



# Monitoring Performance for VOSS

Release 8.0 (VOSS)  
9035657 Rev 02  
February 2020

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

### **Legal Notice**

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

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

### **Trademarks**

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

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

For additional information on Extreme Networks trademarks, please see: [www.extremenetworks.com/company/legal/trademarks](http://www.extremenetworks.com/company/legal/trademarks)

### **Open Source Declarations**

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

# Contents

<b>Chapter 1: About this document</b> .....	10
Purpose.....	10
Conventions.....	10
Text Conventions.....	11
Documentation and Training.....	13
Getting Help.....	13
Providing Feedback to Us.....	14
<b>Chapter 2: New in this Document</b> .....	15
Notice about Feature Support.....	16
<b>Chapter 3: Port performance management</b> .....	17
Digital Diagnostic Monitoring.....	17
Port Performance Management Using CLI.....	17
Viewing DDI Port Information.....	17
Viewing DDI Temperature Information.....	19
Viewing DDI Voltage Information.....	20
Port Performance Management using EDM.....	20
Viewing DDI Information.....	21
<b>Chapter 4: Key Health Indicators (KHI)</b> .....	24
Key Health Indicators Using the CLI.....	24
Displaying KHI Performance Information.....	24
Displaying KHI Control Processor Information.....	33
Clearing KHI Information.....	34
Displaying KHI Fabric Extend ONA Status.....	35
Displaying KHI Fabric Extend ONA Global Information.....	36
Key Health Indicators Using EDM.....	37
Clearing KHI Statistics.....	37
Displaying KHI Port Information.....	38
<b>Chapter 5: IPFIX</b> .....	39
IPFIX Fundamentals.....	39
IPFIX Configuration Using CLI.....	41
Enabling IPFIX Globally.....	41
Displaying IPFIX Global Status.....	43
Configuring the IPFIX Aging Interval.....	43
Configuring the IPFIX Collector.....	44
Displaying IPFIX Collector Information.....	45
Configuring an IPFIX Observation Domain.....	46
Displaying IPFIX Flows.....	46
IPFIX Configuration Using EDM.....	48
Enabling IPFIX Globally.....	48

Configuring the IPFIX Collector.....	48
<b>Chapter 6: Link state change control.....</b>	<b>51</b>
Link State Change Control.....	51
Link State Change Control Using CLI.....	51
Link State Change Control Using EDM.....	53
<b>Chapter 7: Logs and traps.....</b>	<b>54</b>
Logs and Traps.....	54
Logs and Traps Fundamentals.....	54
Log Configuration Using CLI.....	64
Configuring a UNIX System Log and Syslog Host.....	64
Configuring Secure Forwarding.....	67
Installing Root Certificate for Syslog Client.....	69
Configuring Logging.....	69
Configuring the Remote Host Address for Log Transfer.....	71
Configuring System Logging.....	71
Configuring System Message Control.....	72
Extending System Message Control.....	73
Viewing Logs.....	74
Configuring CLI Logging.....	77
Configuring Email Notification.....	78
Log Configuration Using EDM.....	83
Configuring the System Log.....	83
Configuring the System Log Table.....	84
Configuring Email Notification.....	85
SNMP Trap Configuration Using CLI.....	88
Configuring an SNMP Host.....	88
Configuring an SNMP Notify Filter Table.....	89
Configuring SNMP Interfaces.....	90
Enabling SNMP Trap Logging.....	91
SNMP Trap Configuration Using EDM.....	92
Configuring an SNMP Host Target Address.....	92
Configuring Target Table Parameters.....	94
Configuring SNMP Notify Filter Profiles.....	95
Configuring SNMP Notify Filter Profile Table Parameters.....	96
Enabling Authentication Traps.....	96
Viewing the Trap Sender Table.....	97
<b>Chapter 8: MACsec performance.....</b>	<b>98</b>
MACsec Statistics.....	98
Viewing MACsec Statistics Using the CLI.....	100
Viewing MACsec Statistics.....	100
Viewing MACsec Statistics Using EDM.....	101
Viewing MACsec Interface Statistics.....	101
Viewing Secure Channel (SC) Inbound Statistics.....	102

Viewing Secure Channel (SC) Outbound Statistics.....	104
<b>Chapter 9: Remote Monitoring.....</b>	<b>106</b>
Remote Monitoring.....	106
RMON 2.....	110
RMON Configuration Using CLI.....	112
Configuring RMON.....	112
Enabling Remote Monitoring on an Interface.....	117
Displaying RMON Information.....	118
RMON Configuration Using EDM.....	125
Enabling RMON Globally.....	125
Enabling RMON on a Port or VLAN.....	126
Enabling RMON1 History.....	127
Disabling RMON1 History.....	128
Viewing RMON1 History Statistics.....	129
Creating an RMON1 Alarm.....	131
Viewing RMON1 Alarms.....	134
Deleting an RMON1 Alarm.....	134
Creating an RMON1 Event.....	134
Viewing RMON1 Events.....	135
Deleting an Event.....	136
Viewing the RMON Log.....	136
Viewing the Protocol Directory.....	137
Viewing the Data Source for Protocol Distribution Statistics.....	138
Viewing Protocol Distribution Statistics.....	139
Viewing the Host Interfaces Enabled for Monitoring.....	140
Viewing Address Mappings.....	140
Viewing the Data Source for Host Statistics.....	141
Viewing Network Host Statistics.....	142
Viewing Application Host Statistics.....	143
RMON Alarm Variables.....	144
<b>Chapter 10: sFlow.....</b>	<b>163</b>
sFlow Fundamentals.....	163
sFlow Configuration Using CLI.....	166
Configuring the agent-ip and Enabling sFlow Globally.....	166
Configuring an sFlow Collector.....	167
Configuring the Packet Sampling Rate.....	168
Configuring sFlow Maximum Header Size.....	170
Configuring the Counter Sampling Interval.....	171
Viewing sFlow Statistics.....	173
sFlow Configuration Using EDM.....	174
Enabling sFlow and Configuring the Agent IP Address.....	174
Configuring an sFlow Collector.....	175
Configuring the Packet Samples and Counter Samples.....	176

Displaying sFlow Statistics.....	178
<b>Chapter 11: Application Telemetry Fundamentals.....</b>	<b>180</b>
How Application Telemetry Works.....	180
Common Elements Between sFlow and Application Telemetry.....	181
Operational Considerations and Restrictions.....	182
Configuration Overview.....	185
Application Telemetry Configuration Using CLI.....	186
Configuring the Agent IP Address.....	186
Configuring an Analytics Engine and Enabling Application Telemetry Globally.....	187
Viewing Application Telemetry Counters.....	189
Clearing Application Telemetry Counters.....	190
Application Telemetry Configuration Using EDM.....	190
Enabling sFlow and Configuring the Agent IP Address.....	190
Configuring an sFlow Collector.....	191
Enabling Application Telemetry Globally.....	192
Viewing Application Telemetry Counters.....	193
Clearing Application Telemetry Counters.....	193
Viewing Application Telemetry Status.....	194
<b>Chapter 12: Statistics.....</b>	<b>195</b>
Viewing Statistics Using CLI.....	195
Viewing TCP Statistics.....	195
Viewing Port Routing Statistics.....	196
Displaying Bridging Statistics for Specific Ports.....	197
Displaying DHCP-relay Statistics for Specific Ports.....	199
Displaying DHCP-relay Statistics for all Interfaces.....	200
Displaying LACP Statistics for Specific Ports.....	202
Displaying VLACP Statistics for Specific Ports.....	204
Displaying RMON Statistics for Specific Ports.....	206
Displaying Detailed Statistics for Ports.....	208
Displaying IS-IS Statistics and Counters.....	209
Clearing ACL Statistics.....	211
Viewing ACE Statistics.....	212
Viewing MSTP Statistics.....	214
Viewing RSTP Statistics.....	215
Viewing RSTP Port Statistics.....	216
Viewing MLT Statistics.....	218
Viewing vIST Statistics.....	219
Showing RADIUS Server Statistics.....	222
Viewing RMON Statistics.....	224
Showing OSPF Error Statistics on a Port.....	224
Viewing OSPF Interface Statistics.....	225
Viewing OSPF Range Statistics.....	227
Clearing IP OSPF Statistics.....	229

Viewing Basic OSPF Statistics for a Port.....	229
Showing Extended OSPF Statistics.....	230
Viewing Ingress Port-rate Limit Statistics.....	232
Viewing Ingress Policer Statistics.....	232
Viewing the Management Port Statistics.....	233
Viewing IP VRRPv3 Statistics.....	234
Clearing IPv4 MSDP Statistics.....	235
Viewing IPv4 ICMP Statistics.....	236
Clearing IPv6 Statistics.....	237
Viewing IPv6 ICMP Statistics.....	238
Viewing IPv6 DHCP Relay Statistics.....	239
Viewing IPv6 OSPF Statistics.....	240
Clearing IPv6 OSPF Statistics.....	241
Viewing IPv6 Statistics on an Interface.....	242
Displaying IPsec Statistics.....	243
Viewing IPv6 VRRP Statistics.....	250
Showing the EAPoL Status of the Device.....	253
Showing EAPoL Authenticator Statistics.....	253
Viewing EAPoL Session Statistics.....	255
Viewing Non-EAPoL MAC Information.....	256
Viewing Port EAPoL Operation Statistics.....	257
Viewing IP Multicast Threshold Exceeded Statistics.....	259
Viewing NTP Statistics.....	259
Viewing Segmented Management Instance Statistics.....	260
Viewing Statistics Using EDM.....	261
Graphing Chassis Statistics.....	261
Graphing Port Statistics.....	262
Viewing Chassis System Statistics.....	263
Viewing Chassis SNMP Statistics.....	263
Viewing Chassis IP Statistics.....	265
Viewing Chassis ICMP In Statistics.....	267
Viewing Chassis ICMP Out Statistics.....	267
Viewing Chassis TCP Statistics.....	268
Viewing Chassis UDP Statistics.....	269
Viewing Port Interface Statistics.....	270
Viewing Port Ethernet Errors Statistics.....	272
Viewing Port Bridging Statistics.....	274
Viewing Port Spanning Tree Statistics.....	275
Viewing Port Routing Statistics.....	275
Viewing DHCP Statistics for an Interface.....	276
Graphing DHCP Statistics for a Port.....	276
Viewing DHCP Statistics for a Port.....	277
Graphing DHCP Statistics for a VLAN.....	277

Displaying DHCP-relay Statistics for Option 82.....	278
Viewing Port OSPF Statistics.....	279
Viewing LACP Port Statistics.....	280
Viewing Port Policer Statistics.....	281
Displaying File Statistics.....	282
Viewing ACE Port Statistics.....	282
Viewing ACL Statistics.....	283
Clearing ACL Statistics.....	284
Viewing VLAN and Spanning Tree CIST Statistics.....	284
Viewing VLAN and Spanning Tree MSTI Statistics.....	285
Viewing VRRP Interface Statistics.....	286
Viewing VRRP Statistics.....	287
Viewing SMLT Statistics.....	288
Viewing RSTP Status Statistics.....	289
Viewing MLT Interface Statistics.....	290
Viewing MLT Ethernet Error Statistics.....	291
Viewing RIP Statistics.....	293
Viewing OSPF Chassis Statistics.....	294
Graphing OSPF Statistics for a VLAN.....	295
Graphing OSPF Statistics for a Port.....	296
Viewing BGP Global Stats.....	297
Viewing Statistics for a VRF.....	301
Showing RADIUS Server Statistics.....	302
Showing SNMP Statistics.....	303
Enabling RMON Statistics.....	304
Viewing RMON Statistics.....	305
Displaying IS-IS System Statistics.....	307
Displaying IS-IS Interface Counters.....	308
Displaying IS-IS Interface Control Packets.....	309
Graphing IS-IS Interface Counters.....	309
Graphing IS-IS Interface Sending Control Packet Statistics.....	310
Graphing IS-IS Interface Receiving Control Packet Statistics.....	311
Graphing Stat Rate Limit Statistics for a Port.....	312
Viewing IPv6 Statistics for an Interface.....	312
Viewing ICMP Statistics.....	315
Viewing IPv6 OSPF Statistics.....	318
Viewing IPv6 VRRP Statistics.....	318
Viewing IPv6 VRRP Statistics for an Interface.....	319
Configuring IPv6 VRRP Statistics.....	321
Viewing IP VRRPv3 Statistics.....	321
Graphing IPv6 VRRP Statistics.....	322
Graphing IP VRRPv3 Statistics.....	322
Viewing IPv6 DHCP Relay Statistics for a Port.....	325



- Displaying IPsec Interface Statistics..... 325
- Graphing IPsec Interface Statistics..... 328
- Displaying Switch Level Statistics for IPsec-Enabled Interfaces..... 329
- Viewing EAPoL Authenticator Statistics..... 331
- Viewing Multihost Status Information..... 332
- Viewing EAP Session Statistics..... 333
- Viewing NEAP MAC Information..... 333
- Viewing Secure Channel (SC) Outbound Statistics..... 334
- Viewing Secure Channel (SC) Inbound Statistics..... 335
- Viewing MACsec Interface Statistics..... 336
- Viewing Segmented Management Instance Statistics..... 338
- Displaying RADIUS CoA Reauthenticate Statistics..... 338
- Displaying RADIUS CoA Disconnect Statistics..... 340
- Displaying RADIUS CoA Statistics..... 342
- Glossary..... 344**

# Chapter 1: About this document

---

## Purpose

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

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

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

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

This document also provides information about how to prevent faults and improve the performance of the switch. This includes procedures for link state change, key health indicators, and logs and traps.

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

---







## Conventions

This section discusses the conventions used in this guide.

