packets delivered by this MLT to higher level protocols that were not addressed to a multicast or broadcast address at this sublayer.
OutUcastPkts	Specifies the number of packets that higher-level protocols requested be transmitted that were not addressed to a multicast address at this MLT. This total number includes discarded or unsent packets.
InMulticastPkt	Specifies the number of packets delivered to this MLT that were addressed to a multicast address at this sublayer. For a MAC layer protocol, this number includes both Group and Functional addresses.
OutMulticast	Specifies the total number of packets that higher-level protocols requested be transmitted, and that were addressed to a multicast address at this MLT, including those that were discarded or unsent. For a MAC layer protocol, this number includes both Group and Functional addresses.
InBroadcastPkt	Specifies the number of packets delivered to this MLT that were addressed to a broadcast address at this sublayer.
OutBroadcast	Specifies the total number of packets that higher-level protocols requested be transmitted, and that were addressed to a broadcast address at this MLT, including those that were discarded or not sent.

Name	Description
InLsmPkts	Specifies the total number of Link State Messaging (LSM) packets delivered on this MLT.
OutLsmPkts	Specifies the total number of Link State Messaging (LSM) packets transmitted on this MLT.

Viewing MLT Ethernet error statistics

Use MLT Ethernet error statistics to view the error statistics.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > VLAN**.
2. Click **MLT/LACP**.
3. Click the **MultiLink/LACP Trunks** tab.
4. Select an MLT, and then click **Graph**.
5. Click the **Ethernet Errors** tab.

Ethernet Errors field descriptions

Use the data in the following table to use the **Ethernet Errors** tab.

Name	Description
AlignmentErrors	Specifies the frame count frames received on a particular MLT that is not an integral number of octets in length and does not pass the FCS check. The count represented by an instance of this object increments when the alignmentError status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions occur are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.
FCSErrors	Specifies the frame count received on an MLT that is an integral number of octets in length, but does not pass the Frame Check Sequence (FCS) check. The count represented by an instance of this object increments when the FrameCheckError status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions occur are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.
IMacTransmitError	Specifies the frame count for which transmission on a particular MLT fails due to an internal MAC sublayer transmit error. A frame is only counted by an instance of this object if it is not counted by the corresponding instance of either the LateCollisions object, the ExcessiveCollisions object, or the CarrierSenseErrors object.
IMacReceiveError	Specifies the frame count for which reception on a particular MLT fails due to an internal MAC sublayer receive error. A frame is only counted by an

Name	Description
	<p>instance of this object if it is not counted by the corresponding instance of either the FrameTooLongs object, the AlignmentErrors object, or the FCSErrors object.</p> <p>The precise meaning of the count represented by an instance of this object is implementation specific. In particular, an instance of this object can represent receive errors on a particular interface that are not otherwise counted.</p>
CarrierSenseError	<p>Specifies the number of times that the carrier sense condition was lost or never asserted when attempting to transmit a frame on a particular MLT. The count represented by an instance of this object increments at most once for each transmission attempt, even if the carrier sense condition fluctuates during a transmission attempt.</p>
FrameTooLong	<p>Specifies the frame count received on a particular MLT that exceeds the maximum permitted frame size. The count represented by an instance of this object increments when the frameTooLong status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions occur are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.</p>
SQETestError	<p>Specifies the number of times that the SQE test error message is generated by the PLS sublayer for a particular MLT. The SQE test error message is defined in section 7.2.2.2.4 of ANSI/ IEEE 802.3-1985.</p>
DeferredTransmiss	<p>Specifies the frame count for which the first transmission attempt on a particular MLT is delayed because the medium is busy. The count represented by an instance of this object does not include frames involved in collisions.</p>
SingleCollFrames	<p>Specifies a count of successfully transmitted frames on a particular MLT for which transmission is inhibited by exactly one collision. A frame that is counted by an instance of this object is also counted by the corresponding instance of either the ifOutUcastPkts object, the ifOutMulticastPkts object, or the ifOutBroadcastPkts object, and is not counted by the corresponding instance of the MultipleCollisionFrames object.</p>
MultipleCollFrames	<p>Specifies the successfully transmitted frame count on a particular MLT for which transmission is inhibited by more than one collision. A frame that is counted by an instance of this object is also counted by the corresponding instance of either the ifOutUcastPkts object, the ifOutMulticastPkts object, or the ifOutBroadcastPkts object, and is not counted by the corresponding instance of the SingleCollisionFrames object.</p>
LateCollisions	<p>Specifies the number of times that a collision is detected on a particular MLT later than 512 bit-times into the transmission of a packet; 512 corresponds to 51.2 microseconds on a 10 Mb/s system. A late collision included in a count represented by an instance of this object is also considered as a generic collision for purposes of other collision-related statistics.</p>
ExcessiveCollis	<p>Specifies the frame count for which transmission on a particular MLT fails due to excessive collisions.</p>

Viewing RIP statistics

Use statistics to help you monitor Routing Information Protocol (RIP) performance. You can also use statistics in troubleshooting procedures.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > IP**.
2. Click **RIP**.
3. Click the **Status** tab.

Status field descriptions

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

Name	Description
Address	The IP address of the router interface.
RcvBadPackets	The number of RIP response packets received by the RIP process that were subsequently discarded for any reason (examples: a version 0 packet or an unknown command type).
RcvBadRoutes	The number of routes, in valid RIP packets, that were ignored for any reason (examples: unknown address family or invalid metric).
SentUpdates	The number of triggered RIP updates actually sent on this interface. This field explicitly does not include full updates sent containing new information.

Viewing OSPF chassis statistics

Use statistics to help you monitor Open Shortest Path First (OSPF) performance. You can also graph statistics for all OSPF packets transmitted by the switch.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > IP**.
2. Click **OSPF**.
3. Click the **Stats** tab.
4. To create a graph for OSPF statistics, select a column, and then select a graph type.

Stats field descriptions

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

Name	Description
LsdbTblSize	Specifies the number of entries in the link state database table.
TxPackets	Specifies the number of packets transmitted by OSPF.

Name	Description
RxPackets	Specifies the number of packets received by OSPF.
TxDropPackets	Specifies the number of packets dropped before being transmitted by OSPF.
RxDropPackets	Specifies the number of packets dropped before they are received by OSPF.
RxBadPackets	Specifies the number of packets received by OSPF that are bad.
SpfRuns	Specifies the number of SPF calculations performed by OSPF.
BuffersAllocated	Specifies the number of buffers allocated for OSPF.
BuffersFreed	Specifies the number of buffers freed by OSPF.
BufferAllocFailures	Specifies the number of times that OSPF has failed to allocate buffers.
BufferFreeFailures	Specifies the number of times that OSPF has failed to free buffers.
Routes	Specifies the count of OSPF routes.
Adjacencies	Specifies the count of OSPF adjacencies.
Areas	Specifies the count of OSPF areas.

Viewing IPv6 OSPF statistics

View OSPF statistics to analyze trends. You can also graph statistics for all OSPF packets transmitted by the switch.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > IPv6**.
2. Click **OSPF**.
3. Click **Stats**.

Stats field descriptions

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

Name	Description
TxPackets	Shows the count of sent packets.
RxPackets	Shows the count of received packets.
TxDropPackets	Shows the count of sent, dropped packets.
RxDropPackets	Shows the count of received, dropped packets.
RxBadPackets	Shows the count of received, bad packets.
SpfRuns	Shows the count of intra-area route table updates with calculations using this area link-state database.
LastSpfRun	Shows the count of the most recent SPF run.
LsdbTblSize	Shows the size of the link state database table.
BadLsReqs	Shows the count of bad link requests.

Name	Description
SeqMismatches	Shows the count of sequence mismatched packets.

Graphing OSPF statistics for a VLAN

Use statistics to help you monitor OSPF performance on a VLAN. You can also graph statistics for all OSPF packets.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > VLAN**.
2. Click **VLANs**.
3. Select a **VLAN**.
4. Click **IP**.
5. Click the **OSPF** tab.
6. Click **Graph**.
7. Select one or more values.
8. Click the type of graph.

