



# **Performance Management Avaya Virtual Services Platform 9000**

3.2  
NN46250-701, 03.02  
February 2012

## Notice

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

## Documentation disclaimer

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

## Link disclaimer

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

## Warranty

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

## Licenses

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

## Copyright

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

## Third-party components

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

## Trademarks

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

Avaya is a registered trademark of Avaya Inc.

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

## Downloading Documentation

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

## Contact Avaya Support

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

# Contents

<b>Chapter 1: Purpose of this document.....</b>	<b>7</b>
<b>Chapter 2: New in this release.....</b>	<b>9</b>
Features.....	9
Other changes.....	9
<b>Chapter 3: Performance management fundamentals.....</b>	<b>11</b>
Switch management tools.....	11
Dynamic network applications.....	12
Automatic trace.....	12
Digital diagnostic monitoring.....	13
Remote monitoring.....	13
Layer 2 redundancy and High Availability clarification.....	16
Internet Protocol Flow Information eXport.....	17
<b>Chapter 4: Chassis performance management using ACLI.....</b>	<b>23</b>
Viewing system performance.....	23
Configuring CPU-HA.....	23
<b>Chapter 5: Chassis performance management using EDM.....</b>	<b>25</b>
Viewing system performance.....	25
Configuring CPU-HA for Layer 3 redundancy.....	25
Viewing the trap sender table.....	26
<b>Chapter 6: RMON configuration using ACLI.....</b>	<b>27</b>
Configuring RMON.....	27
Viewing RMON settings.....	29
<b>Chapter 7: RMON configuration using EDM.....</b>	<b>31</b>
Enabling RMON globally.....	31
Enabling RMON history.....	32
Disabling RMON history.....	34
Viewing RMON history statistics.....	35
Creating an alarm.....	37
Creating a port history alarm.....	40
Viewing RMON alarms.....	40
Deleting an alarm.....	41
Creating a default RMON event.....	41
Creating a nondefault RMON event.....	42
Viewing RMON events.....	44
Deleting an event.....	45
<b>Chapter 8: IPFIX configuration using ACLI.....</b>	<b>47</b>
Enabling IPFIX globally.....	47
Configuring IPFIX metering using filters.....	47
Configuring IPFIX on a port.....	49
Configuring IPFIX slot parameters.....	50
Configuring collector parameters.....	51
Viewing flow information.....	52
Flushing IPFIX flow information.....	55
Viewing global IPFIX information.....	56

Viewing collector information.....	57
Viewing exporter information.....	57
Viewing IPFIX information for an interface.....	58
<b>Chapter 9: IPFIX configuration using EDM.....</b>	<b>61</b>
Enabling IPFIX globally.....	61
Configuring collector parameters.....	61
Configuring slot parameters.....	63
Configuring IPFIX on a port.....	64
Configuring IPFIX metering using filters.....	65
<b>Chapter 10: Port performance management using ACLI.....</b>	<b>67</b>
Configuring an automatic trace.....	67
Viewing DDI module information.....	69
Viewing DDI temperature information.....	70
Viewing DDI voltage information.....	71
<b>Chapter 11: Port performance management using EDM.....</b>	<b>73</b>
Configuring rate limits.....	73
Enabling learning limits on a port.....	74
Viewing DDI information.....	75
<b>Chapter 12: Viewing statistics using ACLI.....</b>	<b>79</b>
Viewing TCP statistics.....	79
Viewing port routing statistics.....	80
Displaying bridging statistics for specific ports.....	82
Displaying DHCP-relay statistics for specific ports.....	83
Displaying DHCP-relay statistics for all interfaces.....	84
Viewing IPv6 DHCP Relay statistics.....	87
Displaying LACP statistics for specific ports.....	88
Displaying RMON statistics for specific ports.....	89
Displaying detailed statistics for ports.....	91
Displaying policing statistics.....	92
Clearing ACL statistics.....	94
Viewing ACE statistics.....	95
Viewing MSTP statistics.....	96
Viewing RSTP statistics.....	97
Viewing RSTP port statistics.....	98
Viewing MLT statistics.....	100
Showing OSPF error statistics on a port.....	101
Viewing OSPF interface statistics.....	103
Viewing OSPF range statistics.....	105
Viewing basic OSPF statistics for a port.....	107
Showing extended OSPF statistics.....	108
Viewing IPv6 OSPF statistics.....	109
Showing the EAPoL status of the device.....	110
Showing EAPoL authenticator statistics.....	110
Showing EAPoL session statistics.....	111
Showing RADIUS server statistics.....	113
Viewing RMON statistics.....	115
Viewing PCAP statistics.....	116

Viewing IPFIX statistics.....	118
Clearing IPFIX statistics.....	119
Clearing IPv6 statistics.....	120
Viewing multicast routing process statistics.....	121
Viewing IPv6 VRRP statistics.....	123
Viewing ICMP statistics.....	126
Viewing IPv6 statistics on an interface.....	127
<b>Chapter 13: Viewing statistics using EDM.....</b>	<b>129</b>
Graphing chassis statistics.....	129
Graphing port statistics.....	130
Viewing chassis system statistics.....	131
Viewing chassis SNMP statistics.....	131
Viewing chassis IP statistics.....	133
Viewing chassis ICMP In statistics.....	135
Viewing chassis ICMP Out statistics.....	136
Viewing ICMP statistics.....	137
Viewing chassis TCP statistics.....	140
Viewing chassis UDP statistics.....	142
Configuring Switch Fabric statistics capture.....	143
Viewing Switch Fabric statistics.....	144
Viewing port interface statistics.....	147
Viewing port Ethernet errors statistics.....	149
Viewing port bridging statistics.....	151
Viewing port spanning tree statistics.....	152
Viewing port routing statistics.....	153
Viewing IPv6 statistics for an interface.....	153
Viewing DHCP statistics for an interface.....	157
Viewing IPv6 DHCP Relay statistics for a port.....	157
Graphing DHCP statistics for a port.....	158
Viewing DHCP statistics for a port.....	159
Graphing DHCP statistics for a VLAN.....	159
Displaying DHCP-relay statistics for Option 82.....	160
Viewing LACP port statistics.....	163
Viewing port policer statistics.....	164
Displaying file statistics.....	165
Viewing QoS policy statistics.....	166
Graphing QoS policy statistics.....	167
Viewing statistics for a specific QoS policy.....	168
Viewing ACE port statistics.....	169
Viewing ACL statistics.....	170
Clearing ACL statistics.....	171
Viewing VLAN and Spanning Tree CIST statistics.....	172
Viewing VLAN and Spanning Tree MSTI statistics.....	173
Viewing VRRP interface stats.....	174
Viewing VRRP statistics.....	175
Viewing IPv6 VRRP statistics for an interface.....	176
Viewing IPv6 VRRP statistics.....	178

Viewing SMLT statistics.....	178
Viewing RSTP status statistics.....	180
Viewing MLT interface statistics.....	181
Viewing MLT Ethernet error statistics.....	182
Viewing RIP statistics.....	184
Viewing OSPF chassis statistics.....	185
Viewing IPv6 OSPF statistics.....	186
Graphing OSPF statistics for a VLAN.....	187
Graphing OSPF statistics for a port.....	189
Viewing global BGP statistics.....	191
Viewing BGP peer general statistics.....	194
Viewing BGP peer advanced statistics.....	196
Viewing BGP peer receive statistics.....	197
Viewing BGP peer transmit statistics.....	200
Viewing statistics for a VRF.....	202
Viewing EAPoL Authenticator statistics.....	203
Viewing EAPoL diagnostic statistics.....	204
Viewing EAPoL session statistics.....	207
Showing the Authenticator session statistics.....	208
Showing RADIUS server statistics.....	209
Showing SNMP statistics.....	211
Viewing PCAP stats.....	212
Enabling RMON statistics.....	213
Viewing RMON statistics.....	214
Enabling multicast routing process statistics.....	216
Viewing multicast routing process statistics.....	217
Viewing IPFIX hash statistics.....	218
Viewing IPFIX exporter statistics.....	219
<b>Chapter 14: RMON alarm variables.....</b>	<b>221</b>
<b>Chapter 15: Customer service.....</b>	<b>243</b>
Getting technical documentation.....	243
Getting product training.....	243
Getting help from a distributor or reseller.....	243
Getting technical support from the Avaya Web site.....	243
<b>Glossary.....</b>	<b>245</b>

# Chapter 1: Purpose of this document

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

Purpose of this document



# Chapter 2: New in this release

The following sections detail what is new in *Avaya Virtual Services Platform 9000 Performance Management*, NN46250–701, for Release 3.2 .

---

## Features

See the following sections for information about feature-related changes.

### IPv6

Release 3.2 adds support for IPv6 routing. For information about IPv6 statistics, see [Viewing statistics using ACLI](#) on page 79 and [Viewing statistics using EDM](#) on page 129.

---

## Other changes

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

### Glossary

A glossary of related terms is added to the end of this document.

### Grouping of the chapters

The chapters are grouped according to the description of the feature to improve clarity. The Fundamentals chapter is placed in the beginning, followed by the configuration of the feature chapters by using ACLI and EDM interfaces.

### ACLI Commands

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

### Introduction chapter and navigation

Introduction chapter is renamed to Purpose of the document that states the purpose of referring to that document. Navigation is removed to reduce the length of the document.

### Terminology

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

New in this release

# Chapter 3: Performance management fundamentals

Performance management includes the management tools and features that are available to monitor and manage your routing switch. This section provides overviews for Simple Network Management Protocol (SNMP), Remote Monitoring (RMON), and Digital Diagnostic Monitoring.

---

## 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 *Avaya Virtual Services Platform 9000 User Interface Fundamentals, 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 *Avaya Virtual Services Platform 9000 User Interface Fundamentals, 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 *Avaya Virtual Services Platform 9000 User Interface Fundamentals, 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 *Avaya Virtual Services Platform 9000 Quick Start*, NN46250-102.

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

For more information about how to access the Web management interface, see *Avaya Virtual Services Platform 9000 User Interface Fundamentals*, NN46250-103.

---

## Automatic trace

You can configure the Virtual Services Platform 9000 to automatically enable a trace if CP usage reaches a predefined value. If you enable automatic trace, the CP usage increases by up to 30 percent.

If you enable or disable automatic trace, the information is not saved to the configuration file. After the Virtual Services Platform 9000 restarts, automatic trace is disabled.

---

## Digital diagnostic monitoring

Use Digital Diagnostic Monitoring (DDM) to monitor laser operating characteristics such as temperature, voltage, current, and power. This feature works at any time during active laser operation without affecting data traffic. There are two optical transceivers that support DDM: small form-factor pluggable (SFP) transceivers and 10 Gigabit small form-factor pluggable plus (SFP+).

Digital Diagnostic Interface (DDI) is an interface that supports DDM. These devices provide real-time monitoring of individual DDI SFPs and SFP+s 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 SFPs and SFP+s, see *Avaya Virtual Services Platform 9000 — SFP Hardware Components, NN46250-305*.

---

## Remote monitoring

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 1: 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  $\pm 1$  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 2: 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.

### **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.

---

## **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.



---

## 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.

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.

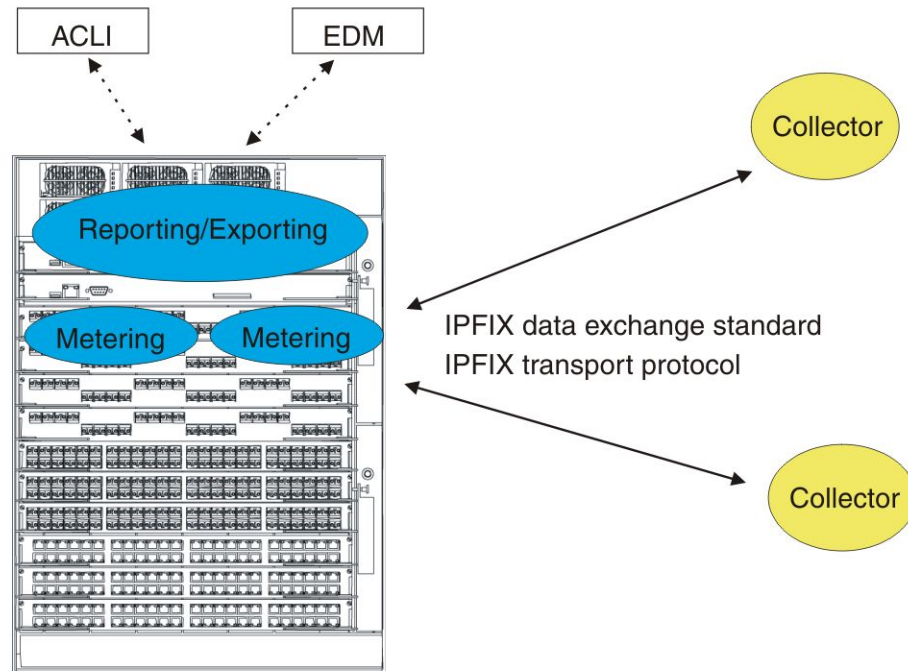
### Applications

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 3: IPFIX collecting and reporting functions**

**Common terminology and concepts**

The following table provides definitions for common IPFIX concepts.

**Table 1: IPFIX concepts and terminology**

Term	Description
Observation point	An observation point is a network location where you can observe IP packets. Examples include a port or a VLAN.
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

Term	Description
	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.
IP traffic flow or flow	A set of 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.
Flow key	<p>Each field that</p> <ul style="list-style-type: none"> <li>• belongs to the packet header (for example, destination IP address)</li> <li>• is a property of the packet itself (for example, packet length)</li> <li>• is derived from packet treatment (for example, AS number)</li> </ul> <p>A field used to define a flow is termed a flow key.</p>
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.
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.

Term	Description
Exporting process	An export process that sends flow records to one or more collecting processes. One or more metering processes generate the flow records.
IPFIX device	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.
IPFIX node	A host that implements the IPFIX protocol; that is, it can contain an exporting process, a collecting process, or both.
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.
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 IPFIX device to a collector. Each template is uniquely identifiable, for example, by using a template ID.
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.
Options data record	The data record that contains values and scope information of the flow measurement parameters that correspond to an options template record.
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.
Template flowset	One or more template records in an export packet.

Term	Description
Options template flowset	One or more options template records in an export packet.
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.

### Virtual Services Platform 9000 implementation

Use IPFIX on the Virtual Services Platform 9000 to perform the following actions:

- Configure IPFIX using either ACLI or Enterprise Device Manager (EDM).
  - 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

Flow records expire after the following conditions:

- export after the activity timer expires
- export after record ages out regardless of activity
- no flow record updates received for an aging time period
- A TCP RST or FIN received for a TCP flow

If the operational state of a port is down, the module marks this information in all the flow records that correspond to this port and exports this information to collector

Virtual Services Platform 9000 supports the following observation points:

- ingress port
- filtered ingress port
- filtered ingress VLAN

The interface module is the observation domain from the perspective of the collector. Each interface module sends a unique value of domain Id (slot number) when it exports a flow record.

Both CP modules save the IPFIX configuration information in local tables but only the master CP sends the commands to the interface module. Because of this, both CP modules always contain the same IPFIX information so if a failover occurs, the backup CP can take over.

# Chapter 4: Chassis performance management using ACLI

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

---

## Viewing system performance

### About this task

For information about how to use Key Health Indicators functionality to view system performance, see *Avaya Virtual Services Platform 9000 Fault Management*, NN46250–703.

---

## Configuring CPU-HA

### Before you begin

- You must log on to Global Configuration mode to complete this procedure.

### About this task

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

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
```





# Chapter 5: 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

### About this task

For information about how to use Key Health Indicators functionality to view system performance, see *Avaya Virtual Services Platform 9000 Fault Management*, NN46250–703.

---

## Configuring CPU-HA for Layer 3 redundancy

### About this task

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

### Procedure

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

---

## Viewing the trap sender table

### About this task

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 tree, open the following folders: **Configuration > Edit**.
  3. Double-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
<b>RecvAddress</b>	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.
<b>SrcAddress</b>	Source IP address to use when sending traps. This IP address will be inserted into the source IP address field in the UDP trap packet.

# Chapter 6: RMON configuration using ACLI

This chapter explains how to use ACLI to configure RMON on the Avaya Virtual Services Platform 9000.

---

## Configuring RMON

### Before you begin

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

### About this task

Configure RMON functions on Virtual Services Platform 9000 to set alarms and capture events.

### Procedure

1. Enable RMON globally:

```
rmon
```

2. 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>]
```

3. 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}]]
```

---

### Example

```
VSP-9012:1(config)#rmon
```

```
VSP-9012:1(config)#rmon alarm 4 rcCliNumAccessViolations.0 10  
absolute rising-threshold 2 event 60000
```