## Text Conventions

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

**Table 1: Notice Icons**

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

**Table 2: Text Conventions**

Convention	Description
Angle brackets ( < > )	<p>Angle brackets ( &lt; &gt; ) indicate that you choose the text to enter based on the description inside the brackets. Do not type the brackets when you enter the command.</p> <p>If the command syntax is <code>cfm maintenance-domain maintenance-level &lt;0-7&gt;</code>, you can enter <code>cfm maintenance-domain maintenance-level 4</code>.</p>
<b>Bold text</b>	<p>Bold text indicates the GUI object name you must act upon.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• Click <b>OK</b>.</li> <li>• On the <b>Tools</b> menu, choose <b>Options</b>.</li> </ul>
Braces ( { } )	<p>Braces ( { } ) indicate required elements in syntax descriptions. Do not type the braces when you enter the command.</p> <p>For example, if the command syntax is <code>ip address {A.B.C.D}</code>, you must enter the IP address in dotted, decimal notation.</p>

*Table continues...*

Convention	Description
Brackets ( [ ] )	<p>Brackets ( [ ] ) indicate optional elements in syntax descriptions. Do not type the brackets when you enter the command.</p> <p>For example, if the command syntax is <code>show clock [detail]</code>, you can enter either <code>show clock</code> or <code>show clock detail</code>.</p>
Ellipses ( ... )	<p>An ellipsis ( ... ) indicates that you repeat the last element of the command as needed.</p> <p>For example, if the command syntax is <code>ethernet/2/1 [ &lt;parameter&gt; &lt;value&gt; ]...</code>, you enter <code>ethernet/2/1</code> and as many parameter-value pairs as you need.</p>
<i>Italic Text</i>	<p>Italics emphasize a point or denote new terms at the place where they are defined in the text. Italics are also used when referring to publication titles that are not active links.</p>
Plain Courier Text	<p>Plain Courier text indicates command names, options, and text that you must enter. Plain Courier text also indicates command syntax and system output, for example, prompts and system messages.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• <code>show ip route</code></li> <li>• <code>Error: Invalid command syntax</code>  <code>[Failed][2013-03-22 13:37:03.303</code>  <code>-04:00]</code></li> </ul>
Separator ( > )	<p>A greater than sign ( &gt; ) shows separation in menu paths.</p> <p>For example, in the Navigation tree, expand the <b>Configuration &gt; Edit</b> folders.</p>
Vertical Line (   )	<p>A vertical line (   ) separates choices for command keywords and arguments. Enter only one choice. Do not type the vertical line when you enter the command.</p> <p>For example, if the command syntax is <code>access-policy by-mac action { allow   deny }</code>, you enter either <code>access-policy by-mac action allow</code> or <code>access-policy by-mac action deny</code>, but not both.</p>

---

## Documentation and Training

To find Extreme Networks product guides, visit our documentation pages at:

Current Product Documentation	<a href="http://www.extremenetworks.com/documentation/">www.extremenetworks.com/documentation/</a>
Archived Documentation (for earlier versions and legacy products)	<a href="http://www.extremenetworks.com/support/documentation-archives/">www.extremenetworks.com/support/documentation-archives/</a>
Release Notes	<a href="http://www.extremenetworks.com/support/release-notes">www.extremenetworks.com/support/release-notes</a>
Hardware/Software Compatibility Matrices	<a href="https://www.extremenetworks.com/support/compatibility-matrices/">https://www.extremenetworks.com/support/compatibility-matrices/</a>
White papers, data sheets, case studies, and other product resources	<a href="https://www.extremenetworks.com/resources/">https://www.extremenetworks.com/resources/</a>

### Training

Extreme Networks offers product training courses, both online and in person, as well as specialized certifications. For more information, visit [www.extremenetworks.com/education/](http://www.extremenetworks.com/education/).

---

## Getting Help

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

#### [Extreme Portal](#)

Search the GTAC (Global Technical Assistance Center) knowledge base, manage support cases and service contracts, download software, and obtain product licensing, training, and certifications.

#### [The Hub](#)

A forum for Extreme Networks customers to connect with one another, answer questions, and share ideas and feedback. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.

#### [Call GTAC](#)

For immediate support: 1-800-998-2408 (toll-free in U.S. and Canada) or +1 408-579-2826. For the support phone number in your country, visit: [www.extremenetworks.com/support/contact](http://www.extremenetworks.com/support/contact)

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

- Your Extreme Networks service contract number and/or serial numbers for all involved Extreme Networks products
- A description of the failure
- A description of any action(s) already taken to resolve the problem
- A description of your network environment (such as layout, cable type, other relevant environmental information)

- Network load at the time of trouble (if known)
- The device history (for example, if you have returned the device before, or if this is a recurring problem)
- Any related RMA (Return Material Authorization) numbers

### Subscribing to Service Notifications

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

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

 **Note:**

You can modify your product selections or unsubscribe at any time.

4. Click **Submit**.

---

## Providing Feedback to Us

Quality is our first concern at Extreme Networks, and we have made every effort to ensure the accuracy and completeness of this document. We are always striving to improve our documentation and help you work better, so we want to hear from you! We welcome all feedback but especially want to know about:

- Content errors or confusing or conflicting information.
- Ideas for improvements to our documentation so you can find the information you need faster.
- Broken links or usability issues.

If you would like to provide feedback to the Extreme Networks Information Development team, you can do so in two ways:

- Use our short online feedback form at <https://www.extremenetworks.com/documentation-feedback/>.
- Email us at [documentation@extremenetworks.com](mailto:documentation@extremenetworks.com).

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

# Chapter 2: New in this Document

The following sections detail what is new in this document.

## Application Telemetry

Application Telemetry is enhanced to support interaction with Analytics Engine. Application Telemetry and sFlow collector configuration has been enhanced to include reachability status.

For more information, see:

- [Application Telemetry Fundamentals](#) on page 180
- [Operational Considerations and Restrictions](#) on page 182
- [Configuring an sFlow Collector](#) on page 167
- [Configuring an Analytics Engine and Enabling Application Telemetry Globally](#) on page 187
- [Viewing Application Telemetry Status](#) on page 194
- [Configuring an sFlow Collector](#) on page 175

## IPFIX

Internet Protocol Flow Information eXport (IPFIX) is an Internet Engineering Task Force (IETF) standard to export IP flow information. IPFIX monitors IPv4 traffic flows and exports flow information, such as source and destination IP address or the number of packets belonging to a specific flow, to a collector. The collector performs network analytics to provide statistical data.

For more information, see:

- [IPFIX Fundamentals](#) on page 39
- [IPFIX Configuration Using CLI](#) on page 41
- [IPFIX Configuration Using EDM](#) on page 48

## New IPv6 Virtualized features

IPv6 Virtualization adds the ability to support the following IPv6 features in Virtual Routing and Forwarding (VRF) and Layer 3 Virtual Services Networks (L3 VSNs):

- Open Shortest Path First for IPv6 (OSPFv3)
- IPv6 Border Gateway Protocol (IPv6 BGP)
- IPv6 route redistribution enhancements

For more information, see [Viewing IPv6 OSPF Statistics](#) on page 240.

---

## Notice about Feature Support

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

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

For information about feature support, see [Release Notes for VOSS](#).

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



# Chapter 3: Port performance management

---

## 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.

The following optical transceivers support DDM:

- 1 Gbps Small Form Factor Pluggable (SFP)
- 10 Gbps Small Form Factor Pluggable plus (SFP+)
- 40 Gbps Quad Small Form Factor Pluggable plus (QSFP+)
- 100 Gbps Quad Small Form Factor Pluggable 28 (QSFP28)

**\* Note:**

Not all hardware platforms support each form factor. For more information on supported form factors, see your hardware documentation.

Digital Diagnostic Interface (DDI) is an interface that supports DDM. These devices provide real-time monitoring of individual DDI transceivers. 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 optical transceivers, see [Extreme Networks Pluggable Transceivers Installation Guide](#)

---

## Port Performance Management Using CLI

This section contains procedures to monitor individual DDI transceivers using the CLI.

---

### Viewing DDI Port Information

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

### About this task

This command displays information for DDI transceivers.

**\* Note:**

Different hardware platforms can support different form factors. For more information, see the hardware documentation for your platform.

### Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View basic manufacturing information and characteristics:

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

3. View configuration information:

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

4. View detailed manufacturing information and characteristics:

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

### Example

```
Switch:1#show pluggable-optical-modules config
=====
Pluggable Optical Module Global Configuration
=====
          ddm-monitor : disabled
 ddm-monitor-interval : 5
          ddm-traps-send : enabled
 ddm-alarm-portdown   : disabled
```

### Variable Definitions

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

Variable	Value
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

## Viewing DDI Temperature Information

### About this task

This command displays information for DDI transceivers.

#### \* Note:

Different hardware platforms can support different form factors. For more information, see the hardware documentation for your platform.

### Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View temperatures:

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

### Example

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

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

PORT NUM	LOW_ALARM THRESHOLD	LOW_WARN THRESHOLD	ACTUAL VALUE	HIGH_WARN THRESHOLD	HIGH_ALARM THRESHOLD	THRESHOLD STATUS
1/2	7.0	1.1250	65.2539	0.0	3.0156	Low Alarm
1/3	7.0	1.1250	65.2539	0.0	3.0156	Low Alarm
1/9	7.0625	0.0	65.2539	0.0	3.0156	Low Alarm
1/15	7.0625	0.0	65.2539	0.0	3.0156	Low Alarm
2/1	7.0625	0.0	65.2539	0.0	3.0156	Low Alarm
2/17	7.0625	0.0	65.2539	0.0	3.0156	Low Alarm
2/40	7.0625	0.0	65.2539	0.0	3.0156	Low Alarm

## Variable Definitions

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

Variable	Value
{slot/port[/sub-port]}[-slot/port[/sub-port]][,...]	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

## Viewing DDI Voltage Information

### About this task

This command displays information for DDI transceivers.

### Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View voltages:

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

### Example

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

```
=====
                          Pluggable Optical Module Voltage(V)
=====
```

PORT NUM	LOW_ALARM THRESHOLD	LOW_WARN THRESHOLD	ACTUAL VALUE	HIGH_WARN THRESHOLD	HIGH_ALARM THRESHOLD	THRESHOLD STATUS
1/2	0.1281	0.0	1.2596	0.5376	1.6396	Normal
1/3	0.0001	0.0	1.2596	0.3072	1.6396	Normal
1/9	0.0006	0.0	1.2596	2.6368	0.0	Normal
1/15	0.0006	0.0	1.2596	2.6368	0.0	Normal
2/1	0.0006	0.0	1.2596	2.6368	0.0	Normal
2/17	0.0006	0.0	1.2596	2.6368	0.0	Normal
2/40	0.0006	0.0	1.2596	2.6368	0.0	Normal

## Variable Definitions

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

Variable	Value
{slot/port[/sub-port]}[-slot/port[/sub-port]] [,...]]	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

## Port Performance Management using EDM

This section contains procedures to monitor individual DDI transceivers using EDM.

## Viewing DDI Information

### About this task

You can view DDI information, for example, port information, temperature, and voltages for DDI transceivers.

#### Note:

Different hardware platforms can support different types of transceivers. For more information, see your hardware documentation.

### Procedure

1. In the Physical Device view, select a port.
2. In the navigation pane, expand the **Configuration > Edit > Port** folders.
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>ConnectorType</b>	Indicates the type of connector.
<b>SupportsDDM</b>	Indicates if the transceiver supports DDM.
<b>DdmStatusMask</b>	Indicates the DDM status. A value other than ddm-ok represents a specific error.
<b>CLEI</b>	Indicates the Telcordia register assignment CLEI code.
<b>VendorName</b>	Indicates the name of the manufacturer.
<b>VendorPartNumber</b>	Indicates the part number for the transceiver.
<b>VendorRevNumber</b>	Indicates the manufacturer revision level for the transceiver.
<b>VendorSN</b>	Indicates the manufacturer serial number for the transceiver.
<b>VendorDateCode</b>	Indicates the manufacturer date code for the transceiver.
<b>Wavelength</b>	Indicates the wavelength in nm. This field is valid for optical transceivers only.
<b>Calibration</b>	Indicates if the calibration is internal or external.
<b>PowerMeasure</b>	Indicates Rx power measurement as average or OMA.
<b>Aux1Monitoring</b>	Indicates if auxiliary monitoring is implemented for the transceiver.

*Table continues...*

Name	Description
<b>Aux2Monitoring</b>	Indicates if auxiliary monitoring is implemented for the transceiver.
<b>TemperatureLowAlarmThreshold</b>	Indicates the low alarm threshold in degrees Celsius.
<b>TemperatureLowWarningThreshold</b>	Indicates the low warning threshold in degrees Celsius.
<b>Temperature</b>	Indicates the current temperature in degrees Celsius of the transceiver.
<b>TemperatureHighWarningThreshold</b>	Indicates the high warning threshold in degrees Celsius.
<b>TemperatureHighAlarmThreshold</b>	Indicates the high alarm threshold in degrees Celsius.
<b>TemperatureStatus</b>	Indicates if any temperature thresholds were exceeded.
<b>VoltageLowAlarmThreshold</b>	Indicates the low alarm threshold in volts.
<b>VoltageLowWarningThreshold</b>	Indicates the low warning threshold in volts.
<b>Voltage</b>	Indicates the current voltage in volts.
<b>VoltageHighWarningThreshold</b>	Indicates the high warning threshold in volts.
<b>VoltageHighAlarmThreshold</b>	Indicates the high alarm threshold in volts.
<b>VoltageStatus</b>	Indicates if any voltage thresholds were exceeded.
<b>BiasLowAlarmThreshold</b>	Indicates the bias current low alarm threshold in mA.
<b>BiasLowWarningThreshold</b>	Indicates the bias current low warning threshold in mA.
<b>Bias</b>	Indicates the laser bias current in mA.
<b>BiasHighWarningThreshold</b>	Indicates the bias current high warning threshold in mA.
<b>BiasHighAlarmThreshold</b>	Indicates the bias current high alarm threshold in mA.
<b>BiasStatus</b>	Indicates if any bias thresholds were exceeded.
<b>TxPowerLowAlarmThreshold</b>	Indicates the low alarm threshold in dBm for the Tx power.
<b>TxPowerLowWarningThreshold</b>	Indicates the low warning threshold in dBm for the Tx power.
<b>TxPowerHighWarningThreshold</b>	Indicates the high warning threshold in dBm for the Tx power.
<b>TxPowerHighAlarmThreshold</b>	Indicates the high alarm threshold in dBm for the Tx power.
<b>TxPowerStatus</b>	Indicates if any Tx power thresholds were exceeded.
<b>RxPowerLowAlarmThreshold</b>	Indicates the low alarm threshold in dBm for the Rx power.
<b>RxPowerLowWarningThreshold</b>	Indicates the low warning threshold in dBm for the Rx power.
<b>RxPower</b>	Indicates the current Rx power in dBm.
<b>RxPowerHighWarningThreshold</b>	Indicates the high warning threshold in dBm for the Rx power.
<b>RxPowerHighAlarmThreshold</b>	Indicates the high alarm threshold in dBm for the Rx power.
<b>RxPowerStatus</b>	Indicates if any Rx power thresholds were exceeded.