OSPF field descriptions

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

Name	Description
VersionMismatches	Indicates the number of version mismatches received by this interface.
AreaMismatches	Indicates the number of area mismatches received by this interface.
AuthTypeMismatches	Indicates the number of AuthType mismatches received by this interface.
AuthFailures	Indicates the number of authentication failures.
NetMaskMistmatches	Indicates the number of net mask mismatches received by this interface.
HelloIntervalMismatches	Indicates the number of hello interval mismatches received by this interface.
DeadIntervalMismatches	Indicates the number of dead interval mismatches received by this interface.
OptionMismatches	Indicates the number of options mismatches received by this interface.
RxHellos	Indicates the number of hello packets received by this interface.

Name	Description
RxDBDescrs	Indicates the number of database descriptor packets received by this interface.
RxLSUpdates	Indicate the number of Link state update packets received by this interface.
RxLsReqs	Indicates the number of Link state request packets received by this interface.
RxLSAcks	Indicates the number of Link state acknowledge packets received by this interface.
TxHellos	Indicates the number of hello packets transmitted by this interface.
TxDBDescrs	Indicates the number of database descriptor packets transmitted by this interface.
TxLSUpdates	Indicate the number of Link state update packets transmitted by this interface.
TxLSReqs	Indicates the number of Link state request packets transmitted by this interface.
TxLSAcks	Indicates the number of Link state acknowledge packets transmitted by this interface.

Graphing OSPF statistics for a port

Use statistics to help you monitor OSPF performance on a VLAN. You can also graph statistics for all OSPF packets.

Procedure

1. On the Device Physical View, select a port.
2. In the navigation pane, expand the following folders: **Configuration > Edit > Port**.
3. Click **IP**.
4. Click the **OSPF** tab.
5. Click **Graph**.
6. Select one or more values.
7. Click the type of graph.

OSPF field descriptions

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

Name	Description
VersionMismatches	Indicates the number of version mismatches received by this interface.

Name	Description
AreaMismatches	Indicates the number of area mismatches received by this interface.
AuthTypeMismatches	Indicates the number of AuthType mismatches received by this interface.
AuthFailures	Indicates the number of authentication failures.
NetMaskMistmatches	Indicates the number of net mask mismatches received by this interface.
HelloIntervalMismatches	Indicates the number of hello interval mismatches received by this interface.
DeadIntervalMismatches	Indicates the number of dead interval mismatches received by this interface.
OptionMismatches	Indicates the number of options mismatches received by this interface.
RxHellos	Indicates the number of hello packets received by this interface.
RxDBDescrs	Indicates the number of database descriptor packets received by this interface.
RxLSUpdates	Indicate the number of Link state update packets received by this interface.
RxLsReqs	Indicates the number of Link state request packets received by this interface.
RxLSAcks	Indicates the number of Link state acknowledge packets received by this interface.
TxHellos	Indicates the number of hello packets transmitted by this interface.
TxDBDescrs	Indicates the number of database descriptor packets transmitted by this interface.
TxLSUpdates	Indicate the number of Link state update packets transmitted by this interface.
TxLSReqs	Indicates the number of Link state request packets transmitted by this interface.
TxLSAcks	Indicates the number of Link state acknowledge packets transmitted by this interface.

Viewing BGP global stats

View BGP global stats.

Procedure

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

2. Click **BGP**.
3. Click the **Global Stats** tab.

Global Stats field descriptions

Use the data in the following table to use the BGP Global Stats tab.

Name	Description
AbsoluteValue	Displays the counter value.
Cumulative	Displays the total value since you opened the Stats tab.
Average/sec	Displays the average value for each second.
Minimum/sec	Displays the minimum value for each second.
Maximum/sec	Displays the maximum value for each second.
LastVal/sec	Displays the last value for each second.
Starts	Displays the number of times the BGP connection started.
Stops	Displays the number of times the BGP connection stopped.
Opens	Displays the number of times BGP opens TCP.
Closes	Displays the number of times BGP closes TCP.
Fails	Displays the number of times TCP attempts failed.
Fatals	Displays the number of times TCP crashes due to fatal error.
ConnExps	Displays the number of times the TCP retry timer expired.
HoldExps	Displays the number of times the hold timer expired.
KeepExps	Displays the number of times the keepalive timer expired.
RxOpens	Displays the number of open instances BGP receives.
RxKeeps	Displays the number of keepalive instances BGP receives.
RxUpdates	Displays the number of update instances BGP receives.
RxNotifys	Displays the number of notification instances BGP receives.
TxOpens	Displays the number of open instances BGP transmitted.
TxKeeps	Displays the number of keepalive instances BGP transmitted.
TxUpdates	Displays the number of updates instances BGP transmits.

Name	Description
TxNotifys	Displays the number of notification instances BGP transmits.
BadEvents	Displays the number of invalid events FSM received.
SyncFails	Displays the number of times FDB sync failed.
TrEvent	Displays the trace event.
RxECodeHeader	Displays the total header errors received.
RxECodeOpen	Displays the total open errors received.
RxECodeUpdate	Displays the total update errors received.
RxECodeHoldtimer	Displays the total holdtimer errors received.
RxECodeFSM	Displays the total FSM errors received.
RxECodeCease	Displays the total cease errors received.
RxHdrCodeNoSync	Displays the header not synchronized errors received.
RxHdrCodeInvalidMsgLen	Displays the header invalid message length errors received.
RxHdrCodeInvalidMsgType	Displays the header invalid message type errors received.
RxOpCodeBadVer	Displays the open errors received for Bad Version.
RxOpCodeBadAs	Displays the open errors received for le Bad AS Number.
RxOpCodeBadRtrID	Displays the open errors received for Bad BGP Rtr ID.
RxOpCodeUnsuppOption	Displays the open errors received for Unsupported Option.
RxOpCodeAuthFail	Displays the open errors received for Auth Failures.
RxOpCodeBadHold	Displays the open errors received for Bad Hold Value.
RxUpdCodeMalformedAttrList	Displays the update errors received for Malformed Attr List.
RxUpdCodeWelKnownAttrUnrecog	Displays the update errors received for Welknown Attr Unrecog.
RxUpdCodeWelknownAttrMiss	Displays the update errors received for Welknown Attr Missing.
RxUpdCodeAttrFlagError	Displays the update errors received for Attr Flag Error.
RxUpdCodeAttrLenError	Displays the update errors received for Attr Len Error.
RxUpdCodeBadORIGINAttr	Displays the update errors received for Bad ORIGIN Attr.

Name	Description
RxUpdCodeASRoutingLoop	Displays the update errors received for AS Routing Loop.
RxUpdCodeBadNHAttr	Displays the update errors received for Bad NEXT-HOP Attr.
RxUpdCodeOptionalAttrError	Displays the update errors receivedfor Optional Attr Error.
RxUpdCodeBadNetworkField	Displays the update errors received for Bad Network Field.
RxUpdCodeMalformedASPath	Displays the update errors received for Malformed AS Path.
TxECodeHeader	Displays the total Header errors transmitted.
TxECodeOpen	Displays the total Open errors transmitted.
TxECodeUpdate	Displays the total Update errors transmitted.
TxECodeHoldtimer	Displays the total Holdtimer errors transmitted.
TxECodeFSM	Displays the total FSM errors transmitted.
TxECodeCease	Displays the total Cease errors transmitted.
TxHdrCodeNoSync	Displays the header Not Synchronized errors transmitted.
TxHdrCodeInvalidMsgLen	Displays the header Invalid msg len errors transmitted.
TxHdrCodeInvalidMsgType	Displays the header Invalid msg type errors transmitted.
TxOpCodeBadVer	Displays the open errors transmitted for Bad Version.
TxOpCodeBadAs	Displays the open errors transmitted for Bad AS Number.
TxOpCodeBadRtrID	Displays the open errors transmitted for Bad BGP Rtr ID.
TxOpCodeUnsuppOption	Displays the open errors transmitted for Unsupported Option.
TxOpCodeAuthFail	Displays the open errors transmitted for Auth Failures.
TxOpCodeBadHold	Displays the open errors transmitted for Bad Hold Value.
TxUpdCodeMalformedAttrList	Displays the update errors transmitted for Malformed Attr List.
TxUpdCodeWelknownAttrUnrecog	Displays the update errors transmitted for Welknown Attr Unrecog.
TxUpdCodeWelknownAttrMiss	Displays the update errors transmitted for Welknown Attr Missing.
TxUpdCodeAttrFlagError	Displays the update errors transmitted for Attr Flag Error.