```
VSP-9012:1(config)#rmon event 60534 community public description  
"Rising Event" log trap
```

## Variable definitions

Use the data in this table to use the `rmon` command.

**Table 2: Variable definitions**

Variable	Value
<pre>alarm &lt;1-65535&gt; WORD &lt;1-1536&gt; &lt;1-3600&gt; {absolute  delta} [falling-threshold &lt;-2147483647-2147483647&gt; event &lt;1-65535&gt;] [owner WORD&lt;1-127&gt;] [rising-threshold &lt;-2147483647-2147483647&gt; event &lt;1-65535&gt;]</pre>	<p>Create an alarm interface.</p> <ul style="list-style-type: none"> <li>• <code>&lt;1-65535&gt;</code> is the interface index number from 1–65535.</li> <li>• <code>WORD &lt;1-1536&gt;</code> is the variable name or OID, case sensitive (string length 1–1536).</li> <li>• <code>{absolute   delta}</code> is the sample type.</li> <li>• <code>rising-threshold &lt;-2147483648-2147483647&gt; [<code>&lt;event:1-65535&gt;</code>]</code> is the rising threshold (–2147483648–2147483647) and the rising event number (1–65535).</li> <li>• <code>falling-threshold &lt;-2147483648-2147483647&gt; [<code>&lt;event:1-65535&gt;</code>]</code> is the falling threshold (–2147483648–2147483647) and the falling event number (1–65535).</li> <li>• <code>owner WORD&lt;1-127&gt;</code> is the name of the owner (string length 1–127).</li> </ul> <p>Use the default operator to reset the RMON alarms to their default configuration:  <code>default rmon alarm &lt;65535&gt;</code></p> <p>Use the no operator to disable RMON alarms:  <code>no rmon alarm [&lt;1-65535&gt;]</code></p>
<pre>event &lt;1-65535&gt; [community WORD&lt;1-127&gt;] [description WORD&lt;1-127&gt;] [log] [owner WORD&lt;1-127&gt;] [trap] [trap_dest {{A.B.C.D}}] [trap_src {{A.B.C.D}}]</pre>	<p>Create an event.</p> <ul style="list-style-type: none"> <li>• <code>&lt;1-65535&gt;</code> is the event index number.</li> <li>• <code>[log]</code> display information about configured traps.</li> <li>• <code>[trap]</code> specify trap source and destination IP addresses.</li> <li>• <code>description WORD&lt;1-127&gt;</code> is the event description (string length 0–127).</li> <li>• <code>owner WORD&lt;1-127&gt;</code> is the name of the owner (string length 1–127).</li> <li>• <code>trap_src {A.B.C.D}</code> is the trap source ip address.</li> </ul>

Variable	Value
	<ul style="list-style-type: none"> <li>• trap_dest {A.B.C.D} is the trap destination ip address.</li> <li>• community WORD&lt;1-127&gt; is the event community (string length 1–127).</li> </ul> <p>Use the no operator to delete a RMON event: no rmon event [&lt;1-65535&gt;] [log ]</p>

## Viewing RMON settings

### About this task

View RMON settings to see information about alarms, statistics, events, or the status of RMON on Virtual Services Platform 9000.

### Procedure

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.

**Table 3: Variable definitions**

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.

# Chapter 7: 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

### About this task

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.

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 tree, open 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.
  8. Click **Apply**.
- 

---

## Options field descriptions

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

Name	Description
<b>Enable</b>	Enables RMON. If you select the Enable check box, the RMON agent starts immediately if the amount of memory

Name	Description
	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.
<b>UtilizationMethod</b>	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.
<b>TrapOption</b>	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.
<b>MemSize</b>	Specifies the RAM size, in bytes, available for RMON to use. The default value is 250 Kilobytes.

---

## Enabling RMON history

### About this task

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. After you configure history characteristics, you cannot modify them; you must delete the history and create another one.

### Procedure

1. In the navigation tree, open the following folders: **Configuration > Serviceability > RMON**.
2. Click **Control**.
3. In the **History** tab, click **Insert**.
4. In the **Port** box, click the ellipsis (...) button.
5. Select a port.
6. 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
<b>Index</b>	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.
<b>Port</b>	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).
<b>BucketsRequested</b>	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.
<b>BucketsGranted</b>	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

Name	Description
	<p>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 current value, the number of associated media-specific entries is allowed to grow.</p>
<b>Interval</b>	<p>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.</p>
<b>Owner</b>	<p>Specifies the entity that configured this entry and is using the assigned resources.</p>

---

## Disabling RMON history

### About this task

Disable RMON history on a port if you do not want to record a statistical sample from that port.

### Procedure

1. In the navigation tree, open 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

### About this task

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. In the Device Physical View, select a port.
  2. In the navigation tree, open the following folders: **Configuration > Graph**
  3. Double-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 4: Variable definitions**

Parameter	Description
<b>Utilization</b>	The best estimate of the mean physical layer network utilization on this interface during the sampling interval, in hundredths of a percent.
<b>Octets</b>	The total number of octets of data (including those in bad packets) received on the network (excluding framing bits but including FCS octets)
<b>Pkts</b>	The number of packets (including bad packets) received during this sampling interval.
<b>BroadcastPkts</b>	The number of good packets received during this sampling interval that were directed to the broadcast address.

Parameter	Description
<b>MulticastPkts</b>	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.
<b>DropEvents</b>	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.
<b>CRCAAlignErrors</b>	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).
<b>UndersizePkts</b>	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.
<b>OversizePkts</b>	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.
<b>Fragments</b>	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.
<b>Collisions</b>	<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</p>

Parameter	Description
	observed on any coax segments to which the repeater is connected.

---

## Creating an alarm

### Before you begin

- You must globally enable RMON.

### About this task

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.

### Procedure

1. In the navigation tree, open 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
<b>Index</b>	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.
<b>Interval</b>	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.
<b>Variable</b>	<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. Alarm variables exist in three formats, depending on the type:</p> <ul style="list-style-type: none"> <li>• A chassis, power supply, or fan-related alarm ends in x where the x index is hard-coded. No further information is required.</li> <li>• 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.</li> <li>• 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).</li> </ul> <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>
<b>SampleType</b>	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

Name	Description
	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.
<b>Value</b>	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.
<b>StartUpAlarm</b>	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 a single falling alarm is generated. You cannot modify this object if the associated alarmStatus object is equal to valid.
<b>Rising Threshold</b>	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.
<b>RisingEventIndex</b>	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. In particular, if this value is zero, no associated event is generated, as zero is not a valid event index. You cannot modify this object if the associated alarmStatus object is equal to valid. The default is 60534.
<b>FallingThreshold</b>	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

Name	Description
	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.
<b>FallingEventIndex</b>	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. In particular, if this value is zero, no associated event is generated, as zero is not a valid event index. You cannot modify this object if the associated alarmStatus object is equal to valid. The default is 60535.
<b>Owner</b>	Specifies the entity that configured this entry and is therefore using the resources assigned to it.
<b>Status</b>	Specifies the status of this alarm entry.

---

## Creating a port history alarm

### About this task

Create a port history alarm to track the number of alarms fired from a particular port.

### Procedure

1. Ensure that you globally enable RMON.  
Enabling RMON globally turns on logging and trapping.
  2. Select the port that has an alarm configured.
  3. Right-click the port.
  4. Choose **Enable Rmon Stats and Enable Rmon History**.
- 

---

## Viewing RMON alarms

### About this task

View the RMON alarm information to see alarm activity.



### Procedure

1. In the navigation tree, open the following folders: **Configuration > Serviceability > RMON**.
  2. Click **Alarms**.
- 

---

## Deleting an alarm

### About this task

Delete an alarm if you no longer want it to appear in the log.

### Procedure

1. In the navigation tree, open the following folders: **Configuration > Serviceability > RMON**.
  2. Click **Alarms**.
  3. Select the alarm you must delete.
  4. Click **Delete**.
- 

---

## Creating a default RMON event

### About this task

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 tree, open 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
<b>Index</b>	Uniquely identifies an entry in the event table. Each such entry defines one event that is generated after the appropriate conditions occur.
<b>Description</b>	Specifies a comment describing this event entry.
<b>Type</b>	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.
<b>Community</b>	Specifies the SNMP community to where you can send SNMP traps.
<b>LastTimeSent</b>	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.
<b>Owner</b>	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.

## Creating a nondefault RMON event

### About this task

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 tree, open 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
<b>Index</b>	Uniquely identifies an entry in the event table. Each such entry defines one event that is generated after the appropriate conditions occur.
<b>Description</b>	Specifies a comment describing this event entry.
<b>Type</b>	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.
<b>Community</b>	Specifies the SNMP community to where you can send SNMP traps.
<b>LastTimeSent</b>	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.
<b>Owner</b>	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

Name	Description
	management stations retrieve those log entries, as they have significance to all management stations connected to this device.

## Viewing RMON events

### About this task

View RMON events to see how many events occurred.

### Procedure

1. In the navigation tree, open 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
<b>Index</b>	Uniquely identifies an entry in the event table. Each such entry defines one event that is generated after the appropriate conditions occur.
<b>Description</b>	Specifies a comment describing this event entry.
<b>Type</b>	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.
<b>Community</b>	Specifies the SNMP community to where you can send SNMP traps.
<b>LastTimeSent</b>	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.
<b>Owner</b>	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

Name	Description
	management stations retrieve those log entries, as they have significance to all management stations connected to this device.

---

## Deleting an event

### About this task

Delete an event after you no longer require the alarm information.

### Procedure

1. In the navigation tree, open 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 8: IPFIX configuration using ACLI

Configure IPFIX to monitor IP flows.

---

## Enabling IPFIX globally

### Before you begin

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

### About this task

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

```
Enable IPFIX:  
ip ipfix enable
```

---

### Example

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

---

## Configuring IPFIX metering using filters

### Before you begin

- You must log on to the Global Configuration mode in ACLI.
- An ACL exists with match criteria.
- You must enable IPFIX globally before you can use it in a filter.

## About this task

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`.

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.

For more information about how to configure filters, see *Avaya Virtual Services Platform 9000 Configuration — QoS and ACL-Based Traffic Filtering* (NN46250–502).

## Procedure

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

```
filter acl set <1-2048> global-action ipfix-enable
```
2. 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.

**Table 5: Variable definitions**

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



Variable	Value
	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

### Before you begin

- You must log on to the GigabitEthernet Interface Configuration mode in ACLI.
- You must enable IPFIX globally before you can use it on a port.

### About this task

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

### Procedure

1. Configure the sampling rate:

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

2. 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.

**Table 6: Variable definitions**

Variable	Value
sampling-rate <1-100000>	Configures the sampling rate for metering on the port, as one in every $n$ packets. The default value is 1, which configures continuous monitoring.

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).

---

## Configuring IPFIX slot parameters

### Before you begin

- You must log on to the Global Configuration mode in ACLI.
- You must enable IPFIX globally before the IPFIX slot configuration will take affect.

### About this task

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

### Procedure

1. Configure the timeout value for flows:  

```
ip ipfix slot {slot[-slot][,...]} active-timeout <1-60>
```
2. Configure the aging interval for flow records:  

```
ip ipfix slot {slot[-slot][,...]} aging-interval <10-3600>
```
3. Configure the frequency to export to the collector:  

```
ip ipfix slot {slot[-slot][,...]} export-interval <10-3600>
```
4. Enable the slot to export flow information:  

```
ip ipfix slot {slot[-slot][,...]} exporter-enable
```
5. 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>
```

---

### 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.

**Table 7: Variable definitions**

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.
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

### Before you begin

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

### About this task

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.

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. 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}]
```
2. Specify the protocol:  

```
ip ipfix collector {slot[-slot][,...]} {A.B.C.D} protocol udp
```
3. 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.

**Table 8: Variable definitions**

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

### About this task

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.

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.

## Procedure

View flow information:

```
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/Vlan   SrcIP/DstIP   Src/Dst   Protocol/   DSCP/   Egress   Start/Last
  Addr      Port         Obsv     Point      TcpFlag Port/Mgid  Time
-----
4/7         16.16.16.1    0         icmp        0        0        Oct 26 14:05:04
7           15.15.15.1    0         Port        none     0        Oct 26 14:06:17

4/5         15.15.15.1    0         udp         0        0        Oct 26 14:05:04
5           16.16.16.1    0         Port        none     0        Oct 26 14:06:17

Total number of Displayed Flows on Slot 4 : 2
```

```
-----
Port/Vlan   SrcMac/DstMac   Byte/Pkt
  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.

**Table 9: Variable definitions**

Variable	Value
byte-count <i>WORD</i> <1-2> <0-4294967295>	Shows the flows that match a number of bytes. Use the format <code>oper{= != &lt;= &gt;=}</code> and <code>byte-count {0-4294967295}</code> ; for example, <code>{&gt;=a}</code> .
dest-addr <i>WORD</i> <1-2> {A.B.C.D}	Shows the flows for a destination address. Use the format <code>oper{= != &lt;= &gt;=}</code> and ip address {A.B.C.D}; for example, <code>{&lt;=A.B.C.D}</code> .
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 <code>oper{= != &lt;= &gt;=}</code> and time {MMddyyyyhhmmss}; for example, <code>{&gt;=a}</code> .
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 <code>oper{= != &lt;= &gt;=}</code> and time {MMddyyyyhhmmss}; for example, <code>{&gt;=a}</code> .
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.
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.
pkt-count <i>WORD</i> <1-2> <0-4294967295>	Shows the flows that match a packet count. Use the format <code>oper{= != &lt;= &gt;=}</code> and <code>pkt-count {0- 4294967295}</code> ; for example, <code>{&gt;=a}</code> .
port <i>WORD</i> <1-2> {slot/port}	Shows the flows for a particular port. Use the format <code>oper{= != &lt;= &gt;=}</code> and {slot/port}; for example, <code>{=a/b}</code> .

Variable	Value
protocol <i>WORD</i> <1-2> <0-255>	Shows the flows for a particular protocol. Use the format <code>oper{= &lt; = &gt;=}</code> and protocol {0-255}; for example, <code>{&gt;=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 <i>WORD</i> <1-2> {A.B.C.D}	Shows the flows for a source address. Use the format <code>oper{= &lt; = &gt;=}</code> and ip address {A.B.C.D}; for example, <code>{&lt;=A.B.C.D}</code> .
TCP-UDP-dest-port <i>WORD</i> <1-2> <0-65535>	Shows the flows for a destination port. Use the format <code>oper{= &lt; = &gt;=}</code> and port {0-65535}; for example, <code>{&gt;=a}</code> .
TCP-UDP-src-port <i>WORD</i> <1-2> <0-65535>	Shows the flows for a source port. Use the format <code>oper{= &lt; = &gt;=}</code> and port {0-65535}; for example, <code>{&gt;=a}</code> .
TOS <i>WORD</i> <1-2> <0-255>	Shows the flows that match a type of service. Use the format <code>oper{= &lt; = &gt;=}</code> and TOS{0-255}; for example, <code>{&gt;=a}</code> .
vlan <i>WORD</i> <1-2> <1-4084>	Shows the flows for a particular VLAN. Use the format <code>oper{= &lt; = &gt;=}</code> and vlan{1-4084}; for example, <code>{!=10}</code> .

---

## Flushing IPFIX flow information

### Before you begin

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

### About this task

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

### Procedure

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.

**Table 10: Variable definitions**

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

### About this task

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

### Procedure

View the global information:

```
show ip ipfix
```

### Example

```
VSP-9012:1#show ip ipfix
=====
IPFIX Global
=====
Global-State : enable
```



---

## Viewing collector information

### About this task

View collector information to verify the collector configuration.

### Procedure

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

### About this task

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**

View the exporter configuration:

```
show ip ipfix exporter [{slot[-slot][, ...]]
```

**Example**

```
VSP-9012:1#show ip ipfix exporter
=====
                        IPFIX Exporter-Info
=====
SlotNum   Admin   Template   Template   Export   Aging   Active
          State  Refresh-Rate Refresh-Rate Period   Period  Timeout
          (in sec)  (# of pkts) (in sec)  (in sec) (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**

**About this task**

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

**Procedure**

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).



# Chapter 9: IPFIX configuration using EDM

Configure IPFIX to monitor IP flows.

---

## Enabling IPFIX globally

### About this task

You must globally enable IPFIX before the software will perform IPFIX functions in filters or on a port. 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 tree, open the following folders: **Configuration > Serviceability**.
  2. Click **IPFIX**.
  3. Click the **Global** tab.
  4. Select **enable**.
  5. Click **Apply**.
- 

---

## Configuring collector parameters

### About this task

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.

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

---

## Configuring slot parameters

### About this task

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

### Procedure

1. In the navigation tree, open 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
<b>SlotNum</b>	Identifies the slot number.
<b>AgingIntv</b>	Configures the flow record aging interval. The default is 15 seconds.
<b>ActiveTimeout</b>	Configures the flow active timeout. The default is 30 minutes.
<b>ExportIntv</b>	Configure the interval at which to export flow information. The default is 50 seconds.
<b>ExportState</b>	Enables the exporter state for the slot. The default is enable.
<b>TempRefIntvSec</b>	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.
<b>TempRefIntvPkts</b>	Configures the template refresh timeout. The template is sent to the collector every x

Name	Description
	seconds or after x exported packets, whichever occurs first. The default is 20 packets.

---

## Configuring IPFIX on a port

### About this task

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

### Procedure

1. In the navigation tree, open 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
<b>Id</b>	Identifies the slot and port.
<b>SampleRate</b>	Configures the sampling rate for metering on the port, as one in every n packets. The default value is 1, which configures continuous monitoring.
<b>Flush</b>	Deletes all records stored in the COP or exports all records and deletes the database after the export finishes.
<b>AllTraffic</b>	Enables or disables IPFIX on all traffic for the specified port. The default is disabled.



---

## Configuring IPFIX metering using filters

### 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

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.

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 tree, open 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 10: Port performance management using ACLI

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

---

## Configuring an automatic trace

### About this task

Configure the Avaya Virtual Services Platform 9000 to automatically enable a trace if CP module utilization reaches a predefined value.

### Procedure

1. Configure automatic trace for a specific module:  
`trace auto module add <0-107> <0-4>`
2. Enable the trace auto-enable feature:  
`trace auto enable`

---

### Example

```
VSP-9012:1>trace auto module add 6 2
```

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

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

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

```

20 - TDP
21 - MAN_DIAG
22 - MAN_TEST

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

## Variable definitions

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

**Table 11: Variable definitions**

Variable	Value
disable	Disables the automatic trace.
enable	Enables the automatic trace.
high-percentage <60–100>	Specify the CP module utilization percentage above which the system enables auto-trace. The default is 90.
high-track-duration <3–10>	Specify the time in seconds to monitor CP module utilization before the system enables auto-trace. The default is 5.
low-percentage <50–90>	Specify the CP module utilization percentage below which the system disables auto-trace. The default is 75.
low-track-duration <3–10>	Specify the time, in seconds, to monitor CP module utilization before the system disables the trace. The default is 5.
module add <0-107> <0-4>	<p>Add a module you want the system to trace.</p> <ul style="list-style-type: none"> <li>• &lt;0-107&gt; identifies the module that you want to add.</li> <li>• &lt;0-4&gt; specifies the level of the module.</li> </ul> <p>Use the <code>show trace modid-list</code> command to see a list of modules with their ID.</p>
module remove <0-107>	<p>Stop the trace for a specific module. &lt;0-107&gt; identifies the module that you want to remove.</p> <p>Use the <code>show trace modid-list</code> command to see a list of modules with their ID.</p>

## Viewing DDI module information

### Before you begin

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

### About this task

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

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

### Procedure

1. View basic SFP and SFP+ manufacturing information and characteristics:  

```
show pluggable-optical-modules basic [{slot/port[-slot/port]
[,...]}]
```
2. View configuration information:  

```
show pluggable-optical-modules config
```
3. View detailed SFP and SFP+ manufacturing information and characteristics:  

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

### Example

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

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

PORT NUM	TYPE	DDM ENABLED	VENDOR NAME	PART NUMBER
3/1	10GbSR	TRUE	Avaya	AA1403015-E6
3/8	10GbSR	TRUE	Avaya	AA1403015-E6
3/9	10GbSR	TRUE	Avaya	AA1403015-E6
3/10	10GbSR	TRUE	Avaya	AA1403015-E6
4/1	GbicSx	TRUE	Avaya	AA1419048-E6
4/25	GbicSx	TRUE	Avaya	AA1419048-E6
4/37	GbicSx	TRUE	Avaya	AA1419048-E6
4/48	GbicSx	TRUE	Avaya	AA1419048-E6
6/23	Gbic1000BaseT	FALSE	Avaya	AA1419043-E6
6/24	Gbic1000BaseT	FALSE	Avaya	AA1419043-E6
6/25	Gbic1000BaseT	FALSE	Avaya	AA1419043-E6
6/26	Gbic1000BaseT	FALSE	Avaya	AA1419043-E6
6/43	GbicSx	TRUE	Avaya	AA1419048-E6

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

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

```

=====
                ddm-monitor : disabled
    ddm-monitor-interval : 5
                ddm-traps-send : enabled
    ddm-alarm-portdown : disabled
=====
    
```

## Variable definitions

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

**Table 12: Variable definitions**

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 temperature information

### Before you begin

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

### About this task

Perform this procedure to view SFP and SFP+ temperatures.