*Table continues...*

Name	Description
<b>Aux1LowAlarmThreshold</b>	Indicates the low alarm threshold auxiliary 1 reading.
<b>Aux1LowWarningThreshold</b>	Indicates the low warning threshold auxiliary 1 reading.
<b>Aux1</b>	Indicates the current auxiliary 1 reading.
<b>Aux1HighWarningThreshold</b>	Indicates the high warning threshold auxiliary 1 reading.
<b>Aux1HighAlarmThreshold</b>	Indicates the high alarm threshold auxiliary 1 reading.
<b>Aux1Status</b>	Indicates if any auxiliary 1 thresholds were exceeded.
<b>Aux2LowAlarmThreshold</b>	Indicates the low alarm threshold auxiliary 2 reading.
<b>Aux2LowWarningThreshold</b>	Indicates the low warning threshold auxiliary 2 reading.
<b>Aux2</b>	Indicates the current auxiliary 2 reading.
<b>Aux2HighWarningThreshold</b>	Indicates the high warning threshold auxiliary 2 reading.
<b>Aux2HighAlarmThreshold</b>	Indicates the high alarm threshold auxiliary 2 reading.
<b>Aux2Status</b>	Indicates if any auxiliary 2 thresholds were exceeded.

 **Note:**

1. Threshold and actual values for TxBias, TxPower, and RxPower are provided for all 4 channels in QSFP+ and QSFP28 optical transceivers.
2. Auxiliary monitoring does not apply to QSFP+s or QSFP28s.

# Chapter 4: Key Health Indicators (KHI)

## About this task

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

### Note:

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

You should capture KHI information during normal operations to provide a baseline for support personnel when detecting fault situations.

---

## Key Health Indicators Using the CLI

Use the procedures in this section to display Key Health Indicator (KHI) information using the CLI.

---

### Displaying KHI Performance Information

Use the following commands to display KHI information about the performance of the switch.

#### Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display buffer performance and utilization statistics:

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

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

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

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

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



**\* Note:**

Depending on the hardware platform, you can display virtual memory history.

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

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

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

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

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

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

### Example

```
Switch:1>show khi performance buffer-pool 1
```

```
Slot:1
  CPP:
    UsedFBuffs: 12
    FreeFBuffs: 3060
    RxQ0FBuffs: 0
    RxQ1FBuffs: 0
    RxQ2FBuffs: 0
    RxQ3FBuffs: 0
    RxQ4FBuffs: 0
    RxQ5FBuffs: 0
    RxQ6FBuffs: 0
    RxQ7FBuffs: 0
    TxQueueFBuffs: 0
    NoFbuff: 0
```

```
Network stack system:
  UsedMbuf: 244
  FreeMbuf: 47606
  SocketMbuf: 19
```

```
Network stack data:
  UsedMbuf: 4
  FreeMbuf: 10748
```

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

```
Switch:1>show khi performance cpu 1
```

```
Slot:1
  Current utilization: 9
  1-minute average utilization: 9
  1-minute high water mark: 14 (06/20/16 06:03:08)
  5-minute average utilization: 8
  5-minute high water mark: 10 (06/19/16 08:35:58)
```

Depending on the switch hardware, any one of the following output can appear for **show khi performance memory** [{slot[-slot][, ...]}].

```
Switch:1>show khi performance memory 1
```

```
Slot:1
  Used: 514560 (KB)
```

## Key Health Indicators (KHI)

```
Free: 521260 (KB)
Current utilization: 49 %
5-minute average utilization: 49 %
5-minute high water mark: 22 (10/08/14 14:48:01)
```

```
Switch:1>show khi performance memory 1
Slot:1
Used: 1684000 (KB)
Free: 2321704 (KB)
Current utilization: 42 %
5-minute average utilization: 41 %
5-minute high water mark: 41 (%)
10-minute average utilization: 41 %
10-minute high water mark: 41 (%)
1-Hour average utilization: 41 %
1-Hour high water mark: 41 (%)
1-Day average utilization: 41 %
1-Day high water mark: 41 (%)
1-Month average utilization: 39 %
1-Year average utilization: 0 %
```

Depending on the hardware platform, you can display virtual memory history using **show khi performance memory history**.

```
Switch:1>show khi performance memory history
Slot:1
Values indicate VMSize in KB
```

Pid	Pname	5-Min	10-Min	1-Hour	1-Day	1-Month	1-Year
4731	logger	1	1	1	1	1	
4747	namServer	20	20	20	20	20	
4748	sockserv	4	4	4	4	4	
4749	oom95	213	213	213	213	213	
4750	oom90	213	213	213	213	213	
4751	imgsync.x	19	19	19	19	19	
4818	logServer	23	23	23	23	23	
4819	trcServer	18	18	18	18	18	
4820	hwsServer	86	86	85	73	45	
4821	cbcp-main.x	789	789	787	786	782	
4822	rssServer	18	18	18	18	18	
4823	dbgServer	21	21	21	21	20	
4824	dbgShell	18	18	18	18	18	
4825	khiCollection	21	21	21	21	21	
4826	coreManager.x	19	19	19	19	19	
4827	filer	18	18	18	18	18	
4828	ssio	672	672	672	672	671	

```

--
4829 hckServer          18      18      18      18      18
--
5045 slamon.sh          3       3       3       3       3
--
4830 fabServer          20      20      20      20      20
--

```

```
Switch:1>show khi performance process 1
Slot:1
```

PID	PPID	PName	VmSize	VmLck	VmRss	VmData	VmStk	VmExe	VmLib
1	0	init	1936	0	656	164	88	32	1556
2	0	kthreadd	0	0	0	0	0	0	0
3	2	migration/0	0	0	0	0	0	0	0
4	2	ksoftirqd/0	0	0	0	0	0	0	0
5	2	watchdog/0	0	0	0	0	0	0	0
6	2	migration/1	0	0	0	0	0	0	0
7	2	ksoftirqd/1	0	0	0	0	0	0	0
8	2	watchdog/1	0	0	0	0	0	0	0
9	2	events/0	0	0	0	0	0	0	0
10	2	events/1	0	0	0	0	0	0	0
11	2	khelper	0	0	0	0	0	0	0
12	2	netns	0	0	0	0	0	0	0
13	2	async/mgr	0	0	0	0	0	0	0
14	2	sync_supers	0	0	0	0	0	0	0
15	2	bdi-default	0	0	0	0	0	0	0
16	2	kblockd/0	0	0	0	0	0	0	0
17	2	kblockd/1	0	0	0	0	0	0	0
18	2	khubd	0	0	0	0	0	0	0
19	2	kmmcd	0	0	0	0	0	0	0
22	2	rpciod/0	0	0	0	0	0	0	0
23	2	rpciod/1	0	0	0	0	0	0	0
24	2	khungtaskd	0	0	0	0	0	0	0
25	2	kswapd0	0	0	0	0	0	0	0
26	2	aio/0	0	0	0	0	0	0	0
27	2	aio/1	0	0	0	0	0	0	0
28	2	nfsiod	0	0	0	0	0	0	0
29	2	mtdblockd	0	0	0	0	0	0	0
38	2	mmcqd	0	0	0	0	0	0	0
55	1	udev	2356	0	832	264	88	96	1672
1351	2	wdd	0	0	0	0	0	0	0
1749	1	portmap	1920	0	416	164	88	16	1556
1762	1	rc	3156	0	1368	128	88	736	1808
1773	1	sshd	4948	0	904	372	88	392	3376
1779	1	syslogd	2476	0	664	172	88	564	1556
1781	1	klogd	2476	0	620	172	88	564	1556
1782	1762	S25vsp	3292	0	1532	264	88	736	1808
4366	1782	rc.appfs.vsp8k	3180	0	1424	152	88	736	1808
4660	2	i2c_wq	0	0	0	0	0	0	0
4672	2	fan_q	0	0	0	0	0	0	0
4700	2	workqueue_0	0	0	0	0	0	0	0
4702	2	workqueue_1	0	0	0	0	0	0	0
4749	4366	start	3176	0	1392	148	88	736	1808
4780	4749	lifecycle	15664	0	4856	5016	88	284	6936
4785	4780	logger	2480	0	580	176	88	564	1556
4794	4780	sockserv	4404	0	1024	72	88	8	3708
4795	4780	oom95	114768	0	107244	106084	88	84	6432
4796	4780	oom90	115032	0	107240	106348	88	84	6432
4797	4780	imgsync.x	12656	0	4332	2952	88	120	6768
4798	4794	logger	2480	0	580	176	88	564	1556
4799	4795	logger	2480	0	696	176	88	564	1556
4800	4797	logger	2480	0	696	176	88	564	1556
4801	4796	logger	2480	0	696	176	88	564	1556

## Key Health Indicators (KHI)

4839	4780	logServer	16228	0	5284	4340	88	1384	7604
4840	4780	trcServer	11264	0	3580	2544	88	124	6432
4841	4780	oobServer	10300	0	3524	1520	88	104	6444
4842	4780	cbcp-main.x	556732	0	447832	505748	88	25184	14080
4843	4780	rssServer	11236	0	3424	2544	88	96	6432
4844	4780	dbgServer	11240	0	3516	2544	88	100	6432
4845	4780	dbgShell	11084	0	3604	2412	88	84	6432
4846	4780	coreManager.x	11056	0	3576	1896	88	124	6612
4847	4780	ssio	256364	0	147604	216088	88	23328	7236
4848	4780	hckServer	11252	0	3560	2544	88	112	6432
4849	4780	remCmdAgent.x	11684	0	3960	2672	88	88	6564
4850	4839	logger	2480	0	696	176	88	564	1556
4851	4841	logger	2480	0	696	176	88	564	1556
4852	4840	logger	2480	0	696	176	88	564	1556
4853	4842	logger	2480	0	696	176	88	564	1556
4854	4844	logger	2480	0	696	176	88	564	1556
4855	4843	logger	2480	0	696	176	88	564	1556
4856	4845	logger	2480	0	700	176	88	564	1556
4857	4847	logger	2480	0	696	176	88	564	1556
4858	4846	logger	2480	0	696	176	88	564	1556
4859	4848	logger	2480	0	696	176	88	564	1556
4860	4849	logger	2480	0	696	176	88	564	1556
4907	4847	logger	2480	0	696	176	88	564	1556
4946	4780	slamon.sh	3152	0	1336	124	88	736	1808
4949	4946	logger	2480	0	580	176	88	564	1556
4973	4946	slamon_second.s	3136	0	1272	108	88	736	1808
4982	4973	ns_exec	4324	0	1020	68	88	8	3696
4989	4982	slac	4944	0	1172	460	88	8	3728

```
Switch:1>show khi performance pthread 1
Slot:1
```

TID	PID	PName	CPU(%)	5MinAvg	CPU(%)	5MinHiWater	CPU(%(time stamp))
1	1	init	0.0	0.0			
2	2	kthreadd	0.0	0.0			
3	3	migration/0	0.0	0.0			
4	4	ksoftirqd/0	0.0	0.0			
5	5	watchdog/0	0.0	0.0			
6	6	migration/1	0.0	0.0			
7	7	ksoftirqd/1	0.0	0.0			
8	8	watchdog/1	0.0	0.0			
9	9	events/0	0.0	0.0			
10	10	events/1	0.1	0.0	0.1	(10/08/14 14:27:31)	
11	11	khelper	0.0	0.0			
12	12	netns	0.0	0.0			
13	13	async/mgr	0.0	0.0			
14	14	sync_supers	0.0	0.0			
15	15	bdi-default	0.0	0.0			
16	16	kblockd/0	0.0	0.0			
17	17	kblockd/1	0.0	0.0			
18	18	khubd	0.0	0.0			
19	19	kmmcd	0.0	0.0			
22	22	rpciod/0	0.0	0.0			
23	23	rpciod/1	0.0	0.0			
24	24	khungtaskd	0.0	0.0			
25	25	kswapd0	0.0	0.0			
26	26	aio/0	0.0	0.0			
27	27	aio/1	0.0	0.0			
28	28	nfsiod	0.0	0.0			
29	29	mtdblockd	0.0	0.0			
38	38	mmcd	0.0	0.0	0.2	(10/08/14 14:27:31)	
55	55	udev	0.0	0.0			
1351	1351	wdd	0.0	0.0			
1749	1749	portmap	0.0	0.0			