Name	Description
TxUpdCodeAttrLenError	Displays the update errors transmitted for Attr Len Error.
TxUpdCodeBadORIGINAttr	Displays the update errors transmitted for Bad ORIGIN Attr.
TxUpdCodeASRoutingLoop	Displays the update errors transmitted for AS Routing Loop
TxUpdCodeBadNHAttr	Displays the update errors transmitted for Bad NEXT-HOP Attr
TxUpdCodeOptionalAttrError	Displays the update errors transmitted for Optional Attr Error.
TxUpdCodeBadNetworkField	Displays the update errors transmitted for Bad Network Field.
TxUpdCodeMalformedASPath	Displays the update errors transmitted for Malformed AS Path.

Viewing BGP peer general statistics

View BGP peer general statistics to manage network performance.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > IP**.
2. Click **BGP**.
3. Select the **Peers** tab.
4. Select the peer for which you want to view statistics.
5. Click **Graph**.
6. Select one or more values.
7. Click the type of graph you want to create.

General Stats field descriptions

Use the data in the following table to use the **General Stats** tab.

Name	Description
InUpdates	Specifies the number of BGP update messages received on this connection. This object must be initialized to zero (0) when the connection is established.
OutUpdates	Specifies the number of BGP update messages transmitted on this connection. This object must be initialized to zero (0) when the connection is established.

Name	Description
InTotalMessages	Specifies the total number of messages receive from the remote peer on this connection. This object must be initialized to zero when the connection is established.
OutTotalMessages	Specifies the total number of messages transmitted to the remote peer on this connection. This object must be initialized to zero when the connection is established.
FsmEstablishedTransitions	Specifies the total number of times the BGP FSM transitioned into the established state.
FsmEstablishedTime	This timer indicates the duration of the peer in the established state or the time since the peer was last in the established state. It is set to zero when a new peer is configured or the router is booted.
InUpdatesElapsedTime	Specifies the elapsed time (in seconds) since the last BGP update message was received from the peer. Each time <code>bgpPeerInUpdates</code> is incremented, the value of this object is set to zero (0).

Viewing BGP peer advanced statistics

View BGP peer advance statistics to manage network performance.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > IP**.
2. Click **BGP**.
3. Select the **Peers** tab.
4. Select the peer for which you want to view statistics.
5. Click **Graph**.
6. Select the **Advanced Stats** tab.

Advanced stats field descriptions

Use the data in the following table to use the **Advanced Stats** tab.

Name	Description
Starts	Shows the number of times peer BGP connection started.
Stops	Shows the number of times peer BGP connection stopped.
Opens	Shows the number of times peer opened TCP.
Closes	Shows the number of times peer closed TCP.

Name	Description
Fails	Shows the number of times a peer TCP attempt failed.
Fatals	Shows the number of times peer TCP crashed due to fatal error.
ConnExps	Shows the number of times the peer TCP retry timer expired.
HoldExps	Shows the number of times the peer hold timer expired.
KeepExps	Shows the number of times the peer keepalive timer expired.
BadEvents	Shows the number of invalid events received by the peer.
SyncFails	Shows the number of times the peer FDB sync failed.
RcvdTooShort	Shows the number of “too short” messages received by the peer.
NoMarker	Shows the number of messages received by the peer with no marker.
Dropped	Shows the number of messages dropped by the peer.
BadMsgTypes	Shows the number of messages received by the peer as “invalid message type.”
TrEvent	Shows the peer trace event.

Viewing BGP peer receive statistics

View BGP peer receive statistics to manage network performance.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > IP**.
2. Click **BGP**.
3. Select the **Peers** tab.
4. Select the peer for which you want to view statistics.
5. Click **Graph**.
6. Select the **Receive Stats** tab.

Receive Stats field descriptions

The following table describes parameters on the **Receive Stats** tab.

Name	Description
RxMsgs	Shows the number of messages received by the peer.
RxInCompPkts	Shows the number of Incomplete messages received by the peer.
RxOpens	Shows the number of opens received by the peer.
RxKeeps	Shows the number of keepalive messages received by the peer.
RxUpdates	Shows the number of updates received by the peer.
RxNotifys	Shows the number of notifications received by the peer.
RxRoutesAdded	Shows the number of routes added into loc_rib by this peer.
RxRoutesReplaced	Shows the number of routes that were replaced by routes received by the peer.
RxNlri	Shows the number of network layer reachability information (NLRI) messages received by the peer.
RxValidUpdates	Shows the number of valid updates received by the peer.
RxECodeHeader	Shows the number of header errors received by the peer.
RxECodeOpen	Shows the number of open errors received by the peer.
RxECodeUpdate	Shows the number of update errors received by the peer.
RxECodeHoldtimer	Shows the number of hold errors received by the peer.
RxECodeFSM	Shows the number of FSM errors received by the peer.
RxECodeCease	Shows the number of cease errors received by the peer.
RxHdrCodeNoSync	Shows the number of header errors received by the peer as: Not Synchronized.
RxHdrCodeInvalidMsgLen	Shows the number of header errors received by the peer as: Invalid Message Length.
RxHdrCodeInvalidMsgType	Shows the number of header errors received by the peer as: Invalid Message Type.
RxOpCodeBadVer	Shows the number of open errors received by the peer as: Bad Version.
RxOpCodeBadAs	Shows the number of open errors received by the peer as: Bad AS.

Name	Description
RxOpCodeBadRtID	Shows the number of open errors received by the peer as: Bad BGP ID.
RxOpCodeUnsuppOption	Shows the number of open errors received by the peer as: Unsupported Options.
RxOpCodeAuthFail	Shows the number of open errors received by the peer as: Authorization Failures.
RxOpCodeBadHold	Shows the number of open errors received by the peer as: Bad Hold Value.
RxUpdCodeMalformedAttrList	Shows the number of update errors received by the peer as: Malformed Attr List.
RxUpdCodeWelknownAttrUnrecog	Shows the number of update errors received by the peer as: Wellknown Attr Unrecog.
RxUpdCodeWelknownAttrMiss	Shows the number of Update errors received by the peer as: Wellknown Attr Missing.
RxUpdCodeAttrFlagError	Shows the number of update errors received by the peer as: Attr Flag Error.
RxUpdCodeAttrLenError	Shows the number of update errors received by the peer as: Attr Length Error.
RxUpdCodeBadORIGINAttr	Shows the number of update errors received by the peer as: Attr Flag Error.
RxUpdCodeASRoutingLoop	Shows the number of update errors received by the peer as: AS Routing Loop.
RxUpdCodeBadNHAttr	Shows the number of update errors received by the peer as: Bad Next-Hop Attr.
RxUpdCodeOptionalAttrError	Shows the number of update errors received by the peer as: Optional Attr Error.
RxUpdCodeBadNetworkField	Shows the number of update errors received by the peer as: Bad Network Field.
RxUpdCodeMalformedASPath	Shows the number of update errors received by the peer as: Malformed AS Path.

Viewing BGP peer transmit statistics

View BGP peer transmit statistics to manage network performance.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > IP**.
2. Click **BGP**.
3. Select the **Peers** tab.
4. Select the peer for which you want to view statistics.

5. Click **Graph**.
6. Select the **Transmit Stats** tab.

Transmit Stats field descriptions

The following table describes parameters on the **Transmit Stats** tab.