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

### Procedure

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          LOW_ALARM LOW_WARN   ACTUAL   HIGH_WARN HIGH_ALARM THRESHOLD
NUM           THRESHOLD THRESHOLD  VALUE   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.

**Table 13: Variable definitions**

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

### Before you begin

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

### About this task

Perform this procedure to view SFP and SFP+ voltages.

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

### Procedure

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.

**Table 14: Variable definitions**

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



# Chapter 11: Port performance management using EDM

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

---

## Configuring rate limits

### About this task

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 tree, open 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
<b>Index</b>	The port number.
<b>TrafficType</b>	The type of traffic being rate limited, either broadcast or multicast traffic.
<b>AllowedRatekbps</b>	This variable is the allowed traffic rate limit for the port.

Name	Description
	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. On gigabit ports and MDAs, there can be up to a 2 percent difference between the configured and actual rate limiting values. For the Avaya Virtual Services Platform 9000, 1–65535 sets the limit in packets for each second.
<b>Enable</b>	Double-click in the field and select to enable (True) or disable (False) rate limiting.

---

## Enabling learning limits on a port

### About this task

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 tree, open 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
<b>PortNum</b>	Shows the slot and port number to configure.
<b>MaxMacCount</b>	Configures the number of entries in the MAC table for the port that causes learning to stop. The default is 1024.

Name	Description
<b>MinMacCount</b>	Configures the number of entries in the MAC table for the port at which learning can resume. The default is 512.
<b>CurrentMacCount</b>	Shows the number of entries currently in the MAC table for the port.
<b>Enable</b>	Enables or disables limit learning for the port.
<b>MacLearning</b>	Shows if MAC learning is enabled or disabled for the port.
<b>ViolationLogTrap</b>	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.
<b>ViolationDownPort</b>	Configures the system to disable the port after a MAC address violation is detected. The default is disable.

---

## Viewing DDI information

### About this task

You can view DDI information (such as module information, temperature, and voltages) for SFPs and SFP+s on the 1 Gb and 10 Gb interface modules.

### Procedure

1. In the Physical Device view, select a port.
2. In the navigation tree, open 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
<b>DdmStatus</b>	Indicates if DDM is enabled.

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

Name	Description
<b>VoltageHighWarningThreshold</b>	Indicates the high warning threshold in volts.
<b>VoltageLowWarningThreshold</b>	Indicates the high warning threshold in volts.
<b>VoltageStatus</b>	Indicates if any voltage thresholds were exceeded.
<b>Bias</b>	Indicates the laser bias current in mA.
<b>BiasHighAlarmThreshold</b>	Indicates the bias current high alarm threshold in mA.
<b>BiasLowAlarmThreshold</b>	Indicates the bias current low alarm threshold in mA.
<b>BiasHighWarningThreshold</b>	Indicates the bias current high warning threshold in mA.
<b>BiasLowWarningThreshold</b>	Indicates the bias current high warning threshold in mA.
<b>BiasStatus</b>	Indicates if any bias thresholds were exceeded.
<b>TxPower</b>	Indicates the current Tx power in mW.
<b>TxPowerHighAlarmThreshold</b>	Indicates the high alarm threshold in mW for the Tx power.
<b>TxPowerLowAlarmThreshold</b>	Indicates the low alarm threshold in mW for the Tx power.
<b>TxPowerHighWarningThreshold</b>	Indicates the high warning threshold in mW for the Tx power.
<b>TxPowerLowWarningThreshold</b>	Indicates the high warning threshold in mW for the Tx power.
<b>TxPowerStatus</b>	Indicates if any Tx power thresholds were exceeded.
<b>RxPower</b>	Indicates the current Rx power in mW.
<b>RxPowerHighAlarmThreshold</b>	Indicates the high alarm threshold in mW for the Rx power.
<b>RxPowerLowAlarmThreshold</b>	Indicates the low alarm threshold in mW for the Rx power.
<b>RxPowerHighWarningThreshold</b>	Indicates the high warning threshold in mW for the Rx power.
<b>RxPowerLowWarningThreshold</b>	Indicates the high warning threshold in mW for the Rx power.
<b>RxPowerStatus</b>	Indicates if any Rx power thresholds were exceeded.
<b>Aux1</b>	Indicates the current auxiliary 1 reading.

Name	Description
<b>Aux1HighAlarmThreshold</b>	Indicates the high alarm threshold auxiliary 1 reading.
<b>Aux1LowAlarmThreshold</b>	Indicates the low alarm threshold auxiliary 1 reading.
<b>Aux1HighWarningThreshold</b>	Indicates the high warning threshold auxiliary 1 reading.
<b>Aux1LowWarningThreshold</b>	Indicates the high warning threshold auxiliary 1 reading.
<b>Aux1Status</b>	Indicates if any auxiliary 1 thresholds were exceeded.
<b>Aux2</b>	Indicates the current auxiliary 2 reading.
<b>Aux2rHighAlarmThreshold</b>	Indicates the high alarm threshold auxiliary 2 reading.
<b>Aux2LowAlarmThreshold</b>	Indicates the low alarm threshold auxiliary 2 reading.
<b>Aux2HighWarningThreshold</b>	Indicates the high warning threshold auxiliary 2 reading.
<b>Aux2LowWarningThreshold</b>	Indicates the high warning threshold auxiliary 2 reading.
<b>Aux2rStatus</b>	Indicates if any auxiliary 2 thresholds were exceeded.

# Chapter 12: 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

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 15: 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.

Field	Description
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 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

### About this task

View port routing statistics to manage network performance.

### Procedure

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      IN_FRAME  IN_FRAME  IN      OUT_FRAME  OUT_FRAME
NUM       UNICAST   MULTICAST DISCARD  UNICAST    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 16: 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

### Before you begin

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

### About this task

Display individual bridging statistics for specific ports to manage network performance.

### Procedure

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   IN_FRAME   IN_FRAME   IN_FRAME   OUT_FRAME  IN_FRAME   OUT_FRAME  IN_DISCARD
NUM    UNICAST    MULTICAST  BROADCAST  IN_FRAME  xSTP BPDU  xSTP BPDU
-----
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
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 17: 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

### Before you begin

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

### About this task

Display individual DHCP-relay statistics for specific ports to manage network performance.

### Procedure

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 <i>WORD</i> <0-16>	Specify a VRF instance by VRF name, where <i>WORD</i> < 0-16> is an integer between 0 and 16.
vrfids <i>WORD</i> <0-255>	Specify the ID of the VRF, where <i>WORD</i> < 0-255> is an integer between 0 and 255..
{slot/port[-slot/port][,...]}	Display all DHCP statistics for a port or a range of ports.

---

## Job aid

The following table describes parameters for the `show interfaces GigabitEthernet statistics dhcp-relay` command output.

**Table 18: 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

### About this task

Display DHCP-relay statistics for all interfaces to manage network performance.

### Procedure

1. Show the number of requests and replied for each interface:

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

2. Show counters for Option 82:

```
show ip dhcp-relay counters option82 [vrf WORD<0-16>] [vrfids
WORD<0-512>]
```

**Example**

```
VSP-9012:1>show ip dhcp-relay counters option82
=====
DHCP Counters Option82 - GlobalRouter
=====
INTERFACE   FOUND DROP  CIRCUIT ADD   REMOVE REMOVE   ADD   REMOVE
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&lt;0-16&gt;</i>	Specify a VRF instance by VRF name, where <i>WORD&lt;0-16&gt;</i> is an integer between 0–16.
vrfids <i>WORD&lt;0-512&gt;</i>	Specify the ID of the VRF, where <i>WORD&lt;0-512&gt;</i> is an integer between 0–255.

**Job aid**

The following table explains the output from the **show ip dhcp-relay counters option82** command.

**Table 19: 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.

Heading	Description
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. To determine the cause of the drop, you must enable trace on level 170. DHCP option 82dropped packets trace
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.
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

View statistics:

```
show ipv6 dhcp-relay counters
```



#### Note:

Use the `sys action reset counters` command to clear DHCP Relay statistics.

### Example

```
VSP-9012:1(config-if)#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 20: 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

### Before you begin

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

### About this task

Display individual LACP statistics for specific ports to manage network performance.

### Procedure

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 21: 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.



Parameter	Description
TX MARKERPDU	The count of transmitted marker protocol data units.
RX MARKERPDU	The count of received marker protocol data units.
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

### Before you begin

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

### About this task

Display individual RMON statistics for specific ports to manage network performance.

### Procedure

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 22: 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.
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

Parameter	Description
	(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

### Before you begin

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

### About this task

Display detailed statistics for specific ports to manage network performance.

### Procedure

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 23: 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 policing statistics

### About this task

View statistics to ensure proper QoS performance.

### Procedure

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 24: 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

### Before you begin

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

### About this task

Clear default ACL statistics if you no longer require previous statistics.

### Procedure

1. Enter the following command to clear default ACL statistics:  
`clear filter acl statistics default [<1-2048>]`
2. Enter the following command to clear global ACL statistics:  
`clear filter acl statistics global [<1-2048>]`
3. Enter the following command to clear all ACL statistics:  
`clear filter acl statistics all`
4. 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
<i>1-2048</i>	Specifies the ACL ID.
<i>1-2000</i>	Specifies the ACE ID.

---

## Viewing ACE statistics

### Before you begin

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

### About this task

View ACE statistics to ensure that the filter operates correctly.

### Procedure

1. View ACE statistics for a specific ACL, ACE, or ACE type:  

```
show filter acl statistics <1-2048> [<1-2000>] [qos]
[security]
```
  2. View all ACE statistics:  

```
show filter acl statistics all
```
  3. View default ACE statistics:  

```
show filter acl statistics default [<1-2048>]
```
  4. 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
<i>1-2048</i>	Specifies the ACL ID.
<i>1-2000</i>	Specifies the ACE ID.

---

## Job aid

The following table describes output for the `show filter acl statistics port` command.

**Table 25: 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

### About this task

Display MSTP statistics to see MSTP related bridge-level statistics.

### Procedure

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 26: 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

### About this task

View Rapid Spanning Tree Protocol statistics to manage network performance.

### Procedure

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 27: 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

### About this task

View RSTP statistics on ports to manage network performance.

### Procedure

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),

Variable	Value
	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 28: 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

### About this task

View MLT statistics to display MultiLinkTrunking statistics for the switch or for the specified MLT ID.

### Procedure

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 29: 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.
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

### Before you begin

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

### About this task

Use statistics to help you monitor Open Shortest Path First (OSPF) performance. You can also use statistics in troubleshooting procedures.

### Procedure

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 30: 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.
AUTHYPEMISMATCH	Indicates the number of AuthType mismatches this interface receives.
AUTH FAILURES	Indicates the number of authentication failures.
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

### About this task

Use statistics to help you monitor Open Shortest Path First (OSPF) performance. You can also use statistics in troubleshooting procedures.

### Procedure

Display OSPF interface statistics:

```
show ip ospf ifstats [detail vrf WORD<0-16> vrfids WORD<0-512>]
[mismatch vrf WORD<0-16> vrfids WORD<0-512>] [vrf WORD<0-16>]
[vrfids WORD<0-512>]
```

### Example

```
VSP-9012:1#show ip ospf ifstats
```

```
=====
                        OSPF Interface Statistics - GlobalRouter
=====
---HELLOS--- ---DBS--- -LS REQ-- --LS UPD--- --LS ACK---
INTERFACE      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.
vrf WORD<0-16>	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 31: 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

### About this task

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

Display the OSPF range statistics:

```
show ip ospf stats [vrf WORD<0-16>] [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 WORD<0-16>	Specifies a VRF instance by VRF name.
vrfids WORD<0-16>	Specifies a VRF or range of VRFs by ID.

---

## Job aid

The following table describes the show command output.

**Table 32: 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.
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

### About this task

Use statistics to help you monitor Open Shortest Path First (OSPF) performance. You can also use statistics in troubleshooting procedures.

### Procedure

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 33: show ports statistics ospf main output description**

Field	Description
PORT NUM	Indicates the port number.

Field	Description
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

### About this task

Use statistics to help you monitor Open Shortest Path First (OSPF) performance. You can also use statistics in troubleshooting procedures.

### Procedure

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 34: 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

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.

Field	Description
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

### About this task

Display the current device configuration.

### Procedure

Display the current device configuration by using the following command:

```
show eapol system
```

### Example

```
VSP-9012:1#show eapol system
                        eap : enabled
                        sess-manage : false
```

---

## Showing EAPoL authenticator statistics

### About this task

Display the authenticator statistics to manage network performance.

## Procedure

Display the authenticator statistics:

```
show eapol auth-stats interface [gigabitEthernet [{slot/port[-
slot/port][,...]}]] [vlan <1-4084>]
```

## Example

```
MSP-9012:1#show eapol auth-stats interface
=====
                        Eap Authenticator Statistics
=====
PORT  TOTAL  TOTAL  START  LOGOFF  RESP_ID  RESP  REQ-ID  REQ  INVALID  LENGTH  FRAME  LAST-SRC
  RX   TX    RCVD   RCVD   RCVD    RCVD  TX    TX    FRAMES  ERROR  VER    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

### About this task

View EAPoL session statistics to manage network performance.

## Procedure

Display the session statistics:

```
show eapol session-stats interface [gigabitEthernet [{slot/
port[-slot/port][, ...]}] [vlan <1-4084>]
```

## Example

```
VSP-9012:1#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-logout
4/2	0	0	0	0		local-server	0 day(s),	00:00:00	supp-logout
4/3	0	0	0	0		local-server	0 day(s),	00:00:00	supp-logout
4/4	0	0	0	0		local-server	0 day(s),	00:00:00	supp-logout
4/5	0	0	0	0		local-server	0 day(s),	00:00:00	supp-logout
4/6	0	0	0	0		local-server	0 day(s),	00:00:00	supp-logout
4/7	0	0	0	0		local-server	0 day(s),	00:00:00	supp-logout
4/8	0	0	0	0		local-server	0 day(s),	00:00:00	supp-logout
4/9	0	0	0	0		local-server	0 day(s),	00:00:00	supp-logout
4/10	0	0	0	0		local-server	0 day(s),	00:00:00	supp-logout
4/11	0	0	0	0		local-server	0 day(s),	00:00:00	supp-logout
4/12	0	0	0	0		local-server	0 day(s),	00:00:00	supp-logout
4/13	0	0	0	0		local-server	0 day(s),	00:00:00	supp-logout
4/14	0	0	0	0		local-server	0 day(s),	00:00:00	supp-logout
4/15	0	0	0	0		local-server	0 day(s),	00:00:00	supp-logout

## 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.
<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 35: 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

### Before you begin

- To clear statistics, you must log on to at least the Privileged EXEC mode in the ACLI.

### About this task

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. Display RADIUS server statistics:  
`show radius-server statistics`
2. Clear server statistics:  
`clear radius statistics`

---

### Example

```
VSP-9012:1#show radius-server statistics
```

## Viewing statistics using ACLI

```
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 36: 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

### About this task

View RMON statistics to manage network performance.

### Procedure

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 37: 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 assigned resources.

## Viewing PCAP statistics

### About this task

View PCAP statistics to manage network performance.

### Procedure

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 38: 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.   <b>Important:</b> 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.
Number of packets accumulated in PCAP engine	This is the number of packets accumulated in the PCAP engine.   <b>Important:</b> 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.
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

### About this task

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.

If you do not specify a slot, all slots appear in the command output.

### Procedure

1. View exporter statistics:  

```
show ip ipfix export [{slot[-slot]][,...]]
```
2. View hashing statistics:  

```
show ip ipfix hash-statistics [{slot[-slot]][,...]]
```

### 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

### Before you begin

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

### About this task

Clear IPFIX statistics to remove the exporter and hash statistics.

### Procedure

1. Clear exporter statistics:  

```
clear ip ipfix stats [{slot[-slot][,...]}]
```
2. Clear hash statistics:  

```
clear ip ipfix hash-stats [{slot[-slot][,...]}]
```

---

### 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

Variable	Value
	specify a slot, you clear the statistics for all slots.

---

## Clearing IPv6 statistics

### Before you begin

- You must log on to the User Exec Mode in ACLI

### About this task

Clear all IPv6 statistics if you do not require previous statistics.

### Procedure

1. Enter the following command to clear all the IPv6 statistics:  

```
clear ipv6 statistics all
```
2. Enter the following command to clear interface statistics:  

```
clear ipv6 statistics interface [general|icmp]
[gigabitethernet <slot/port>|mgmtethernet <slot/port>|vlan
<1-4084>]
```
3. Enter the following command to clear TCP statistics:  

```
clear ipv6 statistics tcp
```
4. 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
<code>&lt;1-4084&gt;</code>	Specifies the VLAN id for clearing IPv6 ICMP VLAN interface statistics
<code>slot/port</code>	Specifies the slot and the port number to clear brouter interface statistics.
<code>slot/port</code>	Specifies the slot and the port number to clear statistics on a management port.



---

## Viewing multicast routing process statistics

### 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.

### About this task

Enable the collection and display of multicast routing process statistics. These statistics are not related to the interface (port) statistics. Rather, the statistics are displayed based on multicast group classification. By default, mroute statistics collection is disabled.

### Procedure

1. Enable statistics collection:  
`ip mroute stats enable`
2. View statistics:  
`show ip mroute stats [WORD<7-160>]`
3. Display statistics at regular intervals:  
`monitor ip mroute stats WORD<7-160>`

You can change the duration or interval for monitoring in the Global Configuration mode.

4. Clear statistics:  
`clear ip mroute stats`

---

### Example

```
VSP-9012:1(config)#show ip mroute stats 239.1.1.1
=====
                          Multicast Stats - GlobalRouter
=====
GroupAddress      SourceCounter      IngressPackets      IngressBytes
  AverageSize      Packets/Second      DropPackets          DropBytes
-----
239.1.1.1         1                   1090179126          89394689233
   82              1225653              0                    0
```

---

## Variable definitions

Use the data in the following table to use the `monitor ip mroute stats` and `show ip mroute stats` commands.

Variable	Value
<i>WORD</i> <7-160>	Specifies the group IP address in the format {A.B.C.D[,E.F.G.H][,...]} . The maximum number of group IP addresses is 10. To view statistics, the group IP address is optional. To monitor statistics, the group IP address is required.

---

## Job aid

The following table explains the output of the `show ip mroute stats` command.

**Table 39: show ip mroute stats command output**

Heading	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 group IP address.
DropBytes	Specifies the number of dropped bytes for the associated group IP address.

---

## Viewing IPv6 VRRP statistics

View IPv6 VRRP statistics to monitor network performance

### Procedure

View statistics for the device and for all interfaces:

```
show ipv6 vrrp statistics [link-local WORD<0-127>]] [vrid <1-255>]
```

### Example

```
VSP-9012:1(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</i> <0-127>	Shows statistics for a specific link-local address.