1762	1762	rc	0.0	0.0
1773	1773	sshd	0.0	0.0
1779	1779	syslogd	0.0	0.0
1781	1781	klogd	0.0	0.0
1782	1782	S25vsp	0.0	0.0
4366	4366	rc.appfs.vsp8k	0.0	0.0
4660	4660	i2c_wq	0.0	0.0
4672	4672	fan_q	0.0	0.0
4700	4700	workqueue_0	0.0	0.0
4702	4702	workqueue_1	0.0	0.0
4749	4749	start	0.0	0.0
4780	4780	lifecycle	0.0	0.0
4781	4780	_Z15nd_ipc_disp	0.0	0.0
4782	4780	_Z18nd_ipc_send	0.0	0.0
4783	4780	_Z21nd_ipc_rece	0.0	0.0
4784	4780	_ZN10nd_tmr_grp	0.0	0.0
4786	4780	dpmXportRxMonit	0.0	0.0
4787	4780	dpmXportTxMonit	0.0	0.0
4788	4780	ltrBulkTimerThr	0.0	0.0
4789	4780	lc_wd_exception	0.0	0.0
4790	4780	lc_hwwd_feed	0.0	0.0
4791	4780	lc_swwd_feed	0.0	0.0
4792	4780	worker_thread	0.0	0.0
4793	4780	lc_master	0.0	0.0
4785	4785	logger	0.0	0.0
4794	4794	sockserv	0.0	0.0
4795	4795	oom95	0.0	0.0
4802	4795	_Z15nd_ipc_disp	0.0	0.0
4803	4795	_Z18nd_ipc_send	0.0	0.0
4804	4795	_Z21nd_ipc_rece	0.0	0.0
4808	4795	_ZN10nd_tmr_grp	0.0	0.0
4796	4796	oom90	0.0	0.0
4805	4796	_Z15nd_ipc_disp	0.0	0.0
4806	4796	_Z18nd_ipc_send	0.0	0.0
4807	4796	_Z21nd_ipc_rece	0.0	0.0
4809	4796	_ZN10nd_tmr_grp	0.0	0.0
4797	4797	imgsync.x	0.0	0.0
4810	4797	_Z15nd_ipc_disp	0.0	0.0
4811	4797	_Z18nd_ipc_send	0.0	0.0
4812	4797	_Z21nd_ipc_rece	0.0	0.0
4813	4797	_ZN10nd_tmr_grp	0.0	0.0
4814	4797	dpmXportRxMonit	0.0	0.0
4815	4797	dpmXportTxMonit	0.0	0.0
4816	4797	ltrBulkTimerThr	0.0	0.0
4798	4798	logger	0.0	0.0
4799	4799	logger	0.0	0.0
4800	4800	logger	0.0	0.0
4801	4801	logger	0.0	0.0
4839	4839	logServer	0.0	0.0
4873	4839	_Z15nd_ipc_disp	0.0	0.0
4874	4839	_Z18nd_ipc_send	0.0	0.0
4875	4839	_Z21nd_ipc_rece	0.0	0.0
4876	4839	_ZN10nd_tmr_grp	0.0	0.0
4840	4840	trcServer	0.0	0.0
4865	4840	_Z15nd_ipc_disp	0.0	0.0
4866	4840	_Z18nd_ipc_send	0.0	0.0
4867	4840	_Z21nd_ipc_rece	0.0	0.0
4868	4840	_ZN10nd_tmr_grp	0.0	0.0
4841	4841	oobServer	0.0	0.0
4861	4841	_Z15nd_ipc_disp	0.0	0.0
4862	4841	_Z18nd_ipc_send	0.0	0.0
4863	4841	_Z21nd_ipc_rece	0.0	0.0
4864	4841	_ZN10nd_tmr_grp	0.0	0.0
4842	4842	cbcp-main.x	0.0	0.0
4908	4842	_Z15nd_ipc_disp	0.0	0.0

0.1 (10/08/14 14:45:12)

## Key Health Indicators (KHI)

4909	4842	_Zl8nd_ipc_send	0.0	0.0	
4910	4842	_Z21nd_ipc_rece	0.1	0.0	
4911	4842	_ZN10nd_tmr_grp	0.0	0.0	
4912	4842	tUsrRoot	0.0	0.0	
4913	4842	tExcTask	0.5	0.4	0.4 (10/08/14 14:47:51)
4914	4842	tExcJobTask	0.0	0.0	
4915	4842	tNetTask	0.1	0.0	
4916	4842	traceOutput	0.0	0.0	
4917	4842	nd_profile_cmd	0.0	0.0	0.3 (10/08/14 14:44:51)
4918	4842	tRlogind	0.1	0.0	
4919	4842	tRshd	0.0	0.0	
4920	4842	tFtpdTask	0.0	0.0	
4921	4842	tFtpdTask	0.1	0.0	
4922	4842	dpmXportRxMonit	0.0	0.0	
4923	4842	dpmXportTxMonit	0.0	0.0	
4924	4842	tndMiscServTask	0.0	0.0	
4925	4842	tLoggerTask	0.0	0.0	
4926	4842	_ZN10CLimServer	0.1	0.0	
4927	4842	BootpServer	0.0	0.0	
4928	4842	tSioMsgRx	0.0	0.0	
4929	4842	chEvmTask	0.0	0.0	
4930	4842	chFsmTask	0.0	0.0	
4931	4842	chServiceTask	0.0	0.0	
4933	4842	tSnmpTmr	0.0	0.0	
4934	4842	tSnmpd	0.0	0.0	
4935	4842	tTacacspTask	0.0	0.0	
4936	4842	tTacacsqTask	0.0	0.0	
4937	4842	tMainTask	4.5	4.2	15.7 (10/08/14 14:48:41)
4938	4842	rtMainTask	0.0	0.0	
4939	4842	tCppSend	0.0	0.0	
4940	4842	tCppInterruptTa	0.4	0.1	0.9 (10/08/14 14:28:21)
4941	4842	cfmMain	0.5	0.3	0.3 (10/08/14 14:27:31)
4942	4842	tTalkClient	0.0	0.0	
4943	4842	tSlaClient	0.0	0.0	
4944	4842	cfmClock	0.0	0.0	
4947	4842	tTrapd	0.0	0.0	
4948	4842	tOspf6SpfTimer	0.0	0.0	
4955	4842	tTrapd	0.0	0.0	
4961	4842	tTdpTimer	0.0	0.0	
4962	4842	chHealthMonitor	0.0	0.0	
4963	4842	tSpfTimer	0.0	0.0	
4965	4842	tIisisTask	0.1	0.0	
4968	4842	tBgpTask	0.0	0.0	
4984	4842	tWebSrv	0.0	0.0	
4995	4842	Http0	0.0	0.0	
4996	4842	Http1	0.0	0.0	
4997	4842	Http2	0.0	0.0	
4998	4842	Http3	0.0	0.0	
4999	4842	Http4	0.0	0.0	
5000	4842	Http5	0.0	0.0	
5001	4842	Http6	0.0	0.0	
5002	4842	Http7	0.0	0.0	
5003	4842	Http8	0.0	0.0	
5004	4842	Http9	0.0	0.0	
5005	4842	Http10	0.0	0.0	
5006	4842	Http11	0.0	0.0	
5007	4842	Http12	0.0	0.0	
5008	4842	Http13	0.0	0.0	
5009	4842	Http14	0.0	0.0	
5010	4842	Http15	0.0	0.0	
5011	4842	Http16	0.0	0.0	
5012	4842	Http17	0.0	0.0	
5013	4842	Http18	0.0	0.0	
5014	4842	Http19	0.0	0.0	
5015	4842	cppTapMain	0.0	0.0	

5072	4842	tShell-cli	0.0	0.0	0.5 (10/08/14 14:27:31)
5074	4842	tTelnetd	0.0	0.0	
5075	4842	smltSlave	0.3	0.0	0.1 (10/08/14 14:30:51)
5084	4842	tTeOut_19637cc0	0.0	0.0	
5085	4842	tTeIn_19637cc0	0.0	0.0	
5086	4842	tShell-cli	0.0	0.0	
4843	4843	rssServer	0.0	0.0	
4869	4843	_Z15nd_ipc_disp	0.0	0.0	
4870	4843	_Z18nd_ipc_send	0.0	0.0	
4871	4843	_Z21nd_ipc_rece	0.0	0.0	
4872	4843	_ZN10nd_tmr_grp	0.0	0.0	
4844	4844	dbgServer	0.0	0.0	
4877	4844	_Z15nd_ipc_disp	0.0	0.0	
4878	4844	_Z18nd_ipc_send	0.0	0.0	
4879	4844	_Z21nd_ipc_rece	0.0	0.0	
4880	4844	_ZN10nd_tmr_grp	0.0	0.0	
4845	4845	dbgShell	0.0	0.0	
4881	4845	_Z15nd_ipc_disp	0.0	0.0	
4882	4845	_Z18nd_ipc_send	0.0	0.0	
4883	4845	_Z21nd_ipc_rece	0.0	0.0	
4885	4845	_ZN10nd_tmr_grp	0.0	0.0	
4846	4846	coreManager.x	0.0	0.0	
4901	4846	_Z15nd_ipc_disp	0.0	0.0	
4902	4846	_Z18nd_ipc_send	0.0	0.0	
4903	4846	_Z21nd_ipc_rece	0.0	0.0	
4904	4846	_ZN10nd_tmr_grp	0.0	0.0	
4847	4847	ssio	0.0	0.0	
4896	4847	_Z15nd_ipc_disp	0.0	0.0	
4897	4847	_Z18nd_ipc_send	0.0	0.0	
4898	4847	_Z21nd_ipc_rece	0.0	0.0	
4899	4847	_ZN10nd_tmr_grp	0.0	0.0	
4900	4847	tUsrRoot	0.0	0.0	
4905	4847	tExcTask	0.2	0.1	0.1 (10/08/14 14:27:31)
4906	4847	tty	0.0	0.0	
5016	4847	dpmXportRxMonit	0.0	0.0	
5017	4847	dpmXportTxMonit	0.0	0.0	
5018	4847	ltrBulkTimerThr	0.1	0.0	
5019	4847	nd_profile_cmd	0.0	0.0	
5020	4847	tMainTask	0.5	0.3	13.5 (10/08/14 14:48:21)
5022	4847	bcmDPC	0.0	0.0	
5023	4847	bcmINTR	2.9	2.6	3.5 (10/08/14 14:28:21)
5024	4847	socdmadesc.0	0.5	0.5	0.5 (10/08/14 14:27:31)
5056	4847	bcmTX	0.0	0.0	0.1 (10/08/14 14:45:51)
5057	4847	bcmXGS3AsyncTX	0.0	0.0	
5058	4847	bcmL2MOD.0	0.0	0.0	0.1 (10/08/14 14:45:31)
5059	4847	bcmCNTR.0	4.7	4.5	4.7 (10/08/14 14:44:40)
5060	4847	bcmL2age.0	0.0	0.0	
5061	4847	bcmRX	0.4	0.2	1.2 (10/08/14 14:27:31)
5062	4847	listener	0.1	0.1	0.7 (10/08/14 14:47:41)
5063	4847	bcmLINK.0	2.3	2.3	2.4 (10/08/14 14:28:21)
5064	4847	tUsrRoot	0.0	0.0	
5065	4847	tRspDebugPollTa	0.0	0.0	
5066	4847	tLcdIntrTask	0.0	0.0	
5067	4847	tTimerTask	0.0	0.0	
5068	4847	tScanSfp	0.1	0.0	
5071	4847	tExcJobTask	0.0	0.0	
4848	4848	hckServer	0.0	0.0	
4884	4848	_Z15nd_ipc_disp	0.0	0.0	
4886	4848	_Z18nd_ipc_send	0.0	0.0	
4887	4848	_Z21nd_ipc_rece	0.0	0.0	
4888	4848	_ZN10nd_tmr_grp	0.0	0.0	
4849	4849	remCmdAgent.x	0.0	0.0	
4889	4849	_Z15nd_ipc_disp	0.0	0.0	
4890	4849	_Z18nd_ipc_send	0.0	0.0	
4891	4849	_Z21nd_ipc_rece	0.0	0.0	

## Key Health Indicators (KHI)

```

4892 4849 _ZN10nd_tmr_grp 0.0 0.0
4893 4849 dpmXportRxMonit 0.0 0.0
4894 4849 dpmXportTxMonit 0.0 0.0
4895 4849 ltrBulkTimerThr 0.0 0.0
4850 4850 logger 0.0 0.0
4851 4851 logger 0.0 0.0
4852 4852 logger 0.0 0.0
4853 4853 logger 0.0 0.0
4854 4854 logger 0.0 0.0
4855 4855 logger 0.0 0.0
4856 4856 logger 0.0 0.0
4857 4857 logger 0.0 0.0 0.1(10/08/14 14:44:40)
4858 4858 logger 0.0 0.0
4859 4859 logger 0.0 0.0
4860 4860 logger 0.0 0.0
4907 4907 logger 0.0 0.0
4946 4946 slamon.sh 0.0 0.0
4949 4949 logger 0.0 0.0
4973 4973 slamon_second.s 0.0 0.0
4982 4982 ns_exec 0.0 0.0
4989 4989 slac 0.0 0.0
4990 4989 slac 0.0 0.0