Name	Description
TxMsgs	Shows the number of messages transmitted by the peer.
TxOpens	Shows the number of opens transmitted by the peer.
TxKeeps	Shows the number of keepalive messages transmitted by the peer.
TxUpdates	Shows the number of updates transmitted by the peer.
TxNotifys	Shows the number of notifications transmitted by the peer.
TxRoutes	Shows the number of network layer reachability information (NLRI) messages transmitted by the peer.
TxECodeHeader	Shows the number of header errors transmitted by the peer.
TxECodeOpen	Shows the number of open errors transmitted by the peer.
TxECodeUpdate	Shows the number of update errors transmitted by the peer.
TxECodeHoldtimer	Shows the number of hold errors transmitted by the peer.
TxECodeFSM	Shows the number of FSM errors transmitted by the peer.
TxECodeCease	Shows the number of cease errors transmitted by the peer.
TxHdrCodeNoSync	Shows the number of header errors transmitted by the peer as: Not Synchronized.
TxHdrCodeInvalidMsgLen	Shows the number of header errors transmitted by the peer as: Invalid Message Length.
TxHdrCodeInvalidMsgType	Shows the number of header errors transmitted by the peer as: Invalid Message Type.
TxOpCodeBadVer	Shows the number of open errors transmitted by the peer as: Bad Version.
TxOpCodeBadAs	Shows the number of open errors transmitted by the peer as: Bad AS.
TxOpCodeBadRtID	Shows the number of open errors transmitted by the peer as: Bad BGP ID.

Name	Description
TxOpCodeUnsuppOption	Shows the number of open errors transmitted by the peer as: Unsupported Options.
TxOpCodeAuthFail	Shows the number of open errors transmitted by the peer as: Authorization Failures.
TxOpCodeBadHold	Shows the number of open errors transmitted by the peer as: Bad Hold Value.
TxUpdCodeMalformedAttrList	Shows the number of update errors transmitted by the peer as: Malformed Attr List.
TxUpdCodeWelknownAttrUnrecog	Shows the number of update errors transmitted by the peer as: Wellknown Attr Unrecog.
TxUpdCodeWelknownAttrMiss	Shows the number of update errors transmitted by the peer as: Wellknown Attr Missing.
TxUpdCodeAttrFlagError	Shows the number of update errors transmitted by the peer as: Attr Flag Error.
TxUpdCodeAttrLenError	Shows the number of update errors transmitted by the peer as: Attr Length Error.
TxUpdCodeBadORIGINAttr	Shows the number of update errors transmitted by the peer as: Attr Flag Error.
TxUpdCodeASRoutingLoop	Shows the number of update errors transmitted by the peer as: AS Routing Loop.
TxUpdCodeBadNHAttr	Shows the number of update errors transmitted by the peer as: Bad Next-Hop Attr.
TxUpdCodeOptionalAttrError	Shows the number of update errors transmitted by the peer as: Optional Attr Error.
TxUpdCodeBadNetworkField	Shows the number of update errors transmitted by the peer as: Bad Network Field.
TxUpdCodeMalformedASPath	Shows the number of update errors transmitted by the peer as: Malformed AS Path.

Viewing statistics for a VRF

View VRF statistics to ensure the instance is performing as expected.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > IP**.
2. Click **VRF**.
3. Select a VRF.
4. Click the **Stats** button.

Stats field descriptions

Use the data in the following table to help you understand the VRF statistics.

Name	Description
StatRouteEntries	Specifies the total number of routes for this VRF.
FIBEntries	Specifies the total number of Forwarding Information Base (FIB) entries for this VRF.

Viewing EAPoL Authenticator statistics

Use EAPoL Authenticator statistics to display the Authenticator Port Access Entity (PAE) statistics for each selected port.

Procedure

1. On the Device Physical View, select the port you want to graph.
A yellow outline appears around the selected ports
If you want to select multiple ports, press Ctrl and hold down the key while you click the ports you want to configure. A yellow outline appears around the selected ports.
2. In the navigation pane, expand the following folders: **Configuration > Graph**, and then click **Port**.
3. Click **EAPOL Stats**.
4. If you selected multiple ports, from the Graph port EAPoL Stats tab Show list, select: Absolute Value, Cumulative, Average/sec, Minimum/sec, Maximum/sec, or LastVal/sec.

EAPOL Stats field descriptions

The following table describes values on the **EAPOL Stats** tab.

Name	Description
FramesRx	Displays the number of valid EAPoL frames of any type received by this Authenticator.
FramesTx	Displays the number of EAPoL frame types of any type transmitted by this Authenticator.
StartFramesRx	Displays the number of EAPoL start frames received by this Authenticator.
LogoffFramesRx	Displays the number of EAPoL Logoff frames received by this Authenticator.
RespIdFramesRx	Displays the number of EAPoL Resp/Id frames received by this Authenticator.
RespFramesRx	Displays the number of valid EAP Response frames (Other than Resp/Id frames) received by this Authenticator.

Name	Description
ReqIdFramesTx	Displays the number of EAPoL Req/Id frames transmitted by this Authenticator.
ReqFramesTx	Displays the number of EAP Req/Id frames (other than Rq/Id frames) transmitted by this Authenticator.
EapLengthErrorFramesRx	Displays the number of EAPoL frames received by this Authenticator in which the Packet Body Length field is invalid.
InvalidEapolFramesRx	Displays the number of EAPoL frames received by this Authenticator in which the frame type is not recognized.

Viewing EAPoL diagnostic statistics

Use EAPoL diagnostic statistics to display the Authenticator PAE diagnostic statistics for each selected port.

Procedure

1. On the Device Physical View, select the port you want to graph.

A yellow outline appears around the selected ports

If you want to select multiple ports, press Ctrl and hold down the key while you select the ports you want to configure. A yellow outline appears around the selected ports.

2. Perform one of the following steps:

- In the navigation pane, expand the following folders: **Configuration > Graph**, and then click **Port**.
- Right-click and from the shortcut menu, choose **Graph**.

The Port dialog box for a single port or for multiple ports appears with the Interface tab visible.

3. Click the **EAPoL Diag** tab.
4. If you selected multiple ports, from the Graph Port EAPoL Stats tab Show list, select: Absolute Value, Cumulative, Average/sec, Minimum/sec, Maximum/sec, or LastVal/sec to view the graph for multiple ports.

EAPoL Diag field descriptions

The following table describes fields on the **EAPoL Diag** tab.

Name	Description
EntersConnecting	Counts the number of times that the Authenticator PAE state machine transitions to the Connecting state from any other state.
EapLogoffsWhileConnecting	Counts the number of times that the Authenticator PAE state machine transitions

Name	Description
	from Connected to Disconnected as a result of receiving an EAPoL-Logoff message.
EntersAuthenticating	Counts the number of times that the Authenticator PAE state machine transitions from Connecting to Authenticating as a result of receiving an EAP-Response/Identity message received from the Supplicant.
AuthSuccessWhileAuthenticating	Counts the number of times that the Authenticator PAE state machine transitions from Authenticating to Authenticated as a result of the Backend authentication state machine indicating successful authentication of the Supplicant.
AuthTimeoutsWhileAuthenticating	Counts the number of times that the Authenticator PAE state machine transitions from Authenticating to Aborting as a result of the Backend authentication state machine indicating authentication timeout.
AuthFailWhileAuthenticating	Counts the number of times that the Authenticator PAE state machine transitions from Authenticating to Held as a result of the Backend authentication state machine indicating authentication failure.
AuthReauthsWhileAuthenticating	Counts the number of times that the Authenticator PAE state machine transitions from Authenticating to Aborting as a result of a reauthentication request.
AuthEapStartsWhileAuthenticating	Counts the number of times that the Authenticator PAE state machine transitions from Authenticating to Aborting as a result of an EAPoL-Start message received from the Supplicant.
AuthEapLogoffWhileAuthenticating	Counts the number of times that the Authenticator PAE state machine transitions from Authenticating to Aborting as a result of an EAPoL-Logoff message received from the Supplicant.
AuthReauthsWhileAuthenticated	Counts the number of times that the Authenticator PAE state machine transitions from Authenticated to Connecting as a result of a reauthentication request.
AuthEapStartsWhileAuthenticated	Counts the number of times that the Authenticator PAE state machine transitions from Authenticated to Connecting as a result of an EAPoL-Start message received from the Supplicant.