Variable	Value
vrid <1-255>	Shows statistics for a specific VRID.

---

## Job aid

The following table describes the output for the `show ipv6 vrrp statistics` command.

**Table 40: 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.
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.

Heading	Description
	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 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

### Before you begin

- You must log on to the User Exec Mode in ACLI.

### About this task

View IPV6 ICMP statistics on an interface for ICMP messages sent over a particular interface.

### Procedure

```
View IPv6 ICMP statistics
show ipv6 interface icmpstatistics
```

---



---

## 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.

### Example

View ICMP statistics:

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

---

## Viewing IPv6 statistics on an interface

View IPv6 statistics to view information about the IPv6 datagrams on an interface.

### Procedure

View statistics:

```
show ipv6 interface statistics [<1-4084>]
```

---

### Example

---

## 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.





# Chapter 13: Viewing statistics using EDM

Use statistics 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 217. 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**.
- 

---

## Graphing chassis statistics

Create graphs of chassis statistics to generate a visual representation of your data.

### Procedure

1. In the Device Physical View, select the chassis.
2. In the navigation tree, open 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. In 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, open 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. In the Device Physical View, select the chassis.
  2. In the navigation tree, open 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
<b>DramUsed</b>	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.
<b>DramFree</b>	The amount in kilobytes of free DRAM.
<b>CpuUtil</b>	Percentage of CPU utilization.

---

## Viewing chassis SNMP statistics

View chassis SNMP statistics to monitor network performance.

### Procedure

1. In the Device Physical View, select the chassis.
2. In the navigation tree, open 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
<b>InPkts</b>	The number of messages delivered to the SNMP entity from the transport service.
<b>OutPkts</b>	The number of SNMP messages passed from the SNMP protocol entity to the transport service.
<b>InTotalReqVars</b>	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.
<b>InTotalSetVars</b>	The number of MIB objects altered successfully by the SNMP protocol entity as the result of receiving valid SNMP Set-Request PDUs.
<b>InGetRequests</b>	The number of SNMP Get-Request PDUs the SNMP protocol accepts and processes.
<b>InGetNexts</b>	The number of SNMP Get-Next PDUs the SNMP protocol accepts and processes.
<b>InSetRequests</b>	The number of SNMP Set-Request PDUs the SNMP protocol accepts and processes.
<b>InGetResponses</b>	The number of SNMP Get-Response PDUs the SNMP protocol accepts and processes.
<b>OutTraps</b>	The number of SNMP Trap PDUs the SNMP protocol generates.
<b>OutTooBig</b>	The number of SNMP PDUs the SNMP protocol generates for which the value of the error-status field is tooBig.
<b>OutNoSuchNames</b>	The number of SNMP PDUs the SNMP protocol generates for which the value of the error-status field is noSuchName.
<b>OutBadValues</b>	The number of SNMP PDUs the SNMP protocol generates for which the value of the error-status field is badValue.
<b>OutGenErrs</b>	The number of SNMP PDUs the SNMP protocol generates for which the value of the error-status field is genErr.
<b>InBadVersions</b>	The number of SNMP messages delivered to the SNMP protocol entity for an unsupported SNMP version.

Name	Description
<b>InBadCommunityNames</b>	The number of SNMP messages delivered to the SNMP protocol entity that used an SNMP community name not known to said entity.
<b>InBadCommunityUses</b>	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.
<b>InASNParseErrs</b>	The number of ASN.1 or BER errors the SNMP protocol encountered when decoding received SNMP messages.
<b>InTooBig</b>	The number of SNMP PDUs delivered to the SNMP protocol entity for which the value of the error-status field is tooBig.
<b>InNoSuchNames</b>	The number of SNMP PDUs delivered to the SNMP protocol entity for which the value of the error-status field is noSuchName.
<b>InBadValues</b>	The number of SNMP PDUs delivered to the SNMP protocol entity for which the value of the error-status field is badValue.
<b>InReadOnly</b>	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 field. This object is provided to detect incorrect implementations of the SNMP.
<b>InGenErrs</b>	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. In the Device Physical View, select the chassis.
  2. In the navigation tree, open 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
<b>InReceives</b>	The number of input datagrams received from interfaces, including those received in error.
<b>InHdrErrors</b>	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.
<b>InAddrErrors</b>	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.
<b>ForwDatagrams</b>	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.
<b>InUnknownProtos</b>	The number of locally addressed datagrams received successfully but discarded because of an unknown or unsupported protocol.
<b>InDiscards</b>	The number of input IP datagrams for which no problems were encountered to prevent their continued processing but that were discarded (for example, for lack of buffer space). This counter does not include any datagrams discarded while awaiting reassembly.
<b>InDelivers</b>	The number of input datagrams successfully delivered to IP user-protocols (including ICMP).
<b>OutRequests</b>	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.
<b>OutDiscards</b>	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).

Name	Description
	This counter includes datagrams counted in ipForwDatagrams if any such packets met this (discretionary) discard criterion.
<b>OutNoRoutes</b>	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.
<b>FragOKs</b>	The number of IP datagrams that were successfully fragmented at this entity.
<b>FragFails</b>	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.
<b>FragCreates</b>	The number of IP datagram fragments that were generated as a result of fragmentation at this entity.
<b>ReasmReqds</b>	The number of IP fragments received that needed to be reassembled at this entity.
<b>ReasmOKs</b>	The number of IP datagrams successfully reassembled.
<b>ReasmFails</b>	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. In the Device Physical View, select the chassis.
  2. In the navigation tree, open 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
<b>SrcQuenchs</b>	The number of ICMP Source Quench messages received.
<b>Redirects</b>	The number of ICMP Redirect messages received.
<b>Echos</b>	The number of ICMP Echo (request) messages received.
<b>EchoReps</b>	The number of ICMP Echo Reply messages received.
<b>Timestamps</b>	The number of ICMP Timestamp (request) messages received.
<b>TimestampReps</b>	The number of ICMP Timestamp Reply messages received.
<b>AddrMasks</b>	The number of ICMP Address Mask Request messages received.
<b>AddrMaskReps</b>	The number of ICMP Address Mask Reply messages received.
<b>ParmProbs</b>	The number of ICMP Parameter Problem messages received.
<b>DestUnreachs</b>	The number of ICMP Destination Unreachable messages received.
<b>TimeExcds</b>	The number of ICMP Time Exceeded messages received.

---

## Viewing chassis ICMP Out statistics

View chassis ICMP Out statistics to monitor network performance.

### Procedure

1. In the Device Physical View, select the chassis.
  2. In the navigation tree, open 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
<b>SrcQuenchs</b>	The number of ICMP Source Quench messages sent.
<b>Redirects</b>	The number of ICMP Redirect messages received. For a host, this object is always zero, because hosts do not send redirects.
<b>Echos</b>	The number of ICMP Echo (request) messages sent.
<b>EchoReps</b>	The number of ICMP Echo Reply messages sent.
<b>Timestamps</b>	The number of ICMP Timestamp (request) messages sent.
<b>TimestampReps</b>	The number of ICMP Timestamp Reply messages sent.
<b>AddrMasks</b>	The number of ICMP Address Mask Request messages sent.
<b>AddrMaskReps</b>	The number of ICMP Address Mask Reply messages sent.
<b>ParmProbs</b>	The number of ICMP Parameter Problem messages sent.
<b>DestUnreachs</b>	The number of ICMP Destination Unreachable messages sent.
<b>TimeExcds</b>	The number of ICMP Time Exceeded messages sent.

---

## Viewing ICMP statistics

### About this task

View ICMP statistics for ICMP configuration information.

### Procedure

1. In the navigation tree, open 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.  
The **Interfaces-ICMPStats** tab is displayed with the required ICMP statistics.
-

## ICMP stats field descriptions

Use the data in the following table to use the ICMP **Statistics** tab.

Name	Description
<b>InMsgs</b>	Specifies the total number of ICMP messages which the entity received. Note that this counter includes all those counted by icmpInErrors.
<b>InErrors</b>	Specifies the number of ICMP messages which the entity received but determined as having ICMP-specific errors (bad ICMP checksums, bad length, etc.).
<b>InDestUnreachs</b>	Specifies the number of ICMP Destination Unreachable messages received by the interface.
<b>InAdminProhibs</b>	Specifies the number of ICMP destination unreachable/communication administratively prohibited messages received by the interface.
<b>InTimeExcds</b>	Specifies the number of ICMP Time Exceeded messages by the interface.
<b>InParmProblems</b>	Specifies the number of ICMP Parameter Problem messages received by the interface.
<b>InPktTooBigs</b>	Specifies the number of ICMP Packet Too Big messages received by the interface.
<b>InEchos</b>	Specifies the number of ICMP Echo (request) messages received by the interface.
<b>InEchoReplies</b>	Specifies the number of ICMP Echo Reply messages received by the interface.
<b>InRouterSolicits</b>	Specifies the number of ICMP Router Solicit messages received by the interface.
<b>InRouterAdvertisements</b>	Specifies the number of ICMP Router Advertisement messages received by the interface.
<b>InNeighborSolicits</b>	Specifies the number of ICMP Neighbor Solicit messages received by the interface.

Name	Description
<b>InNeighborAdvertisements</b>	Specifies the number of ICMP Neighbor Advertisement messages received by the interface.
<b>InRedirects</b>	Specifies the number of ICMP Redirect messages received by the interface.
<b>InGroupMembQueries</b>	Specifies the number of ICMPv6 Group Membership Query messages received by the interface
<b>InGroupMembResponses</b>	Specifies the number of ICPv6 Group Membership Response messages received by the interface.
<b>InGroupMembReductions</b>	Specifies the number of ICMPv6 Group Membership Reduction messages received by the interface.
<b>OutMsgs</b>	Specifies the total number of ICMP messages which this interface attempted to send. Note that this counter includes all those counted by icmpOutErrors
<b>OutErrors</b>	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.
<b>OutDestUnreachs</b>	Specifies the number of ICMP Destination Unreachable messages sent by the interface.
<b>OutAdminProhibs</b>	Specifies the number of ICMP dest unreachable/communication administratively prohibited messages sent.
<b>OutTimeExcds</b>	Specifies the number of ICMP Time Exceeded messages sent by the interface
<b>OutParmProblems</b>	Specifies the number of ICMP Parameter Problem messages sent by the interface
<b>OutPktTooBigs</b>	Specifies the number of ICMP Packet Too Big messages sent by the interface.
<b>OutEchos</b>	Specifies the number of ICMP Echo (request) messages sent by the interface.

Name	Description
<b>OutEchoReplies</b>	Specifies the number of ICMP Echo Reply messages sent by the interface
<b>OutRouterSolicits</b>	Specifies the number of ICMP Router Solicitation messages sent by the interface.
<b>OutRouterAdvertisements</b>	Specifies the number of ICMP Router Advertisement messages sent by the interface
<b>OutNeighborSolicits</b>	Specifies the number of ICMP Neighbor Solicitation messages sent by the interface.
<b>OutNeighborAdvertisements</b>	Specifies the number of ICMP Neighbor Advertisement messages sent by the interface
<b>OutRedirects</b>	Specifies the number of Redirect messages sent. For a host, this object will always be zero, since hosts do not send redirects
<b>OutGroupMembQueries</b>	Specifies the number of ICMPv6 Group Membership Query messages sent
<b>OutGroupMembResponses</b>	Specifies the number of ICMPv6 Group Membership Response messages sent.
<b>OutGroupMembReductions</b>	Specifies the number of ICMPv6 Group Membership Reduction messages sent.

---

## Viewing chassis TCP statistics

View TCP statistics to monitor network performance.

### Procedure

1. In the Device Physical View, select the chassis.
  2. In the navigation tree, open 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
<b>ActiveOpens</b>	The number of times TCP connections made a direct transition to the SYN-SENT state from the CLOSED state.
<b>PassiveOpens</b>	The number of times TCP connections made a direct transition to the SYN-RCVD state from the LISTEN state.
<b>AttemptFails</b>	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.
<b>EstabResets</b>	The number of times TCP connections made a direct transition to the CLOSED state from either the ESTABLISHED state or the CLOSE-WAIT state.
<b>CurrEstab</b>	The number of TCP connections for which the current state is either ESTABLISHED or CLOSE-WAIT.
<b>InSegs</b>	The number of segments received, including those received in error. This count includes segments received on currently established connections.
<b>OutSegs</b>	The number of segments sent, including those on current connections, but excluding those containing only retransmitted octets.
<b>RetransSegs</b>	The number of segments retransmitted that is, the number of TCP segments transmitted containing one or more previously transmitted octets.
<b>InErrs</b>	The number of segments received in error (for example, bad TCP checksums).
<b>OutRsts</b>	The number of TCP segments sent containing the RST flag.
<b>HCInSegs</b>	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.
<b>HCOutSegs</b>	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. In the Device Physical View, select the chassis.
2. In the navigation tree, open 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
<b>NoPorts</b>	The number of received UDP datagrams with no application at the destination port. 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.
<b>InErrs</b>	The number of received UDP datagrams that were not delivered for reasons other than the lack of an application at the destination port.

Name	Description
	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.
<b>InDatagrams</b>	The number of UDP datagrams delivered to UDP users, for devices that can receive more than 1 000 000 UDP datagrams for each second. 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.
<b>OutDatagrams</b>	The number of UDP datagrams sent from this entity. 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.
<b>HCInDatagrams</b>	The number of TCP connections for which the current state is either ESTABLISHED or CLOSE-WAIT.
<b>HCOutDatagrams</b>	The number of UDP datagrams sent from this entity, for devices that can transmit more than 1 million UDP datagrams for each second. 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.

---

## Configuring Switch Fabric statistics capture

Configure the statistic settings for the Switch Fabric modules to determine what to capture.

### Procedure

1. In the navigation tree, open the following folders: **Configuration > Edit**.
  2. Click **Switch Fabric**.
  3. Click the **Switch Fabric Stats Settings** tab.
  4. Specify the Class of Service.
  5. Select the port or ports.
  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
<b>StatsCosId</b>	Specifies the Class of Service on which to collect statistics.
<b>StatsPortId</b>	Specifies the port on which to collect statistics. You must select a data port.
<b>StatsCapture</b>	Turns statistics collection on or off. The default is off.
<b>DeviceRead</b>	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 tree, open 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
<b>StatsPortId</b>	Shows the data port on which statistics are collected.
<b>DropPrec1AcceptedPackets</b>	Shows the accepted number of packets from drop precedence 1. The drop precedence is a function of the front-end policer.



Name	Description
<b>DropPrec1AcceptedBytes</b>	Shows the accepted number of bytes from drop precedence 1.
<b>DropPrec1CongestionMarkedPackets</b>	Shows the congestion marked number of packets from drop precedence 1.
<b>DropPrec1CongestionMarkedBytes</b>	Shows the congestion marked number of bytes from drop precedence 1.
<b>DropPrec1DiscardDroppedPackets</b>	Shows the number of WRED dropped packets from drop precedence 1.
<b>DropPrec1DiscardDroppedBytes</b>	Shows the number of WRED dropped bytes from drop precedence 1.
<b>DropPrec2AcceptedPackets</b>	Shows the accepted number of packets from drop precedence 2. The drop precedence is a function of the front-end policer.
<b>DropPrec2AcceptedBytes</b>	Shows the accepted number of bytes from drop precedence 2.
<b>DropPrec2CongestionMarkedPackets</b>	Shows the congestion marked number of packets from drop precedence 2.
<b>DropPrec2CongestionMarkedBytes</b>	Shows the congestion marked number of bytes from drop precedence 2.
<b>DropPrec2DiscardDroppedPackets</b>	Shows the number of WRED dropped packets from drop precedence 2.
<b>DropPrec2DiscardDroppedBytes</b>	Shows the number of WRED dropped bytes from drop precedence 2.
<b>DropPrec3AcceptedPackets</b>	Shows the accepted number of packets from drop precedence 3. The drop precedence is a function of the front-end policer.
<b>DropPrec3AcceptedBytes</b>	Shows the accepted number of bytes from drop precedence 3.
<b>DropPrec3CongestionMarkedPackets</b>	Shows the congestion marked number of packets from drop precedence 3.
<b>DropPrec3CongestionMarkedBytes</b>	Shows the congestion marked number of bytes from drop precedence 3.
<b>DropPrec3DiscardDroppedPackets</b>	Shows the number of WRED dropped packets from drop precedence 3.
<b>DropPrec3DiscardDroppedBytes</b>	Shows the number of WRED dropped bytes from drop precedence 3.
<b>DropPrec4AcceptedPackets</b>	Shows the accepted number of packets from drop precedence 4. The drop precedence is a function of the front-end policer.