```

```

Switch:1>show khi performance slabinfo
Slot:1

```


Name	Active Objs	Num Objs	Objsize	Objper slab	Pageper slab	Active Slabs	Num Slabs
merc_sock	0	0	384	21	2	0	0
cfq_queue	72	72	112	36	1	2	2
bsg_cmd	0	0	288	14	1	0	0
mqueue_inode_cache	15	15	544	15	2	1	1
nfs_direct_cache	0	0	80	51	1	0	0
nfs_inode_cache	0	0	600	13	2	0	0
fat_inode_cache	0	0	416	19	2	0	0
fat_cache	0	0	24	170	1	0	0
ext2_inode_cache	136	41	480	17	2	8	8
configfs_dir_cache	0	0	56	73	1	0	0
posix_timers_cache	0	0	104	39	1	0	0
rpc_inode_cache	17	17	480	17	2	1	1
UNIX	57	57	416	19	2	3	3
UDP-Lite	0	0	512	16	2	0	0
UDP	32	32	512	16	2	2	2
tw_sock_TCP	32	32	128	32	1	1	1
TCP	28	28	1120	14	4	2	2
eventpoll_pwq	204	204	40	102	1	2	2
sgpool-128	12	12	2560	12	8	1	1
sgpool-64	12	12	1280	12	4	1	1
sgpool-32	12	12	640	12	2	1	1
scsi_data_buffer	170	170	24	170	1	1	1
blkdev_queue	48	48	1288	12	4	4	4
blkdev_requests	60	44	200	20	1	3	3
biovec-256	10	10	3072	10	8	1	1
biovec-128	0	0	1536	21	8	0	0
biovec-64	0	0	768	21	4	0	0
sock_inode_cache	304	304	416	19	2	16	16
skbuff_fclone_cache	460	290	352	23	2	20	20
file_lock_cache	72	72	112	36	1	2	2
net_namespace	24	24	320	12	1	2	2
shmem_inode_cache	1170	1144	448	18	2	65	65
proc_inode_cache	777	768	376	21	2	37	37
sigqueue	56	56	144	28	1	2	2
radix_tree_node	1222	1070	296	13	1	94	94
bdev_cache	34	34	480	17	2	2	2



sysfs_dir_cache	7055	7010	48	85	1	83	83
filp	1700	1520	160	25	1	68	68
inode_cache	3243	3038	352	23	2	141	141
dentry	6210	5398	136	30	1	207	207
buffer_head	280	277	72	56	1	5	5
vm_area_struct	3358	3250	88	46	1	73	73
mm_struct	126	115	448	18	2	7	7
files_cache	72	71	224	18	1	4	4
signal_cache	119	116	480	17	2	7	7
sighand_cache	108	103	1312	12	4	9	9
task_struct	260	250	1248	13	4	20	20
anon_vma	1280	1278	16	256	1	5	5
idr_layer_cache	208	208	152	26	1	8	8
kmalloc-8192	8	8	8192	4	8	2	2
kmalloc-4096	104	99	4096	8	8	13	13
kmalloc-2048	128	115	2048	16	8	8	8
kmalloc-1024	256	256	1024	16	4	16	16
kmalloc-512	288	240	512	16	2	18	18
kmalloc-256	352	351	256	16	1	22	22
kmalloc-128	896	895	128	32	1	28	28
kmalloc-64	5120	5120	64	64	1	80	80
kmalloc-32	896	883	32	128	1	7	7
kmalloc-16	1536	1535	16	256	1	6	6
kmalloc-8	2560	2558	8	512	1	5	5
kmalloc-192	273	273	192	21	1	13	13
kmalloc-96	966	900	96	42	1	23	23

## Variable Definitions

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

Variable	Value
{slot[-slot][,...]}	Specifies the slot number. Valid slot is 1.
history	Specifies virtual memory consumed for each process.
<p> <b>Note:</b></p> <p>Depending on the hardware platform, this parameter appears in <b>show khi performance memory</b>.</p>	

## Displaying KHI Control Processor Information

Use the following commands to display key health information about the type of packets and protocols received on a port.

### Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display statistics for control packets that go to the control processor:

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

**Example**

```
Switch:1>show khi cpp port-statistics 3/1-3/7
```

```
=====
                        KHI CPP Details - Port Statistics
=====
Ports      Packet Type                                Rx Packets  Tx Packets
-----
-
3/1        LLC_TDP(134)                               498         498
3/1        LLC_ISIS(137)                              420         421
3/2        LLC_TDP(134)                               498         498
3/4        Ether2_ARP_Request(10)                    0           1
3/4        Ether2_IPv4_PIM_MC(24)                     0          101
3/4        Ether2_IPv4_OSPF_MC(32)                   318         320
3/4        Ether2_IPv4_OSPF_UC(34)                    5           0
3/4        LLC_TDP(134)                               496         496
3/5        Ether2_ARP_Request(10)                     4           4
3/5        Ether2_ARP_Other(11)                       0           4
3/5        Ether2_IPv4_PIM_MC(24)                     0          103
3/5        Ether2_IPv4_OSPF_MC(32)                    0          235
3/5        LLC_TDP(134)                               374         374
3/7        Ether2_ARP_Request(10)                     0           1
3/7        Ether2_ARP_Other(11)                       1           0
3/7        Ether2_IPv4_PIM_MC(24)                     153         151
3/7        Ether2_IPv4_PIM_UC(26)                      4           0
```

**Variable Definitions**

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

Variable	Value
<code>{slot/port[/sub-port]}</code>	Identifies a single slot and port. If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format <code>slot/port/sub-port</code> .

**Clearing KHI Information**

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

**Procedure**

1. Enter Privileged EXEC mode:
 

```
enable
```
2. Clear CPP statistics:
 

```
clear khi cpp <port-statistics>
```

## Displaying KHI Fabric Extend ONA Status

### About this task

#### \* Note:

This feature only applies to platforms that have an Open Networking Adapter (ONA) connected to it.

Use the following command to display the current status of the Fabric Extend ONA, which includes release information.

### Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display the ONA status:

```
show khi fe-ona status
```

### Example

The following output displays the `show khi fe-ona status` when the ONA is operating normally.

```
Switch:1#show khi fe-ona status
```

```
=====
                        ONA STATUS
=====
ONA Device Status : UP
Running Release Name : v1.0.0.0int006-3-g9749735-dirty
Last Image Upgrade Status : UPGRADE_SUCCESS
Last Image File Used For Upgrde: gdb-secure_ona.tgz
-----
```

The following examples display the output when communication from the switch to the ONA is disrupted. Note that the `ONA Down reason` lists the cause of the failure. The reason changes depending on the context of the failure.

The following output displays when the configuration push from the switch to the ONA fails:

```
Switch:1#show khi fe-ona status
```

```
=====
                        ONA STATUS
=====
ONA Device Status : DOWN
ONA DOWN reason : ONA_CONFIG_DOWNLOAD_FAILED
Running Release Name :
Image Upgrade Status : UNKNOWN
-----
```

The following output displays when the port connecting to the ONA device port is DOWN:

```
Switch:1#show khi fe-ona status
```

```
=====
                        ONA STATUS
=====
```

## Key Health Indicators (KHI)

```
=====
ONA Device Status : DOWN
ONA DOWN reason : ONA_DEVICE_PORT_DOWN
Running Release Name :
Image Upgrade Status : UNKNOWN
Image File Is Being Used For Upgrade :
=====
```

The following output displays when the switch is not receiving LLDP packets from the ONA:

```
Switch:1#show khi fe-ona status
```

```
=====
                        ONA STATUS
=====
ONA Device Status : DOWN
ONA DOWN reason : ONA_LLDP_TIMEOUT
Running Release Name :
Image Upgrade Status : UNKNOWN
=====
```

### \* Note:

On the switch console, the following log message precedes all three of the above cases:

```
CP1 [03/22/71 09:30:15.336:UTC] 0x00378601 00000000 GlobalRouter ONA
WARNING ONA device status detected down
```

---

## Displaying KHI Fabric Extend ONA Global Information

### About this task

#### \* Note:

This feature only applies to platforms that have an Open Networking Adapter (ONA) connected to it.

Use the following command to display Fabric Extend ONA global information such as port numbers, IP addresses, and MTU.

### Procedure

1. Enter Privileged EXEC mode:  
enable
2. Display the ONA global information:  
show khi fe-ona detail

### Example

```
Switch:1#show khi fe-ona detail
=====
                        ONA RUNTIME INFORMATION
=====
ONA Port Number : 1/15
ONA Management Address : 100.1.1.11
Tunnel Source IP Address : 198.51.100.11
ONA LLDP Port Status : Enabled
```

```

ONA Device Port Status : UP
ONA Device Status : UP
MTU : 1000
ONA Network Port Number : 1/35
ONA Mac(ARP) Address : 10:cd:ae:69:b6:50
ONA Source VlanId : 1050
ONA Source VlanIP : 192.0.2.1
ONA Gateway IP : 192.0.2.1
ONA Management IP Mask : 255.255.255.0
ONA Bootmode : 1
ONA Uptime : 0 day(s), 00:00:00
pbit-to-dscp-map p0=16 p1=20 p2=24 p3=30 p4=36 p5=40 p6=48 p7=46
-----

```

**\* Note:**

In the above example, the switch receives LLDP packets with the Management IP address of the ONA over the ONA Port (1/15). The switch extracts the ONA Management IP from the LLDP packet and resolves the ARP of the ONA over the network port (1/35). After the switch resolves the ARP of the ONA IP, the `show khi fe-ona detail` updates the following details:

- ONA Network Port Number
- ONA Mac(ARP) Address
- ONA Source VlanId

Note the following in regard to the `show khi fe-ona detail` output shown above:

- `ONA Source VlanIP : 192.0.2.1`—This is the IP address of the switch VLAN that switches traffic to the ONA network port. In the above output, this is VLAN 1050.
- `ONA Gateway IP : 192.0.2.1`—This is the ONA gateway IP address that the switch gets by querying the ONA. The ONA receives this gateway IP from the DHCP server.

**! Important:**

The `ONA Source VlanIP`, and `ONA Gateway IP` addresses must be the same for the tunnels to come up and the traffic to switch.

---

## Key Health Indicators Using EDM

Use the procedures in this section to display KHI information using EDM.

---

### Clearing KHI Statistics

#### About this task

Clear KHI statistics.

#### Procedure

1. In the Device Physical View tab, select the Device.

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

## CPP Stats Control Field Descriptions

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

Name	Description
PortStatsClear	Clears port statistics.

---

## Displaying KHI Port Information

### About this task

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

### Procedure

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

## CPP Stats Field Descriptions

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

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

# Chapter 5: IPFIX

---

## IPFIX Fundamentals

Internet Protocol Flow Information eXport (IPFIX) is an Internet Engineering Task Force (IETF) standard of export for Internet Protocol flow information.

IPFIX monitors flows that pass an observation point. The switch organizes flows into a flow group, which is contained in an observation domain.

**\* Note:**

This feature is not supported on all hardware platforms. For more information about feature support, see [Release Notes for VOSS](#).

An IPFIX flow is a set of packets that pass an observation point in the network during a certain time interval. Packets that belong to a particular flow have a common set of properties. The switch defines each property using values from the following:

- Source IP address
- Destination IP address
- IP protocol
- L4 source port
- L4 destination port

A packet belongs to a flow if it completely satisfies all defined properties of the flow.

The switch logically organizes flows into a flow group, which corresponds to a single observation point. A flow can belong to only 1 flow group. A flow group is a collection of packet flows that meet match criteria. Examples of flow groups are packets ingressing a specific physical port, or packets with a destination IP address belonging to a specific subnet.

A flow group is contained in an observation domain. The switch assigns the flow group to an observation domain. The observation domain has a unique observation domain ID that you can configure. You can configure only 1 observation domain.

The IPFIX solution consists of the following processes:

- Filtering Rules process: The Filtering Rules process gathers information about flows through different ports, or the observation point. Flow information includes the following:
  - The IPv4 source address.

- The IPv4 destination address.
- The L4 source port.
- The L4 destination port.
- The transport protocol.
- The total number of incoming packets for this flow at the observation point since the metering process (re-)initialization for this observation point.
- The total number of octets in incoming packets for this flow at the observation point since the metering process (re-)initialization for this observation point.
- The absolute timestamp of the first packet of this flow.
- The absolute timestamp of the last packet of this flow.

The Filtering Rules process runs on the switch.

- Exporting process: The Filtering Rules process sends information to the Exporting process. The Exporting process uses the UDP transport protocol for network communication with the Collecting process.

The Exporting process runs on the switch.

- Collecting process: You can view flows and export flow information periodically to a collector. 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 a collector, including the formatting of flow records and the underlying UDP transport protocol.

Use the collected information for network planning, troubleshooting a live network, and monitoring security threats.

The best practice is to use the ExtremeAnalytics™ solution as the collector. The ExtremeAnalytics™ solution provides an enhanced method of collecting IPFIX flow information.

The external collector for the IPFIX solution must support our fix template, which contains the following element IDs defined by Internet Assigned Numbers Authority (IANA) IPFIX assignments:

Element ID	Name	Description
4	protocolIdentifier	The value of the protocol number in the IP packet header.
7	sourceTransportPort	The source port identifier in the transport header.
8	sourceIPv4Address	The IPv4 source address in the IP packet header.
11	destinationTransportPort	The destination port identifier in the transport header.
12	destinationIPv4Address	The IPv4 destination address in the IP packet header.
85	octetTotalCount	The total number of octets in incoming packets for this flow at the observation point since the metering process (re-)initialization for this observation point.

*Table continues...*



Element ID	Name	Description
86	packetTotalCount	The total number of incoming packets for this flow at the observation point since the metering process (re-)initialization for this observation point.
152	flowStartMilliseconds	The absolute timestamp of the first packet of this flow.
153	flowEndMilliseconds	The absolute timestamp of the last packet of this flow.

IPFIX is a push protocol. The Filtering Rules and Exporting processes periodically send IPFIX messages to configured receivers without interaction from the Collecting process.

IPFIX collects IPv4 flow information on the switch and conforms with the following:

- IPFIX supports only 1 collector.
- IPFIX learns only IPv4 flows.
- IPFIX sends and receives only TCP/UDP flows
- IPFIX uses only UDP to export packets.
- IPFIX can detect and process a maximum of 20,000 flows.
- You can configure only the template exporting timer.
- The Out-of-Band (OOB) port does not support IPFIX.
- IPFIX exports TCP/UDP IPv4 flows on IS-IS interfaces that are members of a VLAN. IPFIX does not capture Mac-In-Mac encapsulated flows on IS-IS interfaces.