Name	Description
AuthEapLogoffWhileAuthenticated	Counts the number of times that the Authenticator PAE state machine transitions from Authenticated to Disconnected as a result of an EAPoL-Logoff message received from the Supplicant.
BackendResponses	Counts the number of times that the Backend Authentication state machine sends an Initial-Access request packet to the Authentication server.
BackendAccessChallenges	Counts the number of times that the Backend Authentication state machine receives an Initial-Access challenge packet from the Authentication server.
BackendOtherRequestsToSupplicant	Counts the number of times that the Backend Authentication state machine sends an EAP request packet (other than an Identity, Notification, failure, or success message) to the Supplicant.
BackendNonNakResponsesFromSupplicant	Counts the number of times that the Backend Authentication state machine receives a response from the supplicant to an initial EAP request and the response is something other than EAP-NAK.
BackendAuthSuccesses	Counts the number of times that the Backend Authentication state machine receives an EAP-success message from the Authentication server.
BackendAuthFails	Counts the number of times that the Backend Authentication state machine receives an EAP-failure message from the Authentication server.

Viewing EAPoL session statistics

Use EAPoL session statistics to display the Authenticator PAE statistics for each session that is still in progress and the final values for ports where no session is currently active.

Procedure

1. On the Device Physical View, select a port.

If you want to select multiple ports, press Ctrl and hold down the key while you select the ports you want to configure. A yellow outline appears around the selected ports.

2. Perform one of the following steps:

- In the navigation pane, expand the following folders: **Configuration > Graph**, and then click **Port**.
- Right-click and from the shortcut menu, choose **Graph**.

The Port dialog box for a single port or for multiple ports appears with the Interface tab visible.

3. Click **EAPOL Session**.

4. If you selected multiple ports, from the Graph ports EAPoL Session tab Show options box, select Absolute Value, Cumulative, Average/sec, Minimum/sec, Maximum/sec, or LastVal/sec to view the graph for multiple ports.

EAPOL Session field descriptions

The following table describes parameters on the **EAPoL Sessions** tab.

Name	Description
SessionOctetsRx	Displays the number of octets received in user data frames on this port during the session.
SessionOctetsTx	Displays the number of octets transmitted in user data frames on this port during the session.
SessionFramesRx	Displays the number of user data frames received on this port during the session.
SessionFramesTx	Displays the number of user data frames transmitted on this port during the session.

Showing the Authenticator session statistics

Use Authenticator Session Statistics to display the session statistics objects for the Authenticator PAE associated with each port.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > Security > Data Path**.
2. Click **802.1x-EAPOL**.
3. Select the **Authentication Sessions** tab.

Authentication Sessions field descriptions

The following table describes values on the **Authentication Sessions** tab.

Name	Description
PortNumber	Indicates the Port number associated with this Port.
SessionId	Specifies a unique identifier for the session, in the form of a printable ASCII string of at least three characters.

Name	Description
SessionAuthenticMethod	Indicates the authentication method used to establish the session.
SessionTime	Indicates the duration of the session in seconds.
SessionTerminateCause	Indicates the reason for the session termination.
SessionUserName	Indicates the User-Name that represents the identity of the Supplicant PAE.
LastEapolFrameVersion	Indicates the protocol version number carried in the most recently received EAPOL frame.
LastEapolFrameSource	Indicates the source MAC address carried in the most recently received EAPOL frame.

Showing RADIUS server statistics

Use the server statistics feature to display the number of input and output packets and the number of input and output bytes. Statistics from console ports are available to assist with debugging.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > Security > Control Path**.
2. Click **RADIUS**.
3. Click the **RADIUS Servers Stats** tab.

RADIUS Server Stats field descriptions

Use the data in the following table to use the **RADIUS Server Stats** tab.

Name	Description
AddressType	Specifies either an IPv4 or an IPv6 address. RADIUS supports IPv4 and IPv6 addresses.
Address	Shows the IP address of the RADIUS server.
Used by	Identifies the client.
AccessRequests	Shows the number of access-response packets sent to the server; does not include retransmissions.
AccessAccepts	Shows the number of access-accept packets, valid or invalid, received from the server.
AccessRejects	Shows the number of access-reject packets, valid or invalid, received from the server.
BadResponses	Shows the number of invalid access-response packets received from the server.
PendingRequests	Shows the access-request packets sent to the server that have not yet received a response or that have timed out.

Name	Description
ClientRetries	Shows the number of authentication retransmissions to the server.
AcctOnRequests	Shows the number of accounting on requests sent to the server.
AcctOffRequests	Shows the number of accounting off requests sent to the server.
AcctStartRequests	Shows the number of accounting start requests sent to the server.
AcctStopRequests	Shows the number of accounting stop requests sent to the server.
AcctInterimRequests	Number of Accounting Interim requests sent to the server. Important: The AcctInterimRequests counter increments only if you select AcctIncludeCli from the RADIUS Global tab.
AcctBadResponses	Shows the number of Invalid responses discarded from the server.
AcctPendingRequests	Shows the number of requests waiting to be sent to the server.
AcctClientRetries	Shows the number of retries made to this server.
RoundTripTime	Shows the time difference between the instance when a RADIUS request is sent and the corresponding response is received.
AccessChallenges	Shows the number of RADIUS access-challenges packets sent to this server. This does not include retransmission.
NasIpAddress	Shows the RADIUS client NAS Identifier for this server.

Showing SNMP statistics

Display SNMP statistics to monitor the number of specific error messages, such as the number of messages that were delivered to SNMP but were not allowed.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > Security > Control Path**.
2. Click **General**.
3. Click the **SNMP** tab.

SNMP field descriptions

Use the data in the following table to display SNMP statistics.

Name	Description
OutTooBig	Shows the number of SNMP PDUs that the SNMP protocol entity generated and for which the value of the error-status field is tooBig.
OutNoSuchNames	Shows the number of SNMP PDUs that the SNMP protocol entity generated and for which the value of the error-status is noSuchName.

Name	Description
OutBadValues	Shows the number of SNMP PDUs that SNMP protocol entity generated and for which the value of the error-status field is badValue.
OutGenErrors	Shows the number of SNMP PDUs that the SNMP protocol entity generated and for which the value of the error-status field is genErr.
InBadVersions	Shows the number of SNMP messages that were delivered to the SNMP protocol entity and were for an unsupported SNMP version.
InBadCommunityNames	Shows the number of SNMP messages delivered to the SNMP protocol entity that used an SNMP community name not known to the entity.
InBadCommunityUsers	Shows the number of SNMP messages delivered to the SNMP protocol entity that represented an SNMP operation not allowed by the SNMP community named in the message.
InASNParseErrs	Shows the number of ASN.1 or BER errors encountered by the SNMP protocol entity when decoding received SNMP messages.
InTooBigs	Shows the number of SNMP PDUs that were delivered to the SNMP protocol entity and for which the value of the error-status field is tooBig.
InNoSuchNames	Shows the number of SNMP PDUs that were delivered to the SNMP protocol entity and for which the value of the error-status field is noSuchName.
InBadValues	Shows the number of SNMP PDUs that were delivered to the SNMP protocol entity and for which the value of the error-status field is badValue.
InReadOnlys	Shows the number of valid SNMP PDUs delivered to the SNMP protocol entity and for which the value of the error-status field is "read-only". It is a protocol error to generate an SNMP PDU that contains the value "read-only" in the error-status field; this object is provided as a means of detecting incorrect implementations of the SNMP.
InGenErrors	Shows the number of SNMP PDUs delivered to the SNMP protocol entity and for which the value of the error-status field is "genErr."