Name	Description
<b>DropPrec4AcceptedBytes</b>	Shows the accepted number of bytes from drop precedence 4.
<b>NonWredDroppedPackets</b>	Shows the number of dropped packets due to non Random Early Detection.
<b>NonWredDroppedBytes</b>	Shows the number of dropped bytes due to non Random Early Detection.
<b>DequeuedPackets</b>	Shows the number of packets dequeued once inside the Switch Fabric.
<b>DequeuedBytes</b>	Shows the number of bytes dequeued once inside the Switch Fabric.
<b>DropPrec1DroppedPackets</b>	Shows the numbers of non-WRED dropped packets from drop precedence 1.
<b>DropPrec1DroppedBytes</b>	Shows the number of non-WRED dropped bytes from drop precedence 1.
<b>DropPrec2DroppedPackets</b>	Shows the numbers of non-WRED dropped packets from drop precedence 2.
<b>DropPrec2DroppedBytes</b>	Shows the number of non-WRED dropped bytes from drop precedence 2.
<b>DropPrec3DroppedPackets</b>	Shows the numbers of non-WRED dropped packets from drop precedence 3.
<b>DropPrec3DroppedBytes</b>	Shows the number of non-WRED dropped bytes from drop precedence 3.
<b>DropPrec4CongestionMarkedPackets</b>	Shows the congestion marked number of packets from drop precedence 4.
<b>DropPrec4CongestionMarkedBytes</b>	Shows the congestion marked number of bytes from drop precedence 4.
<b>DropPrec4DiscardDroppedPackets</b>	Shows the number of WRED dropped packets from drop precedence 4.
<b>DropPrec4DiscardDroppedBytes</b>	Shows the number of WRED dropped bytes from drop precedence 4.
<b>DropPrec4DroppedPackets</b>	Shows the numbers of non-WRED dropped packets from drop precedence 4.
<b>DropPrec4DroppedBytes</b>	Shows the number of non-WRED dropped bytes from drop precedence 4.
<b>OverSubscribeTotalDroppedPkts</b>	Shows the number of dropped packets due to free list underflow.
<b>OverSubscribeTotalDroppedBytes</b>	Shows the number of dropped bytes due to free list underflow.

Name	Description
<b>OverSubscribeGuaranteeDroppedPkts</b>	Shows the number of dropped packets due to the global buffer threshold guarantee.
<b>OverSubscribeGuaranteeDroppedBytes</b>	Shows the number of dropped bytes due to the global buffer threshold guarantee.
<b>OutPkts</b>	Shows the number of packets out from the egress interface.
<b>OutBytes</b>	Shows the number of bytes out from the egress interface.

---

## Viewing port interface statistics

View port interface statistics to manage network performance.

### Procedure

1. In the Device Physical View, select a port.
  2. In the navigation tree, open the following folders: **Configuration > Graph**.
  3. Click **Port**.
  4. Click the **Interface** tab.
- 

---

## Interface field descriptions

The following table describes parameters on the **Interface** tab.

Name	Description
<b>InOctets</b>	The number of octets received on the interface, including framing characters.
<b>OutOctets</b>	The number of octets transmitted from the interface, including framing characters.
<b>InUcastPkts</b>	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.
<b>OutUcastPkts</b>	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.

Name	Description
<b>InMulticastPkts</b>	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.
<b>OutMulticastPkts</b>	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.
<b>InBroadcastPkts</b>	The number of packets delivered by this sublayer to a higher sublayer that are addressed to a broadcast address at this sublayer.
<b>OutBroadcastPkts</b>	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.
<b>InDiscards</b>	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.
<b>InErrors</b>	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.
<b>InUnknownProtos</b>	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 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.
<b>InFlowCtrlPkts</b>	The number of flow control packets received by this interface.
<b>OutFlowCtrlPkts</b>	The number of flow control packets transmitted by this interface.
<b>NumStateTransition</b>	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. In the Device Physical View, select a port.
  2. In the navigation tree, open 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
<b>AlignmentErrors</b>	A count 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.
<b>FCSErrors</b>	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.
<b>InternalMacTransmitErrors</b>	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

Name	Description
	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.
<b>InternalMacReceiveErrors</b>	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.
<b>CarrierSenseErrors</b>	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.
<b>FrameTooLongs</b>	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, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.
<b>SQETestErrors</b>	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.
<b>DeferredTransmissions</b>	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.
<b>SingleCollisionFrames</b>	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.

Name	Description
<b>MultipleCollisionFrames</b>	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.
<b>LateCollisions</b>	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.
<b>ExcessiveCollisions</b>	A count of frames for which transmission on a particular interface fails due to excessive collisions.
<b>FrameTooShorts</b>	The number of frames, encountered on this interface, that are too short.
<b>LinkFailures</b>	The number of link failures encountered on this interface.
<b>PacketErrors</b>	The number of packet errors encountered on this interface.
<b>CarrierErrors</b>	The number of carrier errors encountered on this interface.
<b>LinkInactiveErrors</b>	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. In the Device Physical View, select a port.
  2. In the navigation tree, open 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
<b>InUnicastFrames</b>	The number of incoming unicast frames bridged.
<b>InMulticastFrames</b>	The number of incoming multicast frames bridged.
<b>InBroadcastFrames</b>	The number of incoming broadcast frames bridged.
<b>InDiscards</b>	The number of frames discarded by the bridging entity.
<b>OutFrames</b>	The number of outgoing frames bridged.

---

## Viewing port spanning tree statistics

View port spanning tree statistics to manage network performance.

### Procedure

1. In the Device Physical View, select a port.
  2. In the navigation tree, open 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
<b>InConfigBpdus</b>	The number of Config BPDUs received.
<b>InTcnBpdus</b>	The number of Topology Change Notifications BPDUs received.
<b>InBadBpdus</b>	The number of unknown or malformed BPDUs received.
<b>OutConfigBpdus</b>	The number of Config BPDUs transmitted.



Name	Description
<b>OutTcnBpdus</b>	The number of Topology Change Notifications BPDUs transmitted.

---

## Viewing port routing statistics

View port routing statistics to manage network performance.

### Procedure

1. In the Device Physical View, select a port.
  2. In the navigation tree, open 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
<b>InUnicastFrames</b>	The number of incoming unicast frames routed.
<b>InMulticastFrames</b>	The number of incoming multicast frames routed.
<b>InDiscards</b>	The number of frames discarded by the routing entity.
<b>OutUnicastFrames</b>	The number of outgoing unicast frames routed.
<b>OutMulticastFrames</b>	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 tree, open the following folders: **Configuration > IPv6**.
2. Click **IPv6**.

3. Click the **Interfaces** tab.
4. Select an interface.
5. Click **IfStats**.
6. Optionally, you can 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
<b>InReceives</b>	Shows the total number of input datagrams received by the interface, including those received in error.
<b>InHdrErrors</b>	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.
<b>InTooBigErrors</b>	Shows the number of input datagrams that could not be forwarded because their size exceeded the link MTU of the outgoing interface.
<b>InNoRoutes</b>	Shows the number of input datagrams discarded because no route could be found to transmit them to their destination.
<b>InAddrErrors</b>	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.
<b>InUnknownProtos</b>	Shows the number of locally-addressed datagrams received successfully but discarded because of an unknown or

Name	Description
	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.
<b>InTruncatedPkts</b>	Shows the number of input datagrams discarded because the datagram frame did not carry enough data.
<b>InDiscards</b>	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.
<b>InDelivers</b>	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.
<b>OutForwDatagrams</b>	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 Source-Route processing was successful. For a successfully forwarded datagram the counter of the outgoing interface is incremented.
<b>OutRequests</b>	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 <b>OutForwDatagrams</b> .
<b>OutDiscards</b>	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

Name	Description
	<b>OutForwDatagrams</b> if such packets met this (discretionary) discard criterion.
<b>OutFragOKs</b>	Shows the number of IPv6 datagrams that have been successfully fragmented at this output interface.
<b>OutFragFails</b>	Shows the number of IPv6 datagrams that have been discarded because they needed to be fragmented at this output interface but could not be.
<b>OutFragCreates</b>	Shows the number of output datagram fragments that have been generated as a result of fragmentation at this output interface.
<b>ReasmReqds</b>	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.
<b>ReasmOKs</b>	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.
<b>ReasmFails</b>	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.
<b>InMcastPkts</b>	Shows the number of multicast packets received by the interface.
<b>OutMcastPkts</b>	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 tree, open the following folders: **Configuration > IP**.
  2. Click **DHCP Relay**.
  3. Click the **Interface Stats** tab.
- 

---

## Interface Stats field descriptions

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

Name	Description
<b>IfIndex</b>	Identifies the physical interface.
<b>AgentAddr</b>	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.
<b>NumRequests</b>	Shows the number of DHCP and BootP requests on this interface.
<b>NumReplies</b>	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. In the Device Physical view, select a port.
2. In the navigation tree, open the following folders: **Configuration > IPv6**
3. Click the **DHCP Relay** tab.

4. Click the **Interface** tab.
  5. Click **Statistics**.
  6. Select one or more values.
  7. Click the type of graph.
- 

---

## Statistics field descriptions

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

Name	Description
<b>NumRequests</b>	Shows the number of DHCP and BootP requests on this interface.
<b>NumReplies</b>	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. In the Device Physical View, select a port.
  2. In the navigation tree, open 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
<b>NumRequests</b>	The number of DHCP and/or BootP requests on this interface.
<b>NumReplies</b>	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. In the Device Physical view, select a port.
  2. In the navigation tree, 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
<b>NumRequests</b>	The number of DHCP and BootP requests on this interface.
<b>NumReplies</b>	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 tree, open 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
<b>NumRequests</b>	The number of DHCP and BootP requests on this interface.
<b>NumReplies</b>	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 tree, open 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
<b>IfIndex</b>	Shows the name of the interface on which you enabled option 82. Shows the port



Name	Description
	number if the interface is a brouter port or the VLAN number if the interface is a VLAN.
<b>AgentAddr</b>	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.
<b>FoundOp82</b>	Shows the number of packets that the interface received that already had option82 in them.
<b>Dropped</b>	Shows the number of packets the interface dropped because of option 82–related issues. To determine the cause of the drop, you must enable trace on level 170.
<b>CircuitId</b>	Show the value inserted in the packets as the circuit ID. The value is the index of the interface.
<b>AddedCircuitId</b>	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.
<b>RemovedCircuitId</b>	Shows on how many packets (replies from server to client) the circuit id was removed for that interface.
<b>RemotId</b>	Shows the value inserted in the packets as the remote ID. The value is the MAC address of the interface.
<b>AddedRemotId</b>	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.
<b>RemovedRemotId</b>	Shows on 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. In the Device Physical View, select a port.
  2. In the navigation tree, open 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
<b>VersionMismatches</b>	The number of version mismatches received by this interface.
<b>AreaMismatches</b>	The number of area mismatches received by this interface.
<b>AuthTypeMismatches</b>	The number of authentication type mismatches received by this interface.
<b>AuthFailures</b>	The number of authentication failures.
<b>NetmaskMismatches</b>	The number of net mask mismatches received by this interface.
<b>HelloIntervalMismatches</b>	The number of hello interval mismatches received by this interface.
<b>DeadIntervalMismatches</b>	The number of dead interval mismatches received by this interface.
<b>OptionMismatches</b>	The number of option mismatches in the hello interval or dead interval fields received by this interface.
<b>RxHellos</b>	The number of hello packets received by this interface.
<b>RxDBDescrs</b>	The number of database descriptor packets received by this interface.

Name	Description
<b>RxLSUpdates</b>	The number of link state update packets received by this interface.
<b>RxLSReqs</b>	The number of link state request packets received by this interface.
<b>RxLSAcks</b>	The number of link state acknowledge packets received by this interface.
<b>TxHellos</b>	The number of hello packets transmitted by this interface.
<b>TxDBDescrs</b>	The number of database descriptor packets transmitted by this interface.
<b>TxLSUpdates</b>	The number of link state update packets transmitted by this interface.
<b>TxLSReqs</b>	The number of link state request packets transmitted by this interface.
<b>TxLSAcks</b>	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. In the Device Physical View, select a port.
  2. In the navigation tree, open the following folders: **Configuration > Graph**.
  3. Click **Port**.
  4. Click the **LACP** tab.
  5. To change the poll interval, 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
<b>LACPDUsRx</b>	The number of valid LACPDU received on this aggregation port.
<b>MarkerPDUsRx</b>	The number of valid marker PDUs received on this aggregation port.
<b>MarkerResponsePDUsRx</b>	The number of valid marker response PDUs received on this aggregation port.
<b>UnknownRx</b>	The number of frames received that either: <ul style="list-style-type: none"> <li>• carry Slow Protocols Ethernet type values, but contain an unknown PDU.</li> <li>• are addressed to the Slow Protocols group MAC Address, but do not carry the Slow Protocols Ethernet Type.</li> </ul>
<b>IllegalRx</b>	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).
<b>LACPDUsTx</b>	The number of LACPDU transmitted on this aggregation port.
<b>MarkerPDUsTx</b>	The number of marker PDUs transmitted on this aggregation port.
<b>MarkerResponsePDUsTx</b>	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 tree, open 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
<b>TotalPkts</b>	Shows the total number of packets received on the port.
<b>TotalBytes</b>	Shows the total number of bytes received on the port.
<b>YellowBytes</b>	Shows the total number of bytes received on the port that were above the committed rate but below the peak rate.
<b>RedBytes</b>	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 tree, open 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
<b>Slot</b>	Specifies the slot number of the CP module.
<b>FlashBytesUsed</b>	Specifies the number of bytes used in internal flash memory.
<b>FlashBytesFree</b>	Specifies the number of bytes available for use in internal flash memory.

Name	Description
<b>FlashNumFiles</b>	Specifies the number of files in internal flash memory.
<b>ExtflashBytesUsed</b>	Specifies the number of bytes used in external flash memory.
<b>ExtflashBytesFree</b>	Specifies the number of bytes available for use in external flash memory.
<b>ExtflashNumFiles</b>	Specifies the number of files in external flash memory.
<b>UsbBytesUsed</b>	Specifies the number of bytes used on the USB device.
<b>UsbBytesFree</b>	Specifies the number of bytes available on the USB device.
<b>UsbNumFiles</b>	Specifies the number of files on the USB device.

---

## Viewing QoS policy statistics

### About this task

Use policy statistics to better tailor policy parameters to suit customer needs.

### Procedure

1. In the navigation tree, open 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
<b>PolicyId</b>	Shows the global policer ID that corresponds to this local policer.
<b>Slot</b>	Shows the slot number of the chassis on which statistics are collected. Valid slots are IO slots.
<b>TotalPkts</b>	Shows the global policer total packet count.
<b>TotalBytes</b>	Shows the global policer total byte count.
<b>GreenPackets</b>	Shows the total number of packets received that were below the committed rate.

Name	Description
<b>GreenBytes</b>	Shows the total number of bytes received that were below the committed rate.
<b>YellowPackets</b>	Shows the total number of packets received that were above the committed rate but below the peak rate.
<b>YellowBytes</b>	Shows the total number of bytes received that were above the committed rate but below the peak rate.
<b>RedPackets</b>	Shows the total number of packets received that were above the peak rate.
<b>RedBytes</b>	Shows the total number of bytes received that were above the peak rate.

---

## Graphing QoS policy statistics

### About this task

Graph QoS policy statistics to create a visual comparison between data values.

### Procedure

1. In the navigation tree, open 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
<b>TotalPkts</b>	Shows the total packet count.

Name	Description
<b>TotalBytes</b>	Shows the total byte count.
<b>GreenPackets</b>	Shows the total number of packets received that were below the committed rate.
<b>GreenBytes</b>	Shows the total number of bytes received that were below the committed rate.
<b>YellowPackets</b>	Shows the total number of packets received that were above the committed rate but below the peak rate.
<b>YellowBytes</b>	Shows the total number of bytes received that were above the committed rate but below the peak rate.
<b>RedPackets</b>	Shows the total number of packets received that were above the peak rate.
<b>RedBytes</b>	Shows the total number of bytes received that were above the peak rate.

---

## Viewing statistics for a specific QoS policy

### About this task

View statistics for a specific QoS policy.

### Procedure

1. In the navigation tree, open 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
<b>PolicyId</b>	Shows the global policer ID that corresponds to this local policer.



Name	Description
<b>Slot</b>	Shows the slot number of the chassis on which statistics are collected. Valid slots are IO slots.
<b>TotalPkts</b>	Shows the global policer total packet count.
<b>TotalBytes</b>	Shows the global policer total byte count.
<b>GreenPackets</b>	Shows the total number of packets received that were below the committed rate.
<b>GreenBytes</b>	Shows the total number of bytes received that were below the committed rate.
<b>YellowPackets</b>	Shows the total number of packets received that were above the committed rate but below the peak rate.
<b>YellowBytes</b>	Shows the total number of bytes received that were above the committed rate but below the peak rate.
<b>RedPackets</b>	Shows the total number of packets received that were above the peak rate.
<b>RedBytes</b>	Shows the total number of bytes received that were above the peak rate.

---

## Viewing ACE port statistics

### About this task

Use port statistics to ensure that the ACE is operating correctly.

### Procedure

1. In the navigation tree, open the following folders: **Configuration > Security > Data Path**.
  2. Click **Advanced Filters (ACE/ACLs)**.
  3. Click the **ACL** tab.
  4. Select 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
<b>AcId</b>	Specifies the associated ACL index.
<b>AcId</b>	Specifies the ACE index.
<b>MatchCountPkts</b>	Specifies a packet count of the matching packets.
<b>MatchCountOctets</b>	Specifies the number of octets of the matching packets.

---

## Viewing ACL statistics

### About this task

Graph statistics for a specific ACL ID to view default statistics.

### Procedure

1. In the navigation tree, open the following folders: **Configuration > Security > Data Path**.
2. Click **Advanced Filters (ACE/ACLs)**.
3. Click the **ACL** tab.
4. Select 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
<b>AcId</b>	Specifies the ACL ID.

Name	Description
<b>MatchDefaultSecurityPkts</b>	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.
<b>MatchDefaultSecurityOctets</b>	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.
<b>MatchDefaultQosPkts</b>	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.
<b>MatchDefaultQosOctets</b>	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.
<b>MatchGlobalSecurityPkts</b>	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.
<b>MatchGlobalSecurityOctets</b>	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.
<b>MatchGlobalQosPkts</b>	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.
<b>MatchGlobalQosOctets</b>	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

### About this task

Clear ACL statistics when you want to gather a new set of statistics.

### Procedure

1. In the navigation tree, open the following folders: **Configuration > Security > Data Path**.
2. Click **Advanced Filters (ACE/ACLs)**.
3. Click the **ACL** tab.
4. Select a field.