IPFIX processes IPv4 UDP or TCP Mac-in-Mac packet flows that are terminated by the switch. IPFIX does not process Mac-in-Mac packet flows that are only traversing the switch (L2 switching).

- Layer 3 Virtual Services Network (L3 VSN) flow packets on NNI ports are not learned by IPFIX.
- Control packet flows that reach the external CPU are not learned by IPFIX.

---

## IPFIX Configuration Using CLI

This section provides procedures to configure IPFIX using Command Line Interface (CLI).

---

### Enabling IPFIX Globally

#### About this task

Use the following procedure to enable IPFIX globally. IPFIX provides the ability to monitor IPv4 traffic flows.

The default global state is disabled.

## Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enable IPFIX:

```
ip ipfix enable
```

3. **(Optional)** Configure the flow aging interval:

```
ip ipfix aging-interval <1-1800>
```

4. **(Optional)** Configure a value for observation domain:

```
ip ipfix observation-domain <0-4294967295>
```

## Example

Enable IPFIX globally:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#ip ipfix enable
Switch:1(config)#ip ipfix aging-interval 30
Switch:1(config)#ip ipfix observation-domain 1
```

Disable IPFIX globally:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#no ip ipfix enable
```

```
By disabling IPFIX globally, all the processes and traffic
sent to collector(s) will be stopped.Do you agree (y/n) ? y
Switch:1(config)#show ip ipfix
```

```
=====
IPFIX Global
=====
```

```
Global-State : disable
Observation-Domain ID : 1
```

## Variable Definitions

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

Variable	Value
aging-interval <1-1800>	Specifies (in seconds) the flow record aging interval. The <code>aging-interval</code> determines how long a traffic flow that is no longer being received is retained as a flow. The default value is 40 seconds.
enable	Enables IPFIX globally.
observation-domain <0-4294967295>	Specifies a value for the observation domain used to send IPFIX messages. The default is 0 (no observation domain). If the value is 0, data that is sent is not applied to a single observation domain.

## Displaying IPFIX Global Status

### About this task

Use the following procedure to display global status information for IPFIX.

### Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display IPFIX global status:

```
show ip ipfix
```

### Example

```
Switch:1#show ip ipfix
```

```
=====
                                IPFIX Global
=====
      Global-State : enable
  Observation-Domain ID : 1
        Flow Limit : 20000
        Flow Count  : 0
    Aging Interval  : 1000
```

## Configuring the IPFIX Aging Interval

### About this task

Use the following procedure to configure an aging interval for IPFIX. The aging interval determines how long a traffic flow that is no longer being received, is retained as a flow.

### Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure a value for the aging interval:

```
ip ipfix aging-interval <1-1800>
```

### Example

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch:1(config)#ip ipfix aging-interval 40
```

## Variable Definitions

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

**Table 3:**

Variable	Value
aging-interval <1-1800>	Specifies (in seconds) the flow record aging interval. The <b>aging-interval</b> determines how long a traffic flow that is no longer being received is retained as a flow. The default value is 40 seconds.

## Configuring the IPFIX Collector

### About this task

Use the following procedure to configure a collector for IPFIX. We recommend that you use the ExtremeAnalytics™ solution as the collector.

### Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure values for the collector ID, the IP address of the collector, and the IP address of the exporter. Optionally, you can configure values for the source port sending flow information and the destination port receiving flow information:

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

#### **Note:**

You cannot configure collector or exporter IP addresses in the following formats:

- 255.255.255.255
- 127. x.x.x
- 0.x.x.x
- 224.0.0.0 to 239.255.255.255

If you configure a collector or exporter IP address in any of these formats, the following error message is displayed:

```
Error: Invalid IP address
```

3. **(Optional)** Configure a value for the export interval:

```
ip ipfix collector 1 export-interval <1-120>
```

4. **(Optional)** Configure a value for the initial burst of template packets:

```
ip ipfix collector 1 initial-burst <1-10>
```

## Example

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#ip ipfix collector 1 192.0.2.15 exporter-ip 192.0.2.16
dest-port 2 src-port 4
Switch:1(config)#ip ipfix collector 1 export-interval 40
Switch:1(config)#ip ipfix collector 1 initial-burst 4
```

## Variable Definitions

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

Variable	Value
<1-1>	Specifies the IPFIX collector ID.
{A.B.C.D}	Specifies the IP address of the collector.
exporter-ip {A.B.C.D}	Specifies the IP address of the exporter.
dest-port <1-65535>	Specifies the destination port receiving flow information.
src-port <1-65535>	Specifies the source port sending flow information.
export-interval	Specifies, in seconds, the frequency of template packet exports to the collector. The default value is 60 seconds.
initial-burst	Specifies the number of template packets sent when the collector becomes reachable. The default value is 5.

## Displaying IPFIX Collector Information

### About this task

Use the following procedure to display information about the IPFIX collector. The IPFIX collector can store flow information from multiple network devices.

### Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display information about the IPFIX collector:

```
show ip ipfix collector [<1-1>]
```

### Example

```
=====
                                IPFIX Collector-Info
=====
```

ID	Collector IP Address	Exporter IP Address	Source Port	Destination Port	Collector State	Reachable Via	Exporting Interval	Initial Burst
1	20.20.20.2	20.20.20.1	2055	2055	Enabled	-	60	5

```
=====
```

## Variable Definitions

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

Variable	Value
<1-1>	Specifies the IPFIX collector ID.

---

## Configuring an IPFIX Observation Domain

### About this task

An observation domain consists of a collection of flow groups. Use this procedure to assign a unique ID to an observation domain. The default value is 0.

### Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure the observation domain ID:

```
ip ipfix observation-domain <0-4294967295>
```

### Example

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch:1(config)#ip ipfix observation-domain 1
```

## Variable Definitions

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

Variable	Value
observation-domain <0-4294967295>	Specifies a value for the observation domain used to send IPFIX messages. The default is 0 (no observation domain). If the value is 0, data that is sent is not applied to a single observation domain.

---

## Displaying IPFIX Flows

### About this task

Use the following procedure to display information about IPFIX flows. You can display information for all IPFIX flows, or you can specify a single IPFIX flow and display additional information about that flow.

### Procedure

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

## 2. Display information about all IPFIX flows or a specific IPFIX flow:

```
show ip ipfix flows [source-addr {A.B.C.D} dest-addr {A.B.C.D}
source-port <1-65535> dest-port <1-65535> protocol {udp|tcp}]
```

### Example

Display all IPFIX flows:

```
Switch:1>show ip ipfix flows
```

```
=====
IPFIX Flows
=====
```

Source IP	Destination IP	Source Port	Destination Port	Protocol
192.0.2.17	192.0.2.19	63	63	UDP

Display information about a specific IPFIX flow:

```
Switch:1>show ip ipfix flow source-addr 192.0.2.17 dest-addr 192.0.2.19 source-port 63
dest-port 63 protocol udp
```

```
=====
IPFIX Flows
=====
```

Source IP/ Destination IP	Src Port/ Dest Port	Protocol	Start Time/ Observation Time	Byte Count/ Packet Count
192.0.2.17	63	UDP	Thu Sep 20 09:03:18 2018 EDT	32004
192.0.2.19	63		Thu Sep 20 09:03:51 2018 EDT	500

## Variable Definitions

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

Variable	Value
source-addr {A.B.C.D}	Specifies an IP address for the flow source.
dest-addr {A.B.C.D}	Specifies an IP address for the flow destination.
source-port <1-65535>	Specifies a value for the source port.
dest-port <1-65535>	Specifies a value for the destination port.
protocol {udp tcp}	Specifies the transport protocol.

---

## IPFIX Configuration Using EDM

This section provides procedures to configure IPFIX in Enterprise Device Manager (EDM).

---

### Enabling IPFIX Globally

#### About this task

IPFIX allows you to monitor IPv4 traffic flows. Use the following procedure to enable IPFIX globally. The default global state is disabled.

#### Procedure

1. In the navigation pane, expand the **Configuration > IP** folders.
2. Click **IPFIX**.
3. Click the **Globals** tab.
4. To enable IPFIX, click **enable** in the **ConfState** option box.
5. **(Optional)** To configure an observation domain ID, type a value in the **ObservationDomainID** field.
6. **(Optional)** To configure the aging interval, type a value in the **AgingInterval** field.

### Globals field descriptions

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

Name	Description
<b>ConfState</b>	Specifies whether IPFIX is enabled or disabled. The default is disabled.
<b>ObservationDomainID</b>	Specifies a value for the observation domain used to send IPFIX messages. The default is 0 (no observation domain). If the value is 0, data that is sent is not applied to a single observation domain.
<b>AgingInterval</b>	Specifies (in seconds) the flow record aging interval. The AgingInterval determines how long a traffic flow that is no longer being received is retained as a flow. The default value is 40 seconds.

---

### Configuring the IPFIX Collector

#### About this task

Use the following procedure to configure a collector for IPFIX. We recommend that you use the ExtremeAnalytics™ solution as the collector.



## Procedure

1. In the navigation pane, expand the **Configuration > IP** folders.
2. Click **IPFIX**.
3. Click the **Collector** tab.
4. Click **Insert**.
5. In the **Num** field, specify a value for the collector ID number.
6. In the **AddressType** field, specify a value for the IP address type of the collector.
7. In the **Address** field, specify a value for the IP address of the collector.
8. **(Optional)** In the **SrcPort** field, specify a value for the source port sending flow information. The default value is 2055.
9. **(Optional)** In the **DestPort** field, specify a value for the destination port receiving flow information. The default value is 2055.
10. In the **ExporterIpType** field, specify a value for the IP address type of the exporter for the collector.
11. In the **ExporterIp** field, specify the IP address of the exporter for the collector.
12. **(Optional)** In the **ExportInterval** field, specify the frequency of template packets exports to the collector.
13. **(Optional)** In the **InitialBurst** field, specify the number of template packets to send when the collector becomes reachable.
14. Click **Insert**.

## Collector field descriptions

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

Name	Description
<b>Num</b>	Specifies the ID number of the collector
<b>AddressType</b>	Specifies the IP address type of the collector.
<b>Address</b>	Specifies the IP address of the collector.
<b>Protocol</b>	Specifies the protocol used to export data from the exporter to the collector.
<b>SrcPort</b>	Specifies the source port sending flow information.
<b>DestPort</b>	Specifies the destination port receiving flow information.
<b>ExporterIpType</b>	Specifies the IP address type of the exporter for the collector.
<b>ExporterIp</b>	Specifies the IP address of the exporter for the collector.

*Table continues...*

Name	Description
<b>State</b>	Specifies the state of the collector. The default value is enabled.
<b>IsReachable</b>	Specifies whether the collector is reachable. The default value is false (not reachable)
<b>ViaNextHopName</b>	Specifies the next-hop through which the collector is reachable.
<b>ExportInterval</b>	Specifies, in seconds, the frequency of template packets exports to the collector. The default value is 60 seconds.
<b>InitialBurst</b>	Specifies the number of template packets sent when the collector becomes reachable. The default value is 5.

# Chapter 6: Link state change control

---

## Link State Change Control

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

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

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

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

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

- send a trap
- bring down the port

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

---

## Link State Change Control Using CLI

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

### Controlling Link State Changes

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

#### Procedure

1. Enter Global Configuration mode:  

```
enable
```

```
configure terminal
```
2. Configure the interval for link state changes:  

```
link-flap-detect interval <2-600>
```
3. Configure the number of changes allowed during the interval:  

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

4. Enable automatic port disabling:

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

5. Enable sending a trap:

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

**Example**

Enable automatic disabling of the port:

```
Switch:1(config)#link-flap-detect auto-port-down
```

Configure the link-flap-detect interval:

```
Switch:1(config)#link-flap-detect interval 20
```

Enable sending traps:

```
Switch:1(config)#link-flap-detect send-trap
```

**Variable Definitions**

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

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

**Displaying Link State Changes**

Displays link flap detection state changes on a physical port.

**Procedure**

1. Enter Privileged EXEC mode:

```
enable
```

2. Display link state changes:

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

## Example

```
Switch:1>enable
Switch:1#show link-flap-detect

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

## Link State Change Control Using EDM

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

### Controlling Link State Changes

#### About this task

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

#### Procedure

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

### Link Flap Field Descriptions

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

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

# Chapter 7: Logs and traps

---

## Logs and Traps

Use the information in this chapter to help you understand Simple Network Management Protocol (SNMP) traps and log files, and how to perform diagnostic and fault management functions and SNMP traps configurations using Command Line Interface (CLI) and Enterprise Device Manager (EDM).

---

## Logs and Traps Fundamentals

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

### Overview of Traps and Logs

#### System log messaging

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

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

The remote UNIX management workstation performs the following actions:

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

#### Log consolidation

The switch generates a system log file and can forward that file to a syslog server for remote viewing, storage, and analyzing.

The system log captures messages for the following components:

- Extensible Authentication Protocol (EAP)
- Remote Authentication Dial-in User Service (RADIUS)

- Remote Monitoring (RMON)
- Web
- hardware (HW)
- MultiLink Trunking (MLT)
- filter
- Quality of Service (QoS)
- Command line interface (CLI) log
- software (SW)
- Central Processing Unit (CPU)
- Internet Protocol (IP)
- Virtual Local Area Network (VLAN)
- policy
- Simple Network Management Protocol (SNMP) log

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

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

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

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

### System log client over IPv6 transport

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

### Log messages with enhanced secure mode

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

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

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

- **more**

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

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

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

**Table 4: Log commands accessible for various users**

Access level role	Commands
Administrator	The <b>remove</b> and <b>delete</b> commands.
No user at any access level.	The following commands: <ul style="list-style-type: none"> <li>• <b>more</b></li> <li>• <b>edit</b></li> <li>• <b>rename</b></li> <li>• <b>copy</b></li> </ul>
Administrator	All configuration commands can be accessed only by the individual in the administrator role, other than the preceding commands.
Administrator and auditor	All show commands for log files.
All users (Administrator, auditor, security, privilege, operator)	All show commands for log configurations.