Viewing PCAP stats

View PCAP statistics to manage network performance.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **PCAP**.
3. Click the **PcapStat** tab.

PcapStat field descriptions

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

Name	Description
ResetStat	Resets statistics when selected.
PacketCapacityCount	Shows the packet capacity count.
NumberOfPacketsReceived	Shows the number of packets received in the PCAP engine.
NumberOfPacketsAccumulated	Shows the number of packets captured in the PCAP engine.
NumberOfPacketsDroppedInPcapEngine	Shows the number of packets dropped in the PCAP engine by filters.
NumberOfPacketsDroppedInHardware	Shows the number of packets dropped in hardware.

Enabling multicast routing process statistics

Enable the collection of multicast routing process statistics. The statistics display based on VRF and multicast group classification. These statistics are not related to the interface (port) statistics. (To display the collected statistics, you must use ACLI.)

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **Multicast**.
3. Click the **Globals** tab.
4. Select the **MrouteStatsEnabled** check box, and then click **Apply**.

Viewing multicast routing process statistics

These statistics are not related to the interface (port) statistics. Rather, the system displays the statistics based on multicast group classification.

Before you begin

- You must enable the collection of multicast statistics.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**
2. Click **Multicast**.
3. Click the **Routes** tab.
4. Select an S, G mroute.

5. Click **Graph**.
6. **(Optional)** To clear all mroute statistics counters, including the AbsoluteValue counter, click **Reset All Mroute Counters**, and then click **Yes**.

Stats field descriptions

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

Name	Description
SourceCount	Specifies the source number that corresponds to the associated group IP address in the whole multicast route records.
IngressPkts	Specifies the number of normally forwarded packets for the associated group IP address.
IngressBytes	Specifies the number of normally forwarded bytes for the associated group IP address.
AverageSizePerPkt	Specifies the average packet length for the associated group IP address. This information indicates only the forward packet length and is calculated using the following formula: forward packet/forward byte.
DropPkts	Specifies the number of dropped packets for the associated group IP address.
DropBytes	Specifies the number of dropped bytes for the associated VRF and group IP address.
PktsPerSecond	Specifies the average speed. This field is only valid in the monitor output. The value is calculated using the following formula: (current forward packet – last forward packet)/ monitor interval. With the first monitor multicast statistics output, this field is null. Subsequent outputs provide valid values.

Viewing IPFIX hash statistics

View the hashing statistics to view total hash overflows.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > Serviceability**.
2. Click **IPFIX**.
3. Click the **Exporters/Slots** tab.
4. Select an exporter slot.
5. Click **Graph**.

Slot Hash field descriptions

Use the data in the following table to use the **Slot Hash** tab.

Name	Description
HashOverflows	Shows the count of hash overflows for the slot.
HashDrops	Shows the count of hash drops for the slot.

Viewing IPFIX exporter statistics

View the exporter statistics for each slot to see the following information:

- collector IP address
- packets sent since you enabled IPFIX
- bytes sent since you enabled IPFIX
- packets lost within the device
- IPFIX protocol status

Procedure

1. In the navigation pane, expand the following folders: **Configuration > Serviceability**.
2. Click **IPFIX**.
3. Click the **Collectors/Slots** tab.
4. Select a slot.
5. Click **Graph**.

Exporter field descriptions

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

Name	Description
OutPkts	Shows the number of packets sent since you enabled IPFIX on the slot.
OutOctets	Shows the number of bytes sent since you enabled IPFIX on the slot.
PktsLoss	Shows the number of packets (records) lost within the device.

Displaying IS-IS system statistics

Use the following procedure to display Intermediate-System-to-Intermediate-System (IS-IS) system statistics.

Procedure

1. In the navigation pane, choose **Configuration > IS-IS**.
2. Click **Stats**.
3. Click the **System Stats** tab.

System Stats field descriptions

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

Name	Description
CorrLSPs	Indicates the number of corrupted in-memory link-state packets (LSPs) detected. LSPs received from the wire with a bad checksum are silently dropped and not counted.
AuthFails	Indicates the number of authentication key failures recognized by this Intermediate System.
LSPDbaseOloads	Indicates the number of times the LSP database has become overloaded.
ManAddrDropFromAreas	Indicates the number of times a manual address has been dropped from the area.
AtmptToExMaxSeqNums	Indicates the number of times the IS has attempted to exceed the maximum sequence number.
SeqNumSkips	Indicates the number of times a sequence number skip has occurred.
OwnLSPPurges	Indicates the number of times a zero-aged copy of the system's own LSP is received from some other node.
IDFieldLenMismatches	Indicates the number of times a PDU is received with a different value for ID field length to that of the receiving system.
PartChanges	Indicates partition changes.
AbsoluteValue	Displays the counter value.
Cumulative	Displays the total value since you expanded the Stats tab.
Average/sec	Displays the average value for each second.
Minimum/sec	Displays the minimum value for each second.
Maximum/sec	Displays the maximum value for each second.
LastVal/sec	Displays the last value for each second.

Displaying IS-IS interface counters

Use the following procedure to display IS-IS interface counters.

Procedure

1. From the navigation pane, choose **Configuration > IS-IS**.

2. Click **Stats**.
3. Click the **Interface Counters** tab.

Interface Counters field descriptions

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

Name	Description
Index	Shows a unique value identifying the IS-IS interface.
Type	Shows the type of interface.
AdjChanges	Shows the number of times an adjacency state change has occurred on this circuit.
InitFails	Shows the number of times initialization of this circuit has failed. This counts events such as PPP NCP failures. Failures to form an adjacency are counted by isisCircRejAdjs.
RejAdjs	Shows the number of times an adjacency has been rejected on this circuit.
IDFieldLenMismatches	Shows the number of times an IS-IS control PDU with an ID field length different to that for this system has been received.
MaxAreaAddrMismatches	Shows the number of times an IS-IS control PDU with a max area address field different to that for this system has been received.
AuthFails	Shows the number of times an IS-IS control PDU with the correct auth type has failed to pass authentication validation.
LANDesISChanges	Shows the number of times the Designated IS has changed on this circuit at this level. If the circuit is point to point, this count is zero.

Displaying IS-IS interface control packets

Use the following procedure to display IS-IS interface control packets.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > IS-IS**.
2. Click **Stats**.
3. Click the **Interface Control Packets** tab.

Interface Control Packets field descriptions

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

Name	Description
Index	Shows a unique value identifying the Intermediate-System-to-Intermediate-System (IS-IS) interface.

Name	Description
Level	Shows the level at which the system collects the counts.
Direction	Indicates whether the switch is sending or receiving the PDUs.
Hello	Indicates the number of IS-IS Hello frames seen in this direction at this level.
LSP	Indicates the number of IS-IS LSP frames seen in this direction at this level.
CSNP	Indicates the number of IS-IS Complete Sequence Number Packets (CSNP) frames seen in this direction at this level.
PSNP	Indicates the number of IS-IS Partial Sequence Number Packets (PSNP) frames seen in this direction at this level.

Graphing IS-IS interface counters

Use the following procedure to graph IS-IS interface counters.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > IS-IS**.
2. Click **IS-IS**.
3. Click the **Interfaces** tab.
4. Select an existing interface.
5. Click the **Graph** button.

Interface Counters field descriptions

The following table describes the fields in the **Interface Counters** tab.

Name	Description
InitFails	Indicates the number of times initialization of this circuit has failed. This counts events such as PPP NCP failures.
RejAdjs	Indicates the number of times an adjacency has been rejected on this circuit.
IDFieldLenMismatches	Indicates the number of times an Intermediate-System-to-Intermediate-System (IS-IS) control PDU with an ID field length different from that for this system has been received.
MaxAreaAddrMismatches	Indicates the number of times an IS-IS control PDU with a max area address field different from that for this system has been received.
AuthFails	Indicates the number of times an IS-IS control PDU with the correct auth type has failed to pass authentication validation.