5. Click **ClearStats**.
- 

---

## Viewing VLAN and Spanning Tree CIST statistics

### About this task

View CIST port statistics to manage network performance.

### Procedure

1. In the navigation tree, open 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
<b>ForwardTransitions</b>	Number of times this port has transitioned to the forwarding state.
<b>RxMstBpduCount</b>	Number of MSTP BPDUs received on this port.
<b>RxRstBpduCount</b>	Number of RSTP BPDUs received on this port.
<b>RxConfigBpduCount</b>	Number of configuration BPDUs received on this port.
<b>RxTcnBpduCount</b>	Number of TCN BPDUs received on this port.
<b>TxMstBpduCount</b>	Number of MSTP BPDUs transmitted from this port.
<b>TxRstBpduCount</b>	Number of RSTP BPDUs transmitted from this port.
<b>TxConfigBpduCount</b>	Number of configuration BPDUs transmitted from this port.
<b>TxTcnBpduCount</b>	Number of TCN BPDUs transmitted from this port.
<b>InvalidMstBpduRxCount</b>	Number of Invalid MSTP BPDUs received on this port.
<b>InvalidRstBpduRxCount</b>	Number of Invalid RSTP BPDUs received on this port.

Name	Descriptions
<b>InvalidConfigBpduRxCount</b>	Number of invalid configuration BPDUs received on this port.
<b>InvalidTcnBpduRxCount</b>	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.
<b>ProtocolMigrationCount</b>	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

### About this task

View multiple spanning tree instance (MSTI) port statistics to manage network performance.

### Procedure

1. In the navigation tree, open 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
<b>ForwardTransitions</b>	Number of times this port has transitioned to the forwarding state for this specific instance.
<b>ReceivedBPDUs</b>	Number of BPDUs received by this port for this spanning tree instance.
<b>TransmittedBPDUs</b>	Number of BPDUs transmitted on this port for this spanning tree instance.

Name	Description
<b>InvalidBPDUsRcvd</b>	Number of invalid BPDUs received on this port for this spanning tree instance.

---

## Viewing VRRP interface stats

### About this task

View VRRP statistics to manage network performance.

### Procedure

1. In the navigation tree, open 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
<b>BecomeMaster</b>	The number of times that this virtual router state transitions from BACKUP to MASTER.
<b>AdvertiseRcvd</b>	The number of VRRP advertisements received by this virtual router.
<b>AdvertiseIntervalErrors</b>	The number of received VRRP advertisement packets with a different interval is than configured for the local virtual router.
<b>IPtTlErrors</b>	The number of VRRP packets received by the virtual router with IP TTL (Time-To-Live) not equal to 255.
<b>PriorityZeroPktsRcvd</b>	The number of VRRP packets received by the virtual router with a priority of 0.

Name	Description
<b>PriorityZeroPktsSent</b>	The number of VRRP packets sent by the virtual router with a priority of 0'.
<b>InvalidTypePktsRcvd</b>	The number of VRRP packets received by the virtual router with an invalid value in the 'type' field.
<b>AddressListErrors</b>	Packets received address list the address list does not match the locally configured list for the virtual router.
<b>AuthTypeMismatch</b>	The count of authentication type mismatch messages.
<b>PacketLengthErrors</b>	The count of packet length errors.
<b>AuthFailures</b>	The count of authentication failure messages.

---

## Viewing VRRP statistics

### About this task

View VRRP statistics to monitor network performance.

### Procedure

1. In the navigation tree, open 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
<b>ChecksumErrors</b>	The number of VRRP packets received with an invalid VRRP checksum value.
<b>VersionErrors</b>	The number of VRRP packets received with an unknown or unsupported version number.

Name	Description
<b>VrIDErrors</b>	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 tree, open 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
<b>MasterTransitions</b>	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.
<b>RcdAdvertisements</b>	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.
<b>AdvIntervalErrors</b>	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



Name	Description
	management system, and at other times as indicated by the value of DiscontinuityTime.
<b>IpTtlErrors</b>	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.
<b>RcvdPriZeroPackets</b>	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.
<b>SentPriZeroPackets</b>	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.
<b>RcvdInvalidTypePkts</b>	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.
<b>AddressListErrors</b>	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.
<b>PacketLengthErrors</b>	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.
<b>RcvdInvalidAuthentications</b>	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 tree, open 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
<b>ChecksumErrors</b>	Shows the number of VRRP packets received with an invalid VRRP checksum value.
<b>VersionErrors</b>	Shows the number of VRRP packets received with an unknown or unsupported version number.
<b>VrIDErrors</b>	Shows the number of VRRP packets received with an invalid VrID for this virtual router.

---

## Viewing SMLT statistics

### About this task

View SMLT statistics to manage network performance.

### Procedure

1. In the navigation tree, open 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
<b>SmltIstDownCnt</b>	The number of times the session between the two peering switches has gone down since last boot.
<b>SmltHelloTxMsgCnt</b>	The count of transmitted hello messages.
<b>SmltHelloRxMsgCnt</b>	The count of received hello messages.
<b>SmltLearnMacAddrTxMsgCnt</b>	The count of transmitted learned MAC address messages.
<b>SmltLearnMacAddrRxMsgCnt</b>	The count of received learned MAC address messages.
<b>SmltMacAddrAgeOutTxMsgCnt</b>	The count of transmitted aging out MAC address messages.
<b>SmltMacAddrAgeOutRxMsgCnt</b>	The count of received aging out MAC address messages.
<b>SmltMacAddrAgeExpTxMsgCnt</b>	The count of transmitted MAC address age expired messages.
<b>SmltMacAddrAgeExpRxMsgCnt</b>	The count of received MAC address age expired messages.
<b>SmltStgInfoTxMsgCnt</b>	The count of transmitted STG information messages.
<b>SmltStgInfoRxMsgCnt</b>	The count of received STG information messages.
<b>SmltDelMacAddrTxMsgCnt</b>	The count of transmitted MAC address deleted messages.
<b>SmltDelMacAddrRxMsgCnt</b>	The count of received MAC address received messages.
<b>SmltSmltDownTxMsgCnt</b>	The count of transmitted SMLT down messages.
<b>SmltSmltDownRxMsgCnt</b>	The count of received SMLT down messages.
<b>SmltUpTxMsgCnt</b>	The count of transmitted SMLT up messages.

Name	Description
<b>SmltUpRxMsgCnt</b>	The count of received SMLT up messages.
<b>SmltSendMacTblTxMsgCnt</b>	The count of sent send MAC table messages.
<b>SmltSendMacTblRxMsgCnt</b>	The count of received send MAC table messages.
<b>SmltIcmpTxMsgCnt</b>	The count of sent IGMP messages.
<b>SmltIcmpRxMsgCnt</b>	The count of received IGMP messages.
<b>SmltPortDownTxMsgCnt</b>	The count of sent port down messages.
<b>SmltPortDownRxMsgCnt</b>	The count of received port down messages.
<b>SmltReqMacTblTxMsgCnt</b>	The count or sent MAC table request messages.
<b>SmltReqMacTblRxMsgCnt</b>	The count of received MAC table request messages.
<b>SmltRxUnknownMsgTypeCnt</b>	The count of received unknown message type messages.

---

## Viewing RSTP status statistics

### About this task

You can view status statistics for Rapid Spanning Tree Protocol (RSTP).

### Procedure

1. In the navigation tree, open 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
<b>RxRstBpduCount</b>	The number of RSTP BPDUs this port received.

Name	Description
<b>RxConfigBpduCount</b>	The number of configuration BPDUs this port received.
<b>RxTcnBpduCount</b>	The number of TCN BPDUs this port received.
<b>TxRstBpduCount</b>	The number of RSTP BPDUs this port transmitted.
<b>TxConfigBpduCount</b>	The number of Config BPDUs this port transmitted.
<b>TxTcnBpduCount</b>	The number of TCN BPDUs this port transmitted.
<b>InvalidRstBpduRxCount</b>	The number of invalid RSTP BPDUs this port received. A trap is generated on the occurrence of this event.
<b>InvalidConfigBpduRx Count</b>	The number of invalid configuration BPDUs this port received. A trap is generated on the occurrence of this event.
<b>InvalidTcnBpduRxCount</b>	The number of invalid TCN BPDUs this port received. A trap is generated on the occurrence of this event.
<b>ProtocolMigrationCount</b>	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 interface statistics

### About this task

Use MLT interface statistics tab to view interface statistics for the selected MLT.

### Procedure

1. In the navigation tree, open 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
<b>InOctets</b>	The total number of octets received on the MLT interface, including framing characters.
<b>OutOctets</b>	The total number of octets transmitted out of the MLT interface, including framing characters.
<b>InUcastPkts</b>	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.
<b>OutUcastPkts</b>	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.
<b>InMulticastPkt</b>	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.
<b>OutMulticast</b>	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.
<b>InBroadcastPkt</b>	The number of packets delivered to this MLT that were addressed to a broadcast address at this sublayer.
<b>OutBroadcast</b>	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.

---

## Viewing MLT Ethernet error statistics

### About this task

Use MLT Ethernet error statistics to view the error statistics.

### Procedure

1. In the navigation tree, open 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
<b>AlignmentErrors</b>	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.
<b>FCSErrors</b>	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.
<b>IMacTransmitError</b>	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.
<b>IMacReceiveError</b>	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 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 receive errors on a particular interface that are not otherwise counted.
<b>CarrierSenseError</b>	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.
<b>FrameTooLong</b>	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

Name	Description
	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.
<b>SQETestError</b>	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.
<b>DeferredTransmiss</b>	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.
<b>SingleCollFrames</b>	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.
<b>MultipleCollFrames</b>	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.
<b>LateCollisions</b>	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.
<b>ExcessiveCollis</b>	The frame count for which transmission on a particular MLT fails due to excessive collisions.

---

## Viewing RIP statistics

### About this task

Use statistics to help you monitor Routing Information Protocol (RIP) performance. You can also use statistics in troubleshooting procedures.



## Procedure

1. In the navigation tree, open 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
<b>Address</b>	The IP address of the router interface.
<b>RcvBadPackets</b>	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).
<b>RcvBadRoutes</b>	The number of routes, in valid RIP packets, that were ignored for any reason (examples: unknown address family or invalid metric).
<b>SentUpdates</b>	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

### About this task

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 tree, open 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
<b>LsdbTblSize</b>	The number of entries in the link state database table.
<b>TxPackets</b>	The number of packets transmitted by OSPF.
<b>RxPackets</b>	The number of packets received by OSPF.
<b>TxDropPackets</b>	The number of packets dropped before being transmitted by OSPF.
<b>RxDropPackets</b>	The number of packets dropped before they are received by OSPF.
<b>RxBadPackets</b>	The number of packets received by OSPF that are bad.
<b>SpfRuns</b>	The number of SPF calculations performed by OSPF.
<b>BuffersAllocated</b>	The number of buffers allocated for OSPF.
<b>BuffersFreed</b>	The number of buffers freed by OSPF.
<b>BufferAllocFailures</b>	The number of times that OSPF has failed to allocate buffers.
<b>BufferFreeFailures</b>	The number of times that OSPF has failed to free buffers.
<b>Routes</b>	The count of OSPF routes.
<b>Adjacencies</b>	The count of OSPF adjacencies.
<b>Areas</b>	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 tree, open 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
<b>TxPackets</b>	Shows the count of sent packets.
<b>RxPackets</b>	Shows the count of received packets.
<b>TxDropPackets</b>	Shows the count of sent, dropped packets.
<b>RxDropPackets</b>	Shows the count of received, dropped packets.
<b>RxBadPackets</b>	Shows the count of received, bad packets.
<b>SpfRuns</b>	Shows the count of intra-area route table updates with calculations using this area link-state database.
<b>LastSpfRun</b>	Shows the count of the most recent SPF run.
<b>LsdbTblSize</b>	Shows the size of the link state database table.
<b>BadLsReqs</b>	Shows the count of bad link requests.
<b>SeqMismatches</b>	Shows the count of sequence mismatched packets.

---

## Graphing OSPF statistics for a VLAN

### About this task

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 tree, open 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
<b>VersionMismatches</b>	Indicates the number of version mismatches received by this interface.
<b>AreaMismatches</b>	Indicates the number of area mismatches received by this interface.
<b>AuthTypeMismatches</b>	Indicates the number of AuthType mismatches received by this interface.
<b>AuthFailures</b>	Indicates the number of authentication failures.
<b>NetMaskMistmatches</b>	Indicates the number of net mask mismatches received by this interface.
<b>HelloIntervalMismatches</b>	Indicates the number of hello interval mismatches received by this interface.
<b>DeadIntervalMismatches</b>	Indicates the number of dead interval mismatches received by this interface.
<b>OptionMismatches</b>	Indicates the number of options mismatches received by this interface.
<b>RxHellos</b>	Indicates the number of hello packets received by this interface.
<b>RxDBDescrs</b>	Indicates the number of database descriptor packets received by this interface.
<b>RxLSUpdates</b>	Indicate the number of Link state update packets received by this interface.
<b>RxLsReqs</b>	Indicates the number of Link state request packets received by this interface.
<b>RxLSAcks</b>	Indicates the number of Link state acknowledge packets received by this interface.
<b>TxHellos</b>	Indicates the number of hello packets transmitted by this interface.

Name	Description
<b>TxDBDescrs</b>	Indicates the number of database descriptor packets transmitted by this interface.
<b>TxLSUpdates</b>	Indicate the number of Link state update packets transmitted by this interface.
<b>TxLSReqs</b>	Indicates the number of Link state request packets transmitted by this interface.
<b>TxLSAcks</b>	Indicates the number of Link state acknowledge packets transmitted by this interface.

---

## Graphing OSPF statistics for a port

### About this task

Use statistics to help you monitor OSPF performance on a VLAN. You can also graph statistics for all OSPF packets.

### Procedure

1. In the Device Physical View, select a port.
  2. In the navigation tree, open 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
<b>VersionMismatches</b>	Indicates the number of version mismatches received by this interface.
<b>AreaMismatches</b>	Indicates the number of area mismatches received by this interface.

Name	Description
<b>AuthTypeMismatches</b>	Indicates the number of AuthType mismatches received by this interface.
<b>AuthFailures</b>	Indicates the number of authentication failures.
<b>NetMaskMistmatches</b>	Indicates the number of net mask mismatches received by this interface.
<b>HelloIntervalMismatches</b>	Indicates the number of hello interval mismatches received by this interface.
<b>DeadIntervalMismatches</b>	Indicates the number of dead interval mismatches received by this interface.
<b>OptionMismatches</b>	Indicates the number of options mismatches received by this interface.
<b>RxHellos</b>	Indicates the number of hello packets received by this interface.
<b>RxDBDescrs</b>	Indicates the number of database descriptor packets received by this interface.
<b>RxLSUpdates</b>	Indicate the number of Link state update packets received by this interface.
<b>RxLsReqs</b>	Indicates the number of Link state request packets received by this interface.
<b>RxLSAcks</b>	Indicates the number of Link state acknowledge packets received by this interface.
<b>TxHellos</b>	Indicates the number of hello packets transmitted by this interface.
<b>TxDBDescrs</b>	Indicates the number of database descriptor packets transmitted by this interface.
<b>TxLSUpdates</b>	Indicate the number of Link state update packets transmitted by this interface.
<b>TxLSReqs</b>	Indicates the number of Link state request packets transmitted by this interface.
<b>TxLSAcks</b>	Indicates the number of Link state acknowledge packets transmitted by this interface.

---

## Viewing global BGP statistics

### About this task

View global BGP statistics to manage network performance.

### Procedure

1. In the navigation tree, open the following folders: **Configuration > IP**.
  2. Click **BGP**.
  3. Select the **Global Stats** tab to view BGP global statistics.
- 

---

## Global Stats field descriptions

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

Name	Description
<b>Starts</b>	Number of times BGP connection started
<b>Stops</b>	Number of times BGP connection stopped
<b>Opens</b>	Number of times BGP connection opened TCP
<b>Closes</b>	Number of times BGP connection closed TCP
<b>Fails</b>	Number of times a TCP attempt failed
<b>Fatals</b>	Number of times TCP crashed due to fatal error
<b>ConnExps</b>	Number of times the TCP retry timer expired
<b>HoldExps</b>	Number of times the hold timer expired
<b>KeepExps</b>	Number of times the keepalive timer expired
<b>RxOpens</b>	Number of Opens received by BGP
<b>RxKeeps</b>	Number of Keepalive messages received by BGP
<b>RxUpdates</b>	Number of Updates received by BGP
<b>RxNotifys</b>	Number of Notifications received by BGP
<b>TxOpens</b>	Number of transmitted by BGP

Name	Description
<b>TxKeeps</b>	Number of Keepalive messages transmitted by BGP
<b>TxUpdates</b>	Number of Updates transmitted by BGP
<b>TxNotifys</b>	Number of Notifications transmitted by BGP
<b>BadEvents</b>	Number of invalid events received by FSM
<b>SyncFails</b>	Number of times the FDB sync failed
<b>TrEvent</b>	Trace event
<b>RxECodeHeader</b>	number of Header errors received
<b>RxECodeOpen</b>	number of Open errors received
<b>RxECodeUpdate</b>	number of Update errors received
<b>RxRxECodeHoldtimer</b>	number of Hold Timer Expired errors received
<b>RxECodeFSM</b>	number of FSM errors received
<b>RxECodeCease</b>	number of Cease errors received
<b>RxHrCodeNoSync</b>	Number of Header errors received as: Not Synchronized
<b>RxHdrCodeInvalidMsgLen</b>	Number of Header errors received as: Invalid Msg len
<b>RxHdrCodeInvalidMsgType</b>	Number of Header errors received as: Invalid Msg type
<b>RxOpCodeBadVer</b>	Number of Open errors received as: Bad version
<b>RxOpCodeBadAs</b>	Number of Open errors received as: Bad AS number
<b>RxOpCodeBadRtID</b>	Number of Open errors received as: Bad BGP Rtr ID
<b>RxOpCodeUnsuppOption</b>	Number of Open errors received as: Unsupported Option
<b>RxOpCodeAuthFail</b>	Number of Open errors received as: Auth Failure
<b>RxOpCodeBadHold</b>	Number of Open errors received as: Bad Hold Value
<b>RxUpdCodeMalformedAttrList</b>	Number of Update errors received as: Malformed Attr List
<b>RxUpdCodeWellknownAttrUnrecog</b>	Number of Update errors received as: Wellknown Attr Unrecog