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

### SNMP traps

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

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

## Secure Syslog

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

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



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

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

### **TLS client for secure syslog:**

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

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

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

## **Simple Network Management Protocol**

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

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

### **! Important:**

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

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

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

- Get response—An SNMP agent sends this message in response to a get request, get next request, or set request message.
- Trap—SNMP trap is a notification triggered by events at the agent.

## Log Message Format

The log messages for the switch have a standardized format. All system messages are tagged with the following information, except that alarm type and alarm status apply to alarm messages only:

- CPU slot number—Indicates the CP slot where the command is logged.
- timestamp—Records the date and time at which the event occurred. The format is MM/DD/YY hh:mm:ss.uuu, where uuu is milliseconds. Example: [11/01/10 11:41:21.376].
- hostname—The Hostname from which the message is generated.
- event code—Precisely identifies the event reported.
- alarm code—Specifies the alarm code.
- alarm type—Identifies the alarm type (Dynamic or Persistent) for alarm messages.
- alarm status—Identifies the alarm status (set or clear) for alarm messages.
- VRF name—Identifies the Virtual Routing and Forwarding (VRF) instance, if applicable.
- module name—Identifies the software module or hardware from which the log is generated.
- severity level—Identifies the severity of the message.
- sequence number—Identifies a specific CLI command.
- context—Specifies the type of the session used to connect to the switch. If the session is a remote session, the remote IP address is identified.
- user name—Specifies the user name used to login to the switch.
- CLI command—Specifies the commands typed during the CLI session. The system logs anything type during the CLI session as soon as the user presses the `Enter` key.

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

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

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

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

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

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

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

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

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

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

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

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

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

```
CP1 [07/24/14 18:04:08.304] 0x00000670 00000000 GlobalRouter SW INFO Basic license
supports all features on this device
CP1 [07/24/14 18:04:10.651] 0x00034594 00000000 GlobalRouter SW INFO System boot
CP1 [07/24/14 18:04:10.651] 0x00034595 00000000 GlobalRouter SW INFO VSP-8200 System
Software Release 0.0.0.0 B553
CP1 [07/24/14 18:04:10.779] 0x00010774 00000000 GlobalRouter HW INFO Detected 8 284XSQ
chassis
CP1 [07/24/14 18:04:10.779] 0x0001081c 00400010.2 DYNAMIC SET GlobalRouter HW INFO Slot
2 is initializing.
CP1 [07/24/14 18:04:10.780] 0x0001081c 00400010.1 DYNAMIC SET GlobalRouter HW INFO Slot
1 is initializing.
CP1 [07/24/14 18:04:10.810] 0x00010729 00000000 GlobalRouter HW INFO Detected 8284XSQ
Power Supply in slot PS 1. Adding 800 watts to available power
CP1 [07/24/14 18:04:10.811] 0x00010830 00000000 GlobalRouter HW INFO Detected 8242XSQ
module (Serial#: SDNIV84Q2013) in slot 2
```

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

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

The following table describes the system message severity levels.

**Table 5: Severity levels**

Severity level	Definition
EMERGENCY	A panic condition that occurs when the system becomes unusable. A severity level of emergency is usually a condition where multiple applications or servers are affected. You must correct a severity level of emergency immediately.
ALERT	Any condition requiring immediate attention and correction. You must correct a severity level of alert immediately, but this level usually indicates failure of a secondary system, such as an Internet Service Provider connection.
CRITICAL	Any critical conditions, such as a hard drive error.
ERROR	A nonfatal condition occurred. You can be required to take appropriate action. For example, the system generates an error message if it is unable to lock onto the semaphore required to initialize the IP addresses used to transfer the log file to a remote host.

*Table continues...*

Severity level	Definition
WARNING	A nonfatal condition occurred. No immediate action is needed. An indication that an error can occur if action is not taken within a given amount of time.
NOTIFICATION	Significant event of a normal nature. An indication that unusual, but not error, conditions have occurred. No immediate action is required.
INFO	Information only. No action is required.
DEBUG	Message containing information useful for debugging.
FATAL	A fatal condition occurred. The system cannot recover without restarting. For example, a fatal message is generated after the configuration database is corrupted.

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

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

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

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

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

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

**Table 6: Default and system log severity level mapping**

UNIX system error codes	System log severity level	Internal severity level
0	Emergency	Fatal
1	Alert	—

*Table continues...*

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

## Log Files

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

The system saves internal log messages in a circular list in memory, which overwrite older log messages as the log fills. Unlike the log messages in a log file, the internal log messages in memory do not contain encrypted information, which can limit the information available during troubleshooting. Free up the disk space on the flash if the system generates the disk space 75% full alarm. After the disk space utilization returns below 75%, the system clears the alarm, and then starts logging to a file again.

### Log file naming conventions

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

- The log file is named as log.xxxxxxxx.sss format. The prefix of the log file name is log. The six characters after the log file prefix contain the last three bytes of the chassis base MAC address. The next two characters are 01. The last three characters (sss) denote the sequence number of the log file.
- The sequence number of the log file is incremented for each new log file created after the existing log file reaches the maximum configured size.
- At initial system start up when no log file exists, a new log file with the sequence number 000 is created. After a restart, the system finds the newest log file from internal flash based on file timestamps. If the newest log file is on the flash that is used for logging, the system continues to use the newest log file. And once the maximum configured size is reached, system continues to create a new log file with incremental sequence number on the internal flash for logging.

## Log File Transfer

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

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

If log file transfer is unsuccessful, the system keeps the old log files on internal flash. The system attempts to transfer old log files after the new log file reaches the configured maximum size. The

system also attempts to transfer old log files periodically (once in one hundred log writes) if the disk space on the flash is more than 75% full.

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

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

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

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

The system appends a suffix of `.xxxxxxx.sss` to the file name. The first six characters of the suffix contain the last three bytes of the chassis base MAC address. The next two characters are 01. The last three characters (sss) denote the sequence number of the log file. For example, if you configure the name prefix as `mylog`, a possible file name is `mylog.90000001.001`.

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

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

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

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

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

Three parameters exist to configure the log file:

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

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

### Log file transfer using a wildcard filename

File transfers using SFTP require file permissions.

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

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

### Important:

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

## Email Notification

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

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

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

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

- General switch
- Chassis
- Card
- Temperature
- Power supplies
- Fans
- LEDs
- System errors
- Port lock
- Message control
- Operational configuration changes
- Current Uboot
- Port interfaces
- Port statistics

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

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

```
Subject: Logs from LabSwitch - 50712100008
From: <LabSwitch@default.com>
To: <test1@default.com>
```

```
CP1 [08/04/15 21:48:04.527:UTC] 0x00004603 00400003.67108870 DYNAMIC CLEAR GlobalRouter
SNMP INFO 2k card up(CardNum=1 AdminStatus=1 OperStatus=1)
CP1 [08/04/15 21:48:04.527:UTC] 0x00004603 00400003.67108870 DYNAMIC CLEAR GlobalRouter
SNMP INFO 2k card up(CardNum=1 AdminStatus=1 OperStatus=1)
CP1 [08/04/15 21:48:04.527:UTC] 0x00004603 00400003.67108870 DYNAMIC CLEAR GlobalRouter
SNMP INFO 2k card up(CardNum=2 AdminStatus=1 OperStatus=1)
CP1 [08/04/15 21:50:03.511:UTC] 0x00088524 00000000 GlobalRouter SW INFO Boot sequence
successful
```

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

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

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

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

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

---

## Log Configuration Using CLI

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

---

## Configuring a UNIX System Log and Syslog Host

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

### About this task

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

### Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enable the system log:

```
syslog enable
```



## 3. Specify the IP header in syslog packets:

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

## 4. Configure the maximum number of syslog hosts:

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

## 5. Create the syslog host:

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

## 6. Configure the IP address for the syslog host:

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

## 7. Enable the syslog host:

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

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

## 8. View the configuration to ensure it is correct:

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

**Example**

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

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

```
Switch:1(config)# syslog host 7 enable
```

```
Switch:1(config)#show syslog host 7
```

```

        Id : 7
        IpAddr : 192.0.2.1
        UdpPort : 514
        Facility : local7
        Severity : info|warning|error|fatal
    MapInfoSeverity : info
    MapWarningSeverity : warning
    MapErrorSeverity : error
    MapMfgSeverity : notice
    MapFatalSeverity : emergency
        Enable : true
SecureForwardingMode: none
        Tcp Port : 1025
```

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

```

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

## Variable Definitions

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

Variable	Value
enable	Enables the sending of syslog messages on the device. Use the no operator before this parameter, no syslog enable, to disable the sending of syslog messages on the device. The default is enabled.
ip-header-type <circuitless-ip default>	Specifies the IP header in syslog packets to circuitless-ip or default. <ul style="list-style-type: none"> <li>• If the value is default, the IP address of the VLAN is used for syslog packets that are transmitted in-band using input/output (I/O) ports.</li> <li>• If the value is circuitless-ip, then for all syslog messages (in-band or out-of-band), the circuitless IP address is used in the IP header. If you configure multiple circuitless IPs, the first circuitless IP configured is used.</li> </ul>
max-hosts <1-10>	Specifies the maximum number of syslog hosts supported, from 1–10. The default is 5.

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

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

*Table continues...*

Variable	Value
mapwarning {emergency alert critical error warning notice info debug}	Specifies the syslog severity to use for warning messages. The default is warning.
severity <info warning error fatal> [<info warning error fatal>] [<info warning error fatal>] [<info warning error fatal>]	Specifies the severity levels for which to send syslog messages. You can specify up to four severity levels in the same command string. The default is info.
udp-port <514-530>	Specifies the User Datagram Protocol port number on which to send syslog messages to the syslog host. This value is the UNIX system syslog host port number from 514–530. The default is 514.

## Configuring Secure Forwarding

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

### Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Create the syslog host:

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

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

3. Configure an IP address for the syslog host:

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

4. Enable the syslog host:

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

5. Enable syslog globally:

```
syslog enable
```

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

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

7. Set the TCP port:

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

8. Display the secure forwarding configured values:

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

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

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

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

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

**Next steps**

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



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

**Variable Definitions**

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

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

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

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

## Installing Root Certificate for Syslog Client

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

### Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Install a root certificate on the store:

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

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

**\* Note:**

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

3. Uninstall a root certificate from the store:

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

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

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

## Variable Definition

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

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

## Configuring Logging

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

### About this task

**\* Note:**

The platform logs CLILog and SNMPLog as INFO. Normally, if you configure the logging level to WARNING, the system skips all INFO messages. However, if you enable CLILog and SNMPLog the system logs CLI Log and SNMP Log information regardless of the logging level you configure. This is not the case for other INFO messages.

## Procedure

1. Enter Global Configuration mode:  

```
enable
```

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

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

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

```
logging screen
```

## Example

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

## Variable Definitions

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

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

## Configuring the Remote Host Address for Log Transfer

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

### Before you begin

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

### Procedure

- Enter Global Configuration mode:

```
enable
configure terminal
```

- Configure the remote host address for log transfer:

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

### Example

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

## Variable Definitions

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

Variable	Value
1-10	Specifies the file ID to transfer.
address {A.B.C.D}	Specifies the IP address of the host to which to transfer the log file. The remote host must be reachable or the configuration fails.
filename-prefix WORD<0-200>	Specifies the name of the file on the remote host. If you do not configure a name, the current log file name is the default.

## Configuring System Logging

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

### About this task

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

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

### Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enable system logging to a PC card file:

```
boot config flags logging
```

3. Configure the logfile parameters:

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

### Example

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

## Variable Definitions

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

Variable	Value
flags logging	Enables or disables logging to a flash file. The log file is named using the format log.xxxxxxxx.sss. The first six characters after the prefix of the file name log contain the last three bytes of the chassis base MAC address. The next two characters specify the slot number. The last three characters denote the sequence number of the log file.
logfile <64-500> <500-16384> <10-90>	Configures the following logfile parameters: <ul style="list-style-type: none"> <li>• &lt;64-500&gt; specifies the minimum free memory space on the external storage device from 64–500 KB. The switch does not support this parameter.</li> <li>• &lt;500-16384&gt; specifies the maximum size of the log file from 500–16384 KB.</li> <li>• &lt;10-90&gt; specifies the maximum percentage, ranging from 10–90 percent, of space on the external storage device the logfile can use. The switch does not support this parameter.</li> </ul>

## Configuring System Message Control

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

### Procedure

1. Enter Global Configuration mode:



- ```
enable
```
- ```
configure terminal
```
2. Configure system message control action:
 

```
sys msg-control action <both|send-trap|suppress-msg>
```
  3. Configure the maximum number of messages:
 

```
sys msg-control max-msg-num <2-500>
```
  4. Configure the interval:
 

```
sys msg-control control-interval <1-30>
```
  5. Enable message control:
 

```
sys msg-control
```

**Example**

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

**Variable Definitions**

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

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

**Extending System Message Control**

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

**About this task**

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

## Procedure

1. Enter Global Configuration mode:  

```
enable
```

```
configure terminal
```
2. Configure the force message control option:  

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

## Example

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

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

## Variable Definitions

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

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

## Viewing Logs

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

### About this task

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

## Procedure

1. Enter Privileged EXEC mode:  