Name	Description
LANDesISChanges	Indicates the number of times the Designated IS has changed on this circuit at this level. If the circuit is point to point, this count is zero.
AbsoluteValue	Displays the counter value.
Cumulative	Displays the total value since you expanded the Stats tab.
Average/Sec	Displays the average value for each second.
Minimum/Sec	Displays the minimum value for each second.
Maximum/Sec	Displays the maximum value for each second.
Last Val/Sec	Displays the last value for each second.

Graphing IS-IS interface sending control packet statistics

Use the following procedure to graph IS-IS interface receiving control packet statistics.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > IS-IS**.
2. Click **IS-IS**.
3. Click the **Interfaces** tab.
4. Select an existing interface.
5. Click the **Graph** button.
6. Click the **Interface Sending Control Packets** tab.

Interface Sending Control Packets field descriptions

The following table describes the fields in the **Interface Sending Control Packets** tab.

Name	Description
Hello	Indicates the number of IS-IS Hello (IIH) PDUs seen in this direction at this level. Point-to-Point IIH PDUs are counted at the lowest enabled level: at L1 on L1 or L1L2 circuits, and at L2 otherwise.
LSP	Indicates the number of IS-IS LSP frames seen in this direction at this level.
CSNP	Indicates the number of IS-IS Complete Sequence Number Packet (CSNP) frames seen in this direction at this level.
PSNP	Indicates the number of IS-IS Partial Sequence Number Packets (PSNPs) seen in this direction at this level.
AbsoluteValue	Displays the counter value.
Cumulative	Displays the total value since you expanded the Stats tab.
Average/Sec	Displays the average value for each second.

Name	Description
Minimum/Sec	Displays the minimum value for each second.
Maximum/Sec	Displays the maximum value for each second.
Last Val/Sec	Displays the last value for each second.

Graphing IS-IS interface receiving control packet statistics

Use the following procedure to graph IS-IS interface sending control packet statistics.

Procedure

1. From the navigation pane, choose **Configuration > IS-IS**.
2. Click **IS-IS**.
3. Click the **Interfaces** tab.
4. Select an existing interface.
5. Click the **Graph** button.
6. Click the **Interface Receiving Control Packets** tab.

Interface Receiving Control Packets field descriptions

The following table describes the fields in the **Interface Receiving Control Packets** tab.

Name	Description
Hello	Indicates the number of IS-IS Hello PDUs seen in this direction at this level. Point-to-Point IIH PDUs are counted at the lowest enabled level: at L1 on L1 or L1L2 circuits, and at L2 otherwise.
LSP	Indicates the number of IS-IS link-state packet (LSP) frames seen in this direction at this level.
CSNP	Indicates the number of IS-IS Complete Sequence Number Packet (CSNP) frames seen in this direction at this level.
PSNP	Indicates the number of IS-IS Partial Sequence Number Packets (PSNPs) seen in this direction at this level.
AbsoluteValue	Displays the counter value.
Cumulative	Displays the total value since you expanded the Stats tab.
Average/Sec	Displays the average value for each second.
Minimum/Sec	Displays the minimum value for each second.
Maximum/Sec	Displays the maximum value for each second.
Last Val/Sec	Displays the last value for each second.

Displaying SPBM packet drop statistics by port

Use this procedure to display the SPBM drop statistics.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > IS-IS**.
2. Click **SPBM**.
3. Click the **Drop Stats** tab.

Drop Stats field descriptions

Use the data in the following table to use the **Drop Stats** tab.

Name	Description
PortIndex	Shows the slot/port number that identifies the ingress port of the dropped packet.
VlanId	Shows the VLAN ID.
VlanType	Shows the VLAN type as either primary or secondary.
UnknownUcastSrcAddr	Shows the total number unknown source unicast packets.
RpfcUcastSrcAddr	Shows the total number of RPFC source unicast packets.
UnknownUcastDestAddr	Shows the total number of unknown destination unicast packets.
Unknown McastDesAddr	Shows the total number of unknown multicast destination packets.
RpfcMcastSrcAddr	Shows the total number of RPFC multicast source packets.
LastDropMac	Shows the time of the last drop backbone MAC.
IsMacDestAddr	Shows the IS MAC destination address.
LastDropMacHostName	Shows the last drop MAC host name.

Resetting AbsoluteValues counter using EDM

Clear the AbsoluteValues to help monitor the performance of the Avaya Virtual Services Switch 9000.

About this task

EDM resets multicast route (mroute) statistics counters in a different manner than other statistics counters. To reset statistics counters for mroutes, see [Viewing multicast routing process statistics](#) on page 318. For all other statistics counters, click **Clear Counters** to reset the counters.

After you click this button, all Cumulative, Average, Minimum, Maximum, and LastVal columns reset to zero, and automatically begin to recalculate statistical data.

Important:

The **Clear Counters** function does not affect the AbsoluteValue counter for the device. The **Clear Counters** function clears all cached data in EDM except AbsoluteValue. Perform the following steps to reset AbsoluteValues.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit**.
2. Click **Chassis**.
3. Click the **System** tab.
4. In ActionGroup1, select **resetCounters**, and then click **Apply**.

Glossary

American Standard Code for Information Interchange (ASCII)	A code to represent characters in computers. ASCII uses uppercase and lowercase alphabetic letters, numeric digits, and special symbols.
Application Programming Interface (API)	Defines how to access a software-based service. An API is a published specification that describes how other software programs can access the functions of an automated service.
Autonomous System (AS)	A set of routers under a single technical administration, using a single IGP and common metrics to route packets within the AS, and using an EGP to route packets to other ASs.
Autonomous System Number (ASN)	A two-byte number that is used to identify a specific AS.
Avaya command line interface (ACLI)	A textual user interface. When you use ACLI, you respond to a prompt by typing a command. After you enter the command, you receive a system response.
bit error rate (BER)	The ratio of the number of bit errors to the total number of bits transmitted in a specific time interval.
Bootstrap Protocol (BootP)	A User Datagram Protocol (UDP)/Internet Protocol (IP)-based protocol that a booting host uses to configure itself dynamically and without user supervision.
Bridge Protocol Data Unit (BPDU)	A data frame used to exchange information among the bridges in local or wide area networks for network topology maintenance.
Collecting process	A process that receives flow records from one or more exporting processes. The collecting process can process or store received flow records.
Collector	A device that hosts one or more collecting processes.
Control Processor Unit High Availability (CPU-HA)	CPU-HA activates two CP modules simultaneously. The CP modules exchange topology data so, if a failure occurs, either CP module can take precedence in less than one second with the most recent topology data.

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.
Data flowset	One or more records, of the same type, in an export packet. Each record is either a flow data record or an options data record previously defined by a template record or an options template record.
Dynamic Random Access Memory (DRAM)	A read-write random-access memory, in which the digital information is represented by charges stored on the capacitors and must be repeatedly replenished to retain the information.
Enterprise Device Manager (EDM)	A Web-based embedded management system to support single-element management. EDM provides complete configuration management functionality for the supported devices and is supplied to the customer as embedded software in the device.
Exporting process	An export process that sends flow records to one or more collecting processes. One or more metering processes generate the flow records.
External Data Representation (XDR)	An IETF standard, RFC 1832, for the description and encoding of data.
Flow key	A field used to define a flow is termed a flow key. A flow key is each field that belongs to the packet header (for example, destination IP address), is a property of the packet itself (for example, packet length), or is derived from packet treatment (for example, AS number).
Flow record	A flow record contains information about a specific flow that was observed at an observation point. The flow record contains measured properties of the flow, for example, the total number of bytes for all packets in the flow, and characteristic properties of the flow, for example, source IP address.
Flowset	A generic term for a collection of flow records that use a similar structure. In an export packet, one or more flowsets follow the packet header. Three flow sets are available: template flowset, options template flowset, and data flowset.
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.
Frame Check Sequence (FCS)	Frames are used to send upper-layer data and ultimately the user application data from a source to a destination.
graphical user interface (GUI)	A graphical (rather than textual) computer interface.