Name	Description
<b>RxUpdCodeWellknownAttrMiss</b>	Number of Update errors received as: Welknown Attr Missing
<b>RxUpdCodeAttrFlagError</b>	Number of Update errors received as: Attr Flag Error
<b>RxUpdCodeAttrLenError</b>	Number of Update errors received as: Attr Len Error
<b>RxUpdCodeBadORIGINAttr</b>	Number of Update errors received as: Bad ORIGIN Attr
<b>RxUpdCodeASRoutingLoop</b>	Number of Update errors received as: AS Routing Loop

<b>RxUpdCodeBadNHAttr</b>	Number of Update errors received as: Bad NEXT-HOP Attr
<b>RxUpdCodeOptionalAttrError</b>	Number of Update errors received as: Optional Attr Error
<b>RxUpdCodeBadNetworkField</b>	Number of Update errors received as: Bad Network Field
<b>RxUpdCodeMalformedASPath</b>	Number of Update errors received as: Malformed AS Path
<b>TxECodeHeader</b>	number of Header errors transmitted
<b>TxECodeOpen</b>	number of Open errors transmitted
<b>TxECodeUpdate</b>	number of Update errors transmitted
<b>TxECodeHoldtimer</b>	number of Hold Timer Expired errors transmitted
<b>TxECodeFSM</b>	number of FSM errors transmitted
<b>TxECodeCease</b>	number of Cease errors transmitted
<b>TxHdrCodeNoSync</b>	Number of Header errors transmitted as: Not Synchronized
<b>TxHdrCodeInvalidMsgLen</b>	Number of Header errors transmitted as: Invalid Msg len
<b>TxHdrCodeInvalidMsgType</b>	Number of Header errors transmitted as: Invalid Msg type
<b>TxOpCodeBadVer</b>	Number of Open errors transmitted as: Bad version
<b>TxOpCodeBadAs</b>	Number of Open errors transmitted as: Bad AS number

<b>TxOpCodeBadRtrID</b>	Number of Open errors transmitted as: Bad BGP Rtr ID
<b>TxOpCodeUnsuppOption</b>	Number of Open errors transmitted as: Unsupported Option
<b>TxOpCodeAuthFail</b>	Number of Open errors transmitted as: Auth Failure
<b>TxOpCodeBadHold</b>	Number of Open errors transmitted as: Bad Hold Value
<b>TxUpdCodeMalformedAttrList</b>	Number of Update errors transmitted as: Malformed Attr List
<b>TxUpdCodeWelknownAttrUnrecog</b>	Number of Update errors transmitted as: Welknown Attr Unrecog
<b>TxUpdCodeWelknownAttrMiss</b>	Number of Update errors transmitted as: Welknown Attr Missing
<b>TxUpdCodeAttrFlagError</b>	Number of Update errors transmitted as: Attr Flag Error
<b>TxUpdCodeAttrLenError</b>	Number of Update errors transmitted as: Attr Len Error
<b>TxUpdCodeBadORIGINAttr</b>	Number of Update errors transmitted as: Bad ORIGIN Attr
<b>TxUpdCodeASRoutingLoop</b>	Number of Update errors transmitted as: AS Routing Loop
<b>TxUpdCodeBadNHAttr</b>	Number of Update errors transmitted as: Bad NEXT-HOP Attr
<b>TxUpdCodeOptionalAttrError</b>	Number of Update errors transmitted as: Optional Attr Error
<b>TxUpdCodeBadNetworkField</b>	Number of Update errors transmitted as: Bad Network Field
<b>TxUpdCodeMalformedASPath</b>	Number of Update errors transmitted as: Malformed AS Path

---

## Viewing BGP peer general statistics

### About this task

View BGP peer general statistics to manage network performance.

## Procedure

1. In the navigation tree, open 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
<b>InUpdates</b>	Number of updates received by the peer.
<b>OutUpdates</b>	Number of updates transmitted by the peer.
<b>InTotalMessages</b>	Number of total messages received by the peer.
<b>OutTotalMessages</b>	Number of total messages transmitted by the peer.
<b>FsmEstablishedTransitions</b>	The number of times the BGP FSM transitioned into the established state.
<b>FsmEstablishedTime</b>	Number of seconds this peer has been in the Established state or how long since this peer was last in the Established state. This value is set to zero when a new peer is configured or the router is booted.
<b>InUpdateElapsedTime</b>	Elapsed time in seconds since the last BGP UPDATE message was received from the peer.
<b>Starts</b>	Number of times peer BGP connection started.
<b>Stops</b>	Number of times peer BGP connection stopped.
<b>Opens</b>	Number of times peer opened TCP.
<b>Closes</b>	Number of times peer closed TCP.

Name	Description
<b>Fails</b>	Number of times a peer TCP attempt failed.
<b>Fatals</b>	Number of times peer TCP crashed due to fatal error.
<b>ConnExps</b>	Number of times the peer TCP retry timer expired.
<b>HoldExps</b>	Number of times the peer hold timer expired.
<b>KeepExps</b>	Number of times the peer keepalive timer expired.
<b>BadEvents</b>	Number of invalid events received by the peer.
<b>SyncFails</b>	Number of times the peer FDB sync failed.
<b>RcvdTooShort</b>	Number of "too short" messages received by the peer.
<b>NoMarker</b>	Number of messages received by the peer with no marker.
<b>Dropped</b>	Number of messages dropped by the peer.
<b>BadMsgTypes</b>	Number of messages received by the peer as "invalid message type."
<b>TrEvent</b>	Peer trace event.

---

## Viewing BGP peer advanced statistics

### About this task

View BGP peer advance statistics to manage network performance.

### Procedure

1. In the navigation tree, open 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
<b>InUpdates</b>	Specifies the number of BGP update messages received on this connection. This object must be initialized to zero (0) when the connection is established.
<b>OutUpdates</b>	Specifies the number of BGP update messages transmitted on this connection. This object must be initialized to zero (0) when the connection is established.
<b>InTotalMessages</b>	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.
<b>OutTotalMessages</b>	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.
<b>FsmEstablishedTransitions</b>	Specifies the total number of times the BGP FSM transitioned into the established state.
<b>FsmEstablishedTime</b>	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.
<b>InUpdatesElapsedTime</b>	Specifies the elapsed time (in seconds) since the last BGP update message was received from the peer. Each time bgpPeerInUpdates is incremented, the value of this object is set to zero (0).

---

## Viewing BGP peer receive statistics

### About this task

View BGP peer receive statistics to manage network performance.

## Procedure

1. In the navigation tree, open 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
<b>RxMsgs</b>	Number of messages received by the peer.
<b>RxInCompPkts</b>	Number of Incomplete messages received by the peer.
<b>RxOpens</b>	Number of Opens received by the peer.
<b>RxKeeps</b>	Number of Keepalive messages received by the peer.
<b>RxUpdates</b>	Number of Updates received by the peer.
<b>RxNotifys</b>	Number of Notifications received by the peer.
<b>RxRoutesAdded</b>	Number of routes added into loc_rib by this peer.
<b>RxRoutesReplaced</b>	Number of routes that were replaced by routes received by the peer.
<b>RxNlri</b>	Number of network layer reachability information (NLRI) messages received by the peer.
<b>RxValidUpdates</b>	Number of valid Updates received by the peer.
<b>RxECodeHeader</b>	Number of Header errors received by the peer.
<b>RxECodeOpen</b>	Number of Open errors received by the peer.

Name	Description
<b>RxECodeUpdate</b>	Number of Update errors received by the peer.
<b>RxECodeHoldtimer</b>	Number of Hold errors received by the peer.
<b>RxECodeFSM</b>	Number of FSM errors received by the peer.
<b>RxECodeCease</b>	Number of Cease errors received by the peer.
<b>RxHdrCodeNoSync</b>	Number of Header errors received by the peer as: Not Synchronized.
<b>RxHdrCodeInvalidMsgLen</b>	Number of Header errors received by the peer as: Invalid Message Length.
<b>RxHdrCodeInvalidMsgType</b>	Number of Header errors received by the peer as: Invalid Message Type.
<b>RxOpCodeBadVer</b>	Number of Open errors received by the peer as: Bad Version.
<b>RxOpCodeBadAs</b>	Number of Open errors received by the peer as: Bad AS.
<b>RxOpCodeBadRtID</b>	Number of Open errors received by the peer as: Bad BGP ID.
<b>RxOpCodeUnsuppOption</b>	Number of Open errors received by the peer as: Unsupported Options.
<b>RxOpCodeAuthFail</b>	Number of Open errors received by the peer as: Authorization Failures.
<b>RxOpCodeBadHold</b>	Number of Open errors received by the peer as: Bad Hold Value.
<b>RxUpdCodeMalformedAttrList</b>	Number of Update errors received by the peer as: Malformed Attr List.
<b>RxUpdCodeWelknownAttrUnrecog</b>	Number of Update errors received by the peer as: Wellknown Attr Unrecog.
<b>RxUpdCodeWelknownAttrMiss</b>	Number of Update errors received by the peer as: Wellknown Attr Missing.
<b>RxUpdCodeAttrFlagError</b>	Number of Update errors received by the peer as: Attr Flag Error.
<b>RxUpdCodeAttrLenError</b>	Number of Update errors received by the peer as: Attr Length Error.

<b>RxUpdCodeBadORIGINAttr</b>	Number of Update errors received by the peer as: Attr Flag Error.
<b>RxUpdCodeASRoutingLoop</b>	Number of Update errors received by the peer as: AS Routing Loop.
<b>RxUpdCodeBadNHAttr</b>	Number of Update errors received by the peer as: Bad Next-Hop Attr.
<b>RxUpdCodeOptionalAttrError</b>	Number of Update errors received by the peer as: Optional Attr Error.
<b>RxUpdCodeBadNetworkField</b>	Number of Update errors received by the peer as: Bad Network Field.
<b>RxUpdCodeMalformedASPath</b>	Number of Update errors received by the peer as: Malformed AS PATH.

---

## Viewing BGP peer transmit statistics

### About this task

View BGP peer transmit statistics to manage network performance.

### Procedure

1. In the navigation tree, open 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
<b>TxMsgs</b>	Number of messages transmitted by the peer.
<b>TxOpens</b>	Number of Opens transmitted by the peer.



Name	Description
<b>TxKeeps</b>	Number of Keepalive messages transmitted by the peer.
<b>TxUpdates</b>	Number of Updates transmitted by the peer.
<b>TxNotifys</b>	Number of Notifications transmitted by the peer.
<b>TxRoutes</b>	Number of network layer reachability information (NLRI) messages transmitted by the peer.
<b>TxECodeHeader</b>	Number of Header errors transmitted by the peer.
<b>TxECodeOpen</b>	Number of Open errors transmitted by the peer.
<b>TxECodeUpdate</b>	Number of Update errors transmitted by the peer.
<b>TxECodeHoldtimer</b>	Number of Hold errors transmitted by the peer.
<b>TxECodeFSM</b>	Number of FSM errors transmitted by the peer.
<b>TxECodeCease</b>	Number of Cease errors transmitted by the peer.
<b>TxHdrCodeNoSync</b>	Number of Header errors transmitted by the peer as: Not Synchronized.
<b>TxHdrCodeInvalidMsgLen</b>	Number of Header errors transmitted by the peer as: Invalid Message Length.
<b>TxHdrCodeInvalidMsgType</b>	Number of Header errors transmitted by the peer as: Invalid Message Type.
<b>TxOpCodeBadVer</b>	Number of Open errors transmitted by the peer as: Bad Version.
<b>TxOpCodeBadAs</b>	Number of Open errors transmitted by the peer as: Bad AS.
<b>TxOpCodeBadRtID</b>	Number of Open errors transmitted by the peer as: Bad BGP ID.
<b>TxOpCodeUnsuppOption</b>	Number of Open errors transmitted by the peer as: Unsupported Options.
<b>TxOpCodeAuthFail</b>	Number of Open errors transmitted by the peer as: Authorization Failures.

<b>TxOpCodeBadHold</b>	Number of Open errors transmitted by the peer as: Bad Hold Value.
<b>TxUpdCodeMalformedAttrList</b>	Number of Update errors transmitted by the peer as: Malformed Attr List.
<b>TxUpdCodeWelknownAttrUnrecog</b>	Number of Update errors transmitted by the peer as: Wellknown Attr Unrecog.
<b>TxUpdCodeWelknownAttrMiss</b>	Number of Update errors transmitted by the peer as: Wellknown Attr Missing.
<b>TxUpdCodeAttrFlagError</b>	Number of Update errors transmitted by the peer as: Attr Flag Error.
<b>TxUpdCodeAttrLenError</b>	Number of Update errors transmitted by the peer as: Attr Length Error.
<b>TxUpdCodeBadORIGINAttr</b>	Number of Update errors transmitted by the peer as: Attr Flag Error.
<b>TxUpdCodeASRoutingLoop</b>	Number of Update errors transmitted by the peer as: AS Routing Loop.
<b>TxUpdCodeBadNHAttr</b>	Number of Update errors transmitted by the peer as: Bad Next-Hop Attr.
<b>TxUpdCodeOptionalAttrError</b>	Number of Update errors transmitted by the peer as: Optional Attr Error.
<b>TxUpdCodeBadNetworkField</b>	Number of Update errors transmitted by the peer as: Bad Network Field.
<b>TxUpdCodeMalformedASPath</b>	Number of Update errors transmitted by the peer as: Malformed AS Path.

---

## Viewing statistics for a VRF

### About this task

View VRF statistics to ensure the instance is performing as expected.

### Procedure

1. In the navigation tree, open 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 use the **Stats** tab.

Name	Description
<b>RouteEntries</b>	Specifies the number of routes for this VRF.
<b>FIBEntries</b>	Specifies the number of Forwarding Information Base (FIB) entries for this VRF.

---

## Viewing EAPoL Authenticator statistics

### About this task

Use EAPoL Authenticator statistics to display the Authenticator Port Access Entity (PAE) statistics for each selected port.

### Procedure

1. In 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 tree, open 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
<b>FramesRx</b>	Displays the number of valid EAPoL frames of any type received by this Authenticator.

Name	Description
<b>FramesTx</b>	Displays the number of EAPoL frame types of any type transmitted by this Authenticator.
<b>StartFramesRx</b>	Displays the number of EAPoL start frames received by this Authenticator.
<b>LogoffFramesRx</b>	Displays the number of EAPoL Logoff frames received by this Authenticator.
<b>RespIdFramesRx</b>	Displays the number of EAPoL Resp/Id frames received by this Authenticator.
<b>RespFramesRx</b>	Displays the number of valid EAP Response frames (Other than Resp/Id frames) received by this Authenticator.
<b>ReqIdFramesTx</b>	Displays the number of EAPoL Req/Id frames transmitted by this Authenticator.
<b>ReqFramesTx</b>	Displays the number of EAP Req/Id frames (other than Rq/Id frames) transmitted by this Authenticator.
<b>EapLengthErrorFramesRx</b>	Displays the number of EAPoL frames received by this Authenticator in which the Packet Body Length field is invalid.
<b>InvalidEapolFramesRx</b>	Displays the number of EAPoL frames received by this Authenticator in which the frame type is not recognized.

---

## Viewing EAPoL diagnostic statistics

### About this task

Use EAPoL diagnostic statistics to display the Authenticator PAE diagnostic statistics for each selected port.

### Procedure

- In 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.
- Perform one of the following steps:
  - In the navigation tree, open the following folders: **Configuration > Graph**, and then click **Port**.

- 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
<b>EntersConnecting</b>	Counts the number of times that the Authenticator PAE state machine transitions to the Connecting state from any other state.
<b>EapLogoffsWhileConnecting</b>	Counts the number of times that the Authenticator PAE state machine transitions from Connected to Disconnected as a result of receiving an EAPoL-Logoff message.
<b>EntersAuthenticating</b>	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.
<b>AuthSuccessWhile Authenticating</b>	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.
<b>AuthTimeoutsWhile Authenticating</b>	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.

Name	Description
<b>AuthFailWhileAuthenticating</b>	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.
<b>AuthReauthsWhile Authenticating</b>	Counts the number of times that the Authenticator PAE state machine transitions from Authenticating to Aborting as a result of a reauthentication request.
<b>AuthEapStartsWhileAuthenticating</b>	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.
<b>AuthEapLogoffWhileAuthenticating</b>	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.
<b>AuthReauthsWhile Authenticated</b>	Counts the number of times that the Authenticator PAE state machine transitions from Authenticated to Connecting as a result of a reauthentication request.
<b>AuthEapStartsWhileAuthenticated</b>	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.
<b>AuthEapLogoffWhileAuthenticated</b>	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.
<b>BackendResponses</b>	Counts the number of times that the Backend Authentication state machine sends an Initial-Access request packet to the Authentication server.

Name	Description
<b>BackendAccessChallenges</b>	Counts the number of times that the Backend Authentication state machine receives an Initial-Access challenge packet from the Authentication server.
<b>BackendOtherRequestsToSupplicant</b>	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.
<b>BackendNonNakResponsesFromSupplicant</b>	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.
<b>BackendAuthSuccesses</b>	Counts the number of times that the Backend Authentication state machine receives an EAP-success message from the Authentication server.
<b>BackendAuthFails</b>	Counts the number of times that the Backend Authentication state machine receives an EAP-failure message from the Authentication server.

---

## Viewing EAPoL session statistics

### About this task

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. In 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 step:

- In the navigation tree, open the following folders: **Configuration > Graph**, and then click **Port**.
- 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
<b>SessionOctetsRx</b>	Displays the number of octets received in user data frames on this port during the session.
<b>SessionOctetsTx</b>	Displays the number of octets transmitted in user data frames on this port during the session.
<b>SessionFramesRx</b>	Displays the number of user data frames received on this port during the session.
<b>SessionFramesTx</b>	Displays the number of user data frames transmitted on this port during the session.

---

## Showing the Authenticator session statistics

### About this task

Use Authenticator Session Statistics to display the session statistics objects for the Authenticator PAE associated with each port.

### Procedure

1. In the navigation tree, open 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
<b>PortNumber</b>	Indicates the Port number associated with this Port.
<b>SessionId</b>	Specifies a unique identifier for the session, in the form of a printable ASCII string of at least three characters.
<b>SessionAuthenticMethod</b>	Indicates the authentication method used to establish the session.
<b>SessionTime</b>	Indicates the duration of the session in seconds.
<b>SessionTerminateCause</b>	Indicates the reason for the session termination.
<b>SessionUserName</b>	Indicates the User-Name that represents the identity of the Supplicant PAE.
<b>LastEapolFrameVersion</b>	Indicates the protocol version number carried in the most recently received EAPOL frame.
<b>LastEapolFrameSource</b>	Indicates the source MAC address carried in the most recently received EAPOL frame.

---

## Showing RADIUS server statistics

### About this task


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 tree, open 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
<b>AddressType</b>	Specifies either an IPv4 or an IPv6 address. RADIUS supports IPv4 and IPv6 addresses.
<b>Address</b>	The IP address of the RADIUS server.
<b>Used by</b>	Identifies the client.
<b>AccessRequests</b>	Number of access-response packets sent to the server; does not include retransmissions.
<b>AccessAccepts</b>	Number of access-accept packets, valid or invalid, received from the server.
<b>AccessRejects</b>	Number of access-reject packets, valid or invalid, received from the server.
<b>BadResponses</b>	Number of invalid access-response packets received from the server.
<b>PendingRequests</b>	Access-request packets sent to the server that have not yet received a response or that have timed out.
<b>ClientRetries</b>	Number of authentication retransmissions to the server.
<b>AcctOnRequests</b>	Number of accounting On requests sent to the server.
<b>AcctOffRequests</b>	Number of accounting Off requests sent to the server.
<b>AcctStartRequests</b>	Number of accounting Start requests sent to the server.
<b>AcctStopRequests</b>	Number of accounting Stop requests sent to the server.
<b>AcctInterimRequests</b>	Number of Accounting Interim requests sent to the server.   <b>Important:</b> The AcctInterimRequests counter increments only if you select AcctIncludeCli from the RADIUS Global tab.
<b>AcctBadResponses</b>	Number of Invalid Responses discarded from the server.
<b>AcctPendingRequests</b>	Number of requests waiting to be sent to the server.
<b>AcctClientRetries</b>	Number of retries made to this server.

---

## Showing SNMP statistics

### About this task

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 tree, open 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
<b>OutTooBigs</b>	The number of SNMP PDUs that the SNMP protocol entity generated and for which the value of the error-status field is tooBig.
<b>OutNoSuchNames</b>	The number of SNMP PDUs that the SNMP protocol entity generated and for which the value of the error-status is noSuchName.
<b>OutBadValues</b>	The number of SNMP PDUs that SNMP protocol entity generated and for which the value of the error-status field is badValue.
<b>OutGenErrors</b>	The number of SNMP PDUs that the SNMP protocol entity generated and for which the value of the error-status field is genErr.
<b>InBadVersions</b>	The number of SNMP messages that were delivered to the SNMP protocol entity and were for an unsupported SNMP version.
<b>InBadCommunity Names</b>	The number of SNMP messages delivered to the SNMP protocol entity that used an SNMP community name not known to the entity.

Name	Description
<b>InBadCommunity Users</b>	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.
<b>InASNParseErrs</b>	The number of ASN.1 or BER errors encountered by the SNMP protocol entity when decoding received SNMP messages.
<b>InTooBigs</b>	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.
<b>InNoSuchNames</b>	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.
<b>InBadValues</b>	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.
<b>InReadOnlys</b>	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.
<b>InGenErrors</b>	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

### About this task

View PCAP statistics to manage network performance.

### Procedure

1. In the navigation tree, open 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
<b>ResetStat</b>	Resets statistics when selected.
<b>PacketCapacityCount</b>	The packet capacity count.
<b>NumberOfPacketsReceived</b>	The number of packets received in the PCAP engine.
<b>NumberOfPacketsAccumulated</b>	The number of packets captured in the PCAP engine.
<b>NumberOfPacketsDroppedInPcapEngine</b>	The number of packets dropped in the PCAP engine by filters.
<b>NumberOfPacketsDroppedInHardware</b>	The number of packets dropped in hardware.

---

## Enabling RMON statistics

### About this task

Enable Ethernet statistics collection for RMON.

### Procedure

1. In the navigation tree, open 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
<b>Index</b>	Uniquely identifies an entry in the Ethernet Statistics table. The default is 1.
<b>Port</b>	Identifies the source of the data that this etherStats entry is configured to analyze.
<b>Owner</b>	Specifies the entity that configured this entry and therefore uses the assigned resources.

---

## Viewing RMON statistics

### Before you begin

- You must enable RMON statistics collection.

### About this task

Use the following procedure to view RMON statistics for each port.

### Procedure

1. In the Device Physical View, select a port.
  2. In the navigation tree, open 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.

Name	Description
<b>Octets</b>	<p>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 = ..... 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>
<b>Pkts</b>	The number of packets (including bad packets, broadcast packets, and multicast packets) received.
<b>BroadcastPkts</b>	The number of good packets received that were directed to the broadcast address. This number does not include multicast packets.
<b>MulticastPkts</b>	The number of good packets received that were directed to a multicast address. This number does not include packets directed to the broadcast address.
<b>CRCAAlignErrors</b>	The 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).
<b>UndersizePkts</b>	The number of packets received that were less than 64 octets long (excluding framing bits, but including FCS octets) and were otherwise well formed.
<b>OversizePkts</b>	The number of packets received that were longer than 1518 octets (excluding framing bits, but including FCS octets) and were otherwise well formed.
<b>Fragments</b>	The 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).

Name	Description
	It is entirely normal for Fragments to increment because it counts both runs (which are normal occurrences due to collisions) and noise hits.
<b>Collisions</b>	<p>The best estimate of the number of collisions on this Ethernet segment. 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 are transmitting 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 much smaller role when considering 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 only detect collisions when it is transmitting. Thus, probes placed on a station and a repeater reports the same number of collisions.</p> <p>An RMON probe inside a repeater reports 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>

---

## Enabling multicast routing process statistics

### About this task

Enable the collection and display of multicast routing process statistics. These statistics are not related to the interface (port) statistics. Rather, the statistics are displayed based on multicast group classification. By default, mroute statistics collection is disabled.

### Procedure

1. In the Navigation tree, open the following folders: **Configuration > IP**.
  2. Click **Multicast**.
  3. Click the **Globals** tab.
  4. Select **MRouteStatsEnabled**.  
Clear this check box to disable the collection of multicast routing statistics.
  5. Click **Apply**.
-



---

## Viewing multicast routing process statistics

### Before you begin

- You must enable the collection of multicast statistics.

### About this task

These statistics are not related to the interface (port) statistics. Rather, the system displays the statistics based on multicast group classification.

### Procedure

1. In the navigation tree, 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**.mroute statisticsresetAbsoluteValuereset mroute
- 

---

## Stats field descriptions

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

Name	Description
<b>SourceCount</b>	Specifies the source number that corresponds to the associated group IP address in the whole multicast route records.
<b>IngressPkts</b>	Specifies the number of normally forwarded packets for the associated group IP address.
<b>IngressBytes</b>	Specifies the number of normally forwarded bytes for the associated group IP address.
<b>AverageSizePerPkt</b>	Specifies the average packet length for the associated group IP address. This information indicates only the forward packet

Name	Description
	length and is calculated using the following formula: forward packet/forward byte.
<b>DropPkts</b>	Specifies the number of dropped packets for the associated group IP address.
<b>DropBytes</b>	Specifies the number of dropped bytes for the associated VRF and group IP address.
<b>PktsPerSecond</b>	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

### About this task

View the hashing statistics to view total hash overflows.

### Procedure

1. In the navigation tree, open 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
<b>HashOverflows</b>	Shows the count of hash overflows for the slot.

Name	Description
HashDrops	Shows the count of hash drops for the slot.

---

## Viewing IPFIX exporter statistics

### About this task

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 tree, open 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.



# Chapter 14: RMON alarm variables

RMON alarm variables are divided into three categories. Each category has subcategories. The following table lists the alarm variable categories and provides a brief variable description.

**Table 41: 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.
		snmplnBadCommunityNames.0	The total 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.
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

Category	Subcategory	Variable	Definition
			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 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

Category	Subcategory	Variable	Definition
			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 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

Category	Subcategory	Variable	Definition
			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. 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



Category	Subcategory	Variable	Definition
			<p>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.</p>
		dot3StatsFrameTooLongs	<p>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.</p>
		dot3StatsInternalMacReceivesErrors	<p>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.</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 a count of transmission errors on a particular interface that are not otherwise counted.</p>

Category	Subcategory	Variable	Definition
	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 number of fragments by combining them as they are received.

Category	Subcategory	Variable	Definition
		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.

Category	Subcategory	Variable	Definition
		rcStatMltEtherExcessiveCollis	The number of times that excessive collisions are 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.
		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

Category	Subcategory	Variable	Definition
			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.
		snmplnAsnParseErrs.0	The total number of ASN.1 or BER errors encountered by the SNMP protocol entity when it decodes received SNMP messages.
		rcStgPortlnBadBpdus	The number of bad BPDUs received by this port.
		dot1dTpPortlnDiscards	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	iflnOctets	The total number of octets received on the interface, including framing characters.
		iflnMulticastPkts	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.
		iflnBroadcastPkts	The number of packets, delivered by this sublayer to a higher (sub) layer, that are addressed to a broadcast address at this sublayer.

Category	Subcategory	Variable	Definition
		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 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

Category	Subcategory	Variable	Definition
			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.
		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.
		etherStatsCRCAlignErrors	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)

Category	Subcategory	Variable	Definition
			and were otherwise well formed.
		etherStatsFragments	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). It is entirely normal for etherStatsFragments to increment because it counts both runts (which are normal occurrences due to collisions) and noise hits.
		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.



Category	Subcategory	Variable	Definition
		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.
		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.

Category	Subcategory	Variable	Definition
		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	snmplnPkts.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.
		snmplnBadVersions.0	The total number of SNMP messages delivered to the SNMP protocol entity that were intended for an unsupported SNMP version.
		snmplnBadCommunityUses.0	The total number of SNMP messages delivered to the SNMP protocol entity that represented an SNMP operation that was not allowed by the SNMP community named in the message.
		snmplnTooBigs.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.
		snmplnNoSuchNames.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.
		snmplnBadValues. 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.

Category	Subcategory	Variable	Definition
		snmplnReadOnlys.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.
		snmplnGenErrs.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.
		snmplnTotalReqVars.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.
		snmplnTotalSetVars.0	The total number of MIB objects altered successfully by the SNMP protocol entity as the result of receiving valid SNMP Set-Request PDUs.
		snmplnGetRequests.0	The total number of SNMP Get-Request PDUs accepted and processed by the SNMP protocol entity.
		snmplnGetNexts.0	The total number of SNMP Get-Next PDUs accepted and processed by the SNMP protocol entity.
		snmplnSetRequests.0	The total number of SNMP Set-Request PDUs accepted and processed by the SNMP protocol entity.
		snmplnGetResponses.0	The total number of SNMP Get-Response PDUs accepted and processed by the SNMP protocol entity.

Category	Subcategory	Variable	Definition
		snmpInTraps.0	The total number of SNMP Trap PDUs accepted and processed by the SNMP protocol entity.
		snmpOutTooBigs.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.
		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.

Category	Subcategory	Variable	Definition
		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 the local bridging function, including bridge management frames.
		dot1dTpLearnedEntryDiscards.0	The total number of Forwarding Database entries learned but

Category	Subcategory	Variable	Definition
			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.
		ospfVirtIfState	The number of times this OSPF interface.
		ospfVirtIfEvents	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.

Category	Subcategory	Variable	Definition
	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.
		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



Category	Subcategory	Variable	Definition
			interface detected by the high-count (64-bit) register.
		rcStatMltIfExtnIfHCOctets	The total number of octets transmitted from the MLT interface, including framing characters.
		rcStatMltIfExtnIfHCOUcastPkts	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.
		rcStatMltIfExtnIfHCOmulticast	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.
		rcStatMltIfExtnIfHCObroadcast	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.



# Chapter 15: Customer service

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

---

## Getting technical documentation

To download and print selected technical publications and release notes directly from the Internet, go to [www.avaya.com/support](http://www.avaya.com/support).

---

## Getting product training

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

---

## Getting help from a distributor or reseller

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

---

## Getting technical support from the Avaya Web site

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



## Glossary

<b>American Standard Code for Information Interchange (ASCII)</b>	A code for representing characters in computers. ASCII uses uppercase and lowercase alphabetic letters, numeric digits, and special symbols.
<b>Autonomous System (AS)</b>	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.
<b>Autonomous System Number (ASN)</b>	A two-byte number that is used to identify a specific AS.
<b>bit error rate (BER)</b>	The ratio of the number of bit errors to the total number of bits transmitted in a given time interval.
<b>Bootstrap Protocol (BootP)</b>	A User Datagram Protocol (UDP)/Internet Protocol (IP)-based protocol that a booting host uses to configure itself dynamically and without user supervision.
<b>Control Processor Unit High Availability (CPU-HA)</b>	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 1 second with the most recent topology data.
<b>Collecting process</b>	A process that receives flow records from one or more exporting processes. The collecting process can process or store received flow records.
<b>Collector</b>	A device that hosts one or more collecting processes.
<b>cyclic redundancy check (CRC)</b>	Ensures frame integrity is maintained during transmission. The CRC performs a computation on frame contents before transmission and on the receiving device. The system discards frames that do not pass the CRC.
<b>Data flowset</b>	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.
<b>Dynamic Random Access Memory (DRAM)</b>	A read-write random-access memory, in which the digital information is represented by charges stored on the capacitors and must be repeatedly replenished in order to retain the information.

<b>Exporting process</b>	An export process that sends flow records to one or more collecting processes. One or more metering processes generate the flow records.
<b>Flow key</b>	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).
<b>Flow record</b>	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.
<b>Flowset</b>	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.
<b>forwarding database (FDB)</b>	A database that maps a port for every MAC address. If a packet is sent to a specific MAC address, the switch refers to the forwarding database for the corresponding port number and sends the data packet through that port.
<b>Frame Check Sequence (FCS)</b>	Frames are used to send upper-layer data and ultimately the user application data from a source to a destination.
<b>graphical user interface (GUI)</b>	A graphical (rather than textual) computer interface.
<b>Internet Control Message Protocol (ICMP)</b>	A collection of error conditions and control messages exchanged by IP modules in both hosts and gateways.
<b>Internet Group Management Protocol (IGMP)</b>	A host membership protocol used to arbitrate membership in multicast services.
<b>Internet Protocol Flow Information eXport (IPFIX)</b>	An IETF standard that improves the Netflow V9 protocol. IPFIX monitors IP flows.
<b>Internet Protocol Flow Information eXport (IPFIX) device</b>	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.
<b>Internet Protocol Flow Information eXport (IPFIX) node</b>	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.

<b>Internet Protocol traffic flow or flow</b>	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.
<b>interswitch trunking (IST)</b>	A feature that uses one or more parallel point-to-point links to connect two aggregation switches. The two aggregation switches use this channel to share information and operate as a single logical switch. Only one interswitch trunk can exist on each Split Multilink Trunking (SMLT) aggregation switch.
<b>Layer 2</b>	The Data Link Layer of the OSI model. Examples of Layer 2 protocols are Ethernet and Frame Relay.
<b>Layer 3</b>	The Network Layer of the OSI model. Example of a Layer 3 protocol is Internet Protocol (IP).
<b>link-state advertisement (LSA)</b>	Packets that contain state information about directly connected links (interfaces) and adjacencies. Each Open Shortest Path First (OSPF) router generates the packets.
<b>link-state database (LSDB)</b>	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.
<b>Link Aggregation Control Protocol (LACP)</b>	A protocol that exists between two endpoints to bundle links into an aggregated link group for bandwidth increase and link redundancy.
<b>Link Aggregation Control Protocol Data Units (LACPDU)</b>	Link aggregation control protocol data unit (LACPDU) is used for exchanging information among LACP-enabled devices.
<b>Logical Link Control (LLC)</b>	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.
<b>management information base (MIB)</b>	Defines system operations and parameters used for the Simple Network Management Protocol (SNMP).
<b>media</b>	A substance that transmits data between ports; usually fiber optic cables or category 5 unshielded twisted pair (UTP) copper wires.
<b>Metering process</b>	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

multiplexing

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 given layer and that consists of protocol-control information of the given layer and possibly user data of that layer.



<b>Protocol Independent Multicast, Source Specific (PIM-SSM)</b>	Uses only shortest-path trees to provide multicast services based on subscription to a particular (source, group) channel.
<b>Protocol Independent Multicast, Sparse Mode (PIM-SM)</b>	Adds a Rendezvous Point router to avoid multicast-data flooding. Use PIM-SM when receivers for multicast data are sparsely distributed throughout the network.
<b>quality of service (QoS)</b>	Use QoS features to reserve resources in a congested network. For example, you can configure a higher priority to IP deskphones, which need a fixed bit rate, and, split the remaining bandwidth between data connections if calls in the network are important than the file transfers.
<b>Random Access Memory (RAM)</b>	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.
<b>remote login (rlogin)</b>	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.
<b>remote monitoring (RMON)</b>	A remote monitoring standard for Simple Network Management Protocol (SNMP)-based management information bases (MIB). The Internetwork Engineering Task Force (IETF) proposed the RMON standard to provide guidelines for remote monitoring of individual LAN segments.
<b>shortest path first (SPF)</b>	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.
<b>small form factor pluggable (SFP)</b>	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.
<b>small form factor pluggable plus (SFP+)</b>	SFP+ transceivers are similar to SFPs in physical appearance but SFP+ transceivers provide Ethernet at 10 gigabit per second (Gb/s).
<b>spanning tree</b>	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.

<b>Spanning Tree Group (STG)</b>	A collection of ports in one spanning tree instance.
<b>Template flowset</b>	One or more options template records in an export packet.
<b>Template record</b>	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.
<b>time-to-live (TTL)</b>	The field in a packet used to determine the valid duration for the packet; the TTL determines the packet lifetime. The system discards a packet with a TTL of zero.
<b>traffic profile</b>	The temporal properties of a traffic stream, such as rate.
<b>Trivial File Transfer Protocol (TFTP)</b>	A protocol that governs transferring files between nodes without protection against packet loss.
<b>trunk</b>	A logical group of ports that behaves like a single large port.
<b>User Datagram Protocol (UDP)</b>	In TCP/IP, a packet-level protocol built directly on the Internet Protocol layer. TCP/IP host systems use UDP for application-to-application programs.
<b>Virtual Router Redundancy Protocol (VRRP)</b>	A protocol used in static routing configurations, typically at the edge of the network. This protocol operates on multiple routers on an IP subnet and elects a primary gateway router. When the primary router fails, a backup router is quickly available to take its place.