Intermediate System to Intermediate System (IS-IS)	<p>Intermediate System to Intermediate System(IS-IS) is a link-state, interior gateway protocol. ISO terminology refers to routers as Intermediate Systems (IS), hence the name Intermediate System to Intermediate System (IS-IS). IS-IS operation is similar to Open Shortest Path First (OSPF).</p> <p>In Shortest Path Bridging MAC (SPBM) networks, IS-IS discovers network topology and builds shortest path trees between network nodes that IS-IS uses for forwarding unicast traffic and determining the forwarding table for multicast traffic. SPBM employs IS-IS as the interior gateway protocol and implements additional Type-Length-Values (TLVs) to support additional functionality.</p>
Internet Control Message Protocol (ICMP)	<p>A collection of error conditions and control messages exchanged by IP modules in both hosts and gateways.</p>
Internet Group Management Protocol (IGMP)	<p>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.</p>
Internet Protocol Flow Information eXport (IPFIX)	<p>An IETF standard that improves the Netflow V9 protocol. IPFIX monitors IP flows.</p>
Internet Protocol Flow Information eXport (IPFIX) device	<p>A device that hosts at least one observation point, a metering process, and an exporting process. Typically, corresponding observation points, metering processes, and exporting processes are located at the same device, for example, at a router.</p>
Internet Protocol Flow Information eXport (IPFIX) node	<p>A host that implements the Internet Protocol Flow Information eXport (IPFIX) protocol; that is, it can contain an exporting process, a collecting process, or both.</p>
Internet Protocol traffic flow or flow	<p>A set of Internet Protocol (IP) packets that pass an observation point in the network during a certain time interval. All packets that belong to a particular flow have a group of common properties. In the Avaya IPFIX implementation, IP SRC, IP DST, IP Protocol, SrcPort, Dst port and observation point uniquely define a flow.</p>
interswitch trunking (IST)	<p>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.</p>
Layer 2	<p>Layer 2 is the Data Link Layer of the OSI model. Examples of Layer 2 protocols are Ethernet and Frame Relay.</p>

Layer 3	Layer 3 is the Network Layer of the OSI model. An example of a Layer 3 protocol is Internet Protocol (IP).
Link Aggregation Control Protocol (LACP)	A protocol that exists between two endpoints to bundle links into an aggregated link group for bandwidth increase and link redundancy.
Link Aggregation Control Protocol Data Units (LACPDU)	Link aggregation control protocol data unit (LACPDU) is used for exchanging information among LACP-enabled devices.
link-state advertisement (LSA)	Packets that contain state information about directly connected links (interfaces) and adjacencies. Each Open Shortest Path First (OSPF) router generates the packets.
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.
Logical Link Control (LLC)	A protocol used in LANs to transmit protocol data units between two end stations. This LLC layer addresses and arbitrates data exchange between two endpoints.
management information base (MIB)	The MIB defines system operations and parameters used for the Simple Network Management Protocol (SNMP).
media	A substance that transmits data between ports; usually fiber optic cables or category 5 unshielded twisted pair (UTP) copper wires.
Metering process	A process that generates flow records. An input to the process is packets observed at an observation point and packet treatment at the observation point. The metering process consists of a set of functions that includes packet header capturing, time stamping, sampling, classifying, and maintaining flow records. The maintenance of flow records can include creating new records, updating existing records, computing flow statistics, deriving further flow properties, detecting flow expiration, passing flow records to the exporting process, and deleting flow records.
multiplexing	Carriage of multiple channels over a single transmission medium; a process where a dedicated circuit is shared by multiple users. Typically, data streams intersperse on a bit or byte basis (time division), or separate by different carrier frequencies (frequency division).
nanometer (nm)	One billionth of a meter (10^{-9} meter). A unit of measure commonly used to express the wavelengths of light.

NonVolatile Random Access Memory (NVRAM)	Random Access Memory that retains its contents after electrical power turns off.
Observation domain	The set of observation points that is the largest set of flow information that can be aggregated at the metering process. Each observation domain uses a unique ID for the export process to identify the IPFIX messages it generates. For example, a router interface module can comprise several interfaces with each interface being an observation point. Every observation point is associated with an observation domain.
Observation point	An observation point is a network location where you can observe IP packets. Examples include a port or a VLAN.
Options data record	The data record that contains values and scope information of the flow measurement parameters that correspond to an options template record.
Options template flowset	One or more options template records in an export packet.
Options template record	A record that defines the structure and interpretation of fields in an options data record, including defining the scope within which the options data record is relevant.
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.
policing	Ensures that a traffic stream follows the domain service-provisioning policy or service-level agreement (SLA).
Port Access Entity (PAE)	Software that controls each port on the switch. The PAE, which resides on the device, supports authenticator functionality. The PAE works with the Extensible Authentication Protocol over LAN (EAPoL).
Protocol Data Units (PDUs)	A unit of data that is specified in a protocol of a specific layer and that consists of protocol-control information of the specific layer and possibly user data of that layer.
Protocol Independent Multicast, Source Specific (PIM-SSM)	PIM-SSM is a multicast routing protocol for IP networks. PIM-SSM uses only shortest-path trees to provide multicast services based on subscription to a particular (source, group) channel. PIM-SSM eliminates the need for starting with a shared tree by immediately joining a source through the shortest path tree. This method enables PIM-SSM to avoid using a rendezvous point (RP) and RP-based shared tree, which can be a potential bottleneck.
Protocol Independent	PIM-SM is a multicast routing protocol for IP networks. PIM-SM provides multicast routing for multicast groups that can span wide-area and inter-

Multicast, Sparse Mode (PIM-SM)	domain networks, where receivers are not densely populated. PIM-SM sends multicast traffic only to those routers that belong to a specific multicast group and that choose to receive the traffic. PIM-SM adds a Rendezvous Point router to avoid multicast-data flooding. Use PIM-SM when receivers for multicast data are sparsely distributed throughout the network.
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.
Random Access Memory (RAM)	Memory into which you can write and read data. A solid state memory device used for transient memory stores. You can enter and retrieve information from storage position.
remote login (rlogin)	An application that provides a terminal interface between hosts (usually UNIX) that use the TCP/IP network protocol. Unlike Telnet, rlogin assumes the remote host is, or behaves like, a UNIX host.
remote monitoring (RMON)	A remote monitoring standard for Simple Network Management Protocol (SNMP)-based management information bases (MIB). The Internetwork Engineering Task Force (IETF) proposed the RMON standard to provide guidelines for remote monitoring of individual LAN segments.
sFlow agent	Provides the interface for the sFlow instance. The agent maintains the measurement session with, and sends sFlow datagrams to, the sFlow collector.
sFlow collector	Receives sFlow datagrams from one or more sFlow agents.
sFlow datagram	A User Datagram Protocol (UDP) packet that contains the measurement information. The sFlow datagram also includes information about the source and process.
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.

shortest path first (SPF)	A class of routing protocols that use Dijkstra's algorithm to compute the shortest path through a network, according to specified metrics, for efficient transmission of packet data.
small form-factor pluggable (SFP)	A hot-swappable input and output enhancement component used with Avaya products to allow gigabit Ethernet ports to link with other gigabit Ethernet ports over various media types.
small form-factor pluggable plus (SFP+)	SFP+ transceivers are similar to SFPs in physical appearance but SFP+ transceivers provide Ethernet at 10 gigabits per second (Gbps).
spanning tree	A simple, fully-connected active topology formed from the arbitrary physical topology of connected bridged Local Area Network components by relaying frames through selected bridge ports. The protocol parameters and states that are used and exchanged to facilitate the calculation of the active topology and to control the bridge relay function.
Spanning Tree Group (STG)	A collection of ports in one spanning-tree instance.
Template flowset	One or more options template records in an export packet.
Template record	An ordered list (for example, of <type, length>pairs) that identifies the structure and semantics of a particular set of information to communicate from an Internet Protocol Flow Information eXport (IPFIX) device to a collector. Each template is uniquely identifiable, for example, by using a template ID.
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.
traffic profile	The temporal properties of a traffic stream, such as rate.
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.
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.
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.