```
enable
```
2. Show log information:

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

## Example

Display log file information:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#show logging file
CP1 [02/06/15 22:38:20.678:UTC] 0x00270428 00000000 GlobalRouter SW INFO Lifecy
cle: Start
CP1 [02/06/15 22:38:21.770:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s sockserv started, pid:4794
CP1 [02/06/15 22:38:21.771:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s oom95 started, pid:4795
CP1 [02/06/15 22:38:21.771:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s oom90 started, pid:4796
CP1 [02/06/15 22:38:21.772:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s imgsync.x started, pid:4797
CP1 [02/06/15 22:38:22.231:UTC] 0x0026452f 00000000 GlobalRouter SW INFO No pat
ch set.
CP1 [02/06/15 22:38:22.773:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s logServer started, pid:4840
CP1 [02/06/15 22:38:22.774:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s trcServer started, pid:4841
CP1 [02/06/15 22:38:22.774:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s oobServer started, pid:4842
CP1 [02/06/15 22:38:22.775:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s cbcp-main.x started, pid:4843
CP1 [02/06/15 22:38:22.776:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s rssServer started, pid:4844
CP1 [02/06/15 22:38:22.777:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s dbgServer started, pid:4845
CP1 [02/06/15 22:38:22.777:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s dbgShell started, pid:4846
CP1 [02/06/15 22:38:22.778:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s coreManager.x started, pid:4847
CP1 [02/06/15 22:38:22.779:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s ssio started, pid:4848
CP1 [02/06/15 22:38:22.780:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s hckServer started, pid:4849
CP1 [02/06/15 22:38:22.780:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s remCmdAgent.x started, pid:4850
CP1 [02/06/15 22:38:24.717:UTC] 0x000006cc 00000000 GlobalRouter SW INFO rcStar
t: FIPS Power Up Self Test SUCCESSFUL - 0
CP1 [02/06/15 22:38:24.718:UTC] 0x000006c2 00000000 GlobalRouter SW INFO rcStar
t: Security Stack Init SUCCESSFUL - 0
CP1 [02/06/15 22:38:24.718:UTC] 0x000006c3 00000000 GlobalRouter SW INFO rcStar
t: IPSEC Init SUCCESSFUL
CP1 [02/06/15 22:38:24.718:UTC] 0x000006bf 00000000 GlobalRouter SW INFO rcStar
t: Security Stack Log init SUCCESSFUL - 0
CP1 [02/06/15 22:38:26.111:UTC] 0x000005c0 00000000 GlobalRouter SW INFO Licens
eLoad = ZERO, loading premier license for developer debugging
IO1 [02/06/15 22:38:26.960:UTC] 0x0011054a 00000000 GlobalRouter COP-SW INFO De
tected Master CP in slot 1

--More-- (q = quit)

Switch:1(config)#show logging file module SNMP
CP1 [02/06/15 22:39:58.530:UTC] 0x00004595 00000000 GlobalRouter SNMP INFO Boot
ed with file
CP1 [02/06/15 22:39:59.547:UTC] 0x00004603 00400003.67108870 DYNAMIC CLEAR Glob
```

```
alRouter SNMP INFO 2k card up(CardNum=1 AdminStatus=1 OperStatus=1)
CP1 [02/06/15 22:39:59.547:UTC] 0x00004603 00400003.67108870 DYNAMIC CLEAR Glob
alRouter SNMP INFO 2k card up(CardNum=2 AdminStatus=1 OperStatus=1)
CP1 [02/06/15 22:39:59.547:UTC] 0x00004603 00400003.67108870 DYNAMIC CLEAR Glob
alRouter SNMP INFO 2k card up(CardNum=3 AdminStatus=1 OperStatus=1)
CP1 [02/06/15 22:40:45.839:UTC] 0x000045e5 00400005 DYNAMIC SET GlobalRouter SN
MP INFO Sending Cold-Start Trap
```

## Variable Definitions

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

Variable	Value
alarm	Displays alarm log entries.
CPU WORD <0-100>	Filters and lists the logs according to the CPU that generated the message. Specify a string length of 0-25 characters. To specify multiple filters, separate each CPU by the vertical bar ( ), for example, CPU1 CPU2.
detail	Displays CLI and SNMP logging information.
event-code WORD<0-10>	Specifies a number that precisely identifies the event reported.
module WORD<0-100>	Filters and lists the logs according to module. Specifies a string length of 0-100 characters. Categories include SNMP, EAP, RADIUS, RMON, WEB, HW, MLT, FILTER, QOS, CLILOG, SW, CPU, IP, VLAN, IPMC, and SNMPLOG. To specify multiple filters, separate each category by the vertical bar ( ), for example,  FILTER QOS.
name-of-file WORD<1-99>	<p>Displays the valid logs from this file. For example, /intflash/logcopy.txt. You cannot use this command on the current log file, the file into which the messages are currently logged. Specify a string length of 1 to 99 characters.</p> <p>If you enable enhanced secure mode, the system encrypts the entire log file. After you use the <code>show log file name-of-file WORD&lt;1-99&gt;</code> command, the system takes the encrypted log file name as input, then decrypts it, and prints the output to the screen. You can then redirect the decrypted output to a file that you can store onto the flash.</p> <p>If enhanced secure mode is disabled, the system only encrypts the proprietary portion of the log file.</p>
save-to-file WORD<1-99>	Redirects the output to the specified file and removes all encrypted information. You cannot use the tail option with the save-to-file option. Specify a string length of 1-99 characters.
severity WORD<0-25>	Filters and lists the logs according to severity. Choices include INFO, ERROR, WARNING, and FATAL. To specify multiple filters, separate each severity by the vertical bar ( ), for example, ERROR WARNING FATAL.
tail	Shows the last results first.
vrf WORD<0-32>	Specifies the name of a VRF instance to show log messages that only pertain to that VRF.

## Configuring CLI Logging

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

### About this task

#### \* Note:

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

### Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enable CLI logging:

```
cliolog enable
```

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

```
no cliolog enable
```

4. Ensure that the configuration is correct:

```
show cliolog
```

5. View the CLI log:

```
show logging file module cliolog
```

### Example

Enable CLI logging, and view the CLI log:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#cliolog enable
Switch:1(config)#show logging file module cliolog
CP1 [02/13/13 17:27:25.956] 0x002c0600 00000000 GlobalRouter CLILOG INFO 1 CONSOLE
rwa show snmp-server host
CP1 [02/13/13 17:28:10.100] 0x002c0600 00000000 GlobalRouter CLILOG INFO 2 CONSOLE
rwa show snmp-server notif
CP1 [02/13/13 17:28:45.732] 0x002c0600 00000000 GlobalRouter CLILOG INFO 3 CONSOLE
rwa snmp-server force-trap
CP1 [02/13/13 17:29:30.628] 0x002c0600 00000000 GlobalRouter CLILOG INFO 4 CONSOLE
rwa show logging file modug
CP1 [02/14/13 19:39:11.648] 0x002c0600 00000000 GlobalRouter CLILOG INFO 5 CONSOLE
rwa ena
CP1 [02/14/13 19:39:13.420] 0x002c0600 00000000 GlobalRouter CLILOG INFO 6 CONSOLE
rwa conf t
CP1 [02/14/13 19:49:21.044] 0x002c0600 00000000 GlobalRouter CLILOG INFO 7 CONSOLE
rwa filter acl 2 enable
```

## Logs and traps

```
CP1 [02/14/13 19:50:08.540] 0x002c0600 00000000 GlobalRouter CLILOG INFO 8 CONSOLE
rwa filter acl 2 type inpo1
CP1 [02/14/13 19:50:38.444] 0x002c0600 00000000 GlobalRouter CLILOG INFO 9 CONSOLE
rwa filter acl 2 type inpoe
CP1 [02/14/13 19:50:52.968] 0x002c0600 00000000 GlobalRouter CLILOG INFO 10 CONSOLE
rwa filter acl enable 2
CP1 [02/14/13 19:51:08.908] 0x002c0600 00000000 GlobalRouter CLILOG INFO 11 CONSOLE
rwa filter acl 2 enable
CP1 [02/15/13 06:50:25.972] 0x002c0600 00000000 GlobalRouter CLILOG INFO 14 CONSOLE
rwa ena
CP1 [02/15/13 06:50:30.288] 0x002c0600 00000000 GlobalRouter CLILOG INFO 15 CONSOLE
rwa conf t
CP1 [02/15/13 06:50:39.412] 0x002c0600 00000000 GlobalRouter CLILOG INFO 16 CONSOLE
rwa show vlan basic
CP1 [02/15/13 06:51:09.488] 0x002c0600 00000000 GlobalRouter CLILOG INFO 17 CONSOLE
rwa show isis spbm
CP1 [02/15/13 06:56:00.992] 0x002c0600 00000000 GlobalRouter CLILOG INFO 19 CONSOLE
rwa spbm 23 b-vid 2 primar1
CP1 [02/15/13 06:56:59.092] 0x002c0600 00000000 GlobalRouter CLILOG INFO 20 CONSOLE
rwa show isis
CP1 [02/15/13 07:10:54.928] 0x002c0600 00000000 GlobalRouter CLILOG INFO 21 CONSOLE
rwa show isis interface
CP1 [02/15/13 07:12:33.404] 0x002c0600 00000000 GlobalRouter CLILOG INFO 22 CONSOLE
rwa show isis spbm
CP1 [02/15/13 07:45:28.596] 0x002c0600 00000000 GlobalRouter CLILOG INFO 23 CONSOLE
rwa ena
CP1 [02/15/13 07:45:30.236] 0x002c0600 00000000 GlobalRouter CLILOG INFO 24 CONSOLE
rwa conf t
CP1 [02/15/13 07:46:29.456] 0x002c0600 00000000 GlobalRouter CLILOG INFO 25 CONSOLE
rwa interface gigabitEther0
CP1 [02/15/13 07:47:28.476] 0x002c0600 00000000 GlobalRouter CLILOG INFO 26 CONSOLE
rwa encapsulation dot1q
--More-- (q = quit)
```

## Variable Definitions

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

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

## Configuring Email Notification

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

### About this task

The SMTP feature is disabled by default.

### Before you begin

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

## Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

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

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

**\* Note:**

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

3. Configure email recipients:

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

**\* Note:**

You must configure at least one recipient.

4. Configure the SMTP server hostname or IPv4 address:

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

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

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

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

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

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

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

8. Enable the SMTP client:

```
smtp enable
```

9. Configure an SMTP domain name:

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

10. Verify the configuration:

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

## Example

Configure the SMTP client to use TCP port 26 to communicate with an SMTP server that is using port 26. Add two receiver email addresses, configure the server information using an IPv4 address,

and enable the SMTP feature. Finally, configure an SMTP domain name, and then verify the configuration.

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

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

```
Switch:1(config)#smtp event-id add 0x0000c5ec
Switch:1(config)#show smtp event-id
=====
SMTP Event IDs Information
=====
Log Event IDs: (total: 51)
0x000045e3,0x00004602,0x00004603,0x0000c5ec,0x000106ce,0x000106cf
0x000106d0,0x000106d1,0x000106d2,0x000106d4,0x000106d8,0x000106d9
0x000106da,0x000106f8,0x000106f9,0x000106fb,0x00010775,0x00010776
0x000107f5,0x000107f6,0x000305c8,0x000305ca,0x000305f1,0x00030637
0x00040506,0x00040507,0x00040508,0x00040509,0x000646da,0x000646db
0x00088524,0x000d8580,0x000d8586,0x000d8589,0x000e4600,0x000e4601
0x000e4602,0x000e4603,0x000e4604,0x000e4605,0x000e4606,0x000e4607
0x000e4608,0x000e4609,0x001985a0,0x00210587,0x00210588,0x00210595
0x00210596,0x0027458a,0x0027458d

Default Event IDs: (total: 50)
0x000045e3,0x00004602,0x00004603,0x000106ce,0x000106cf,0x000106d0
0x000106d1,0x000106d2,0x000106d4,0x000106d8,0x000106d9,0x000106da
0x000106f8,0x000106f9,0x000106fb,0x00010775,0x00010776,0x000107f5
0x000107f6,0x000305c8,0x000305ca,0x000305f1,0x00030637,0x00040506
0x00040507,0x00040508,0x00040509,0x000646da,0x000646db,0x00088524
0x000d8580,0x000d8586,0x000d8589,0x000e4600,0x000e4601,0x000e4602
0x000e4603,0x000e4604,0x000e4605,0x000e4606,0x000e4607,0x000e4608
0x000e4609,0x001985a0,0x00210587,0x00210588,0x00210595,0x00210596
0x0027458a,0x0027458d


Remove From Default: (total: 0)

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



## Variable Definitions

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

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

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

Variable	Value
add <i>WORD</i> <3-1274>	<p>Adds an email address to the recipient list. The recipients receive the email notification generated by the switch.</p> <p>You must configure at least one email recipient and can create a maximum of five email recipients. You can specify multiple addresses in a single command by separating them with a comma.</p> <p>You cannot use quotation marks (") or commas (,) in email addresses. Other restrictions for the format of the email address follow RFC 5321.</p> <p>The maximum length for the address is 254 characters.</p>
remove <i>WORD</i> <3-1274>	<p>Removes an email address from the recipient list. The recipients receive the email notification generated by the switch. You can specify multiple addresses in a single command by separating them with a comma.</p> <p>You cannot use quotation marks (") or commas (,) in email addresses. Other restrictions for the format of the email address follow RFC 5321.</p> <p>The maximum length for the address is 254 characters.</p>

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

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

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

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

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

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

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

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

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

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

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

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

## Log Configuration Using EDM

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

### Configuring the System Log

#### About this task

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

#### Procedure

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

### System Log Field Descriptions

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

Name	Description
<b>Enable</b>	Enables or disables the syslog feature. If you select this variable, this feature sends a message to a server on a network that is configured to receive and store diagnostic messages from this device. You can configure the type of messages sent. The default is enabled.
<b>MaxHosts</b>	Specifies the maximum number of remote hosts considered active and can receive messages from the syslog service. The range is 0–10 and the default is 5.

*Table continues...*

Name	Description
<b>OperState</b>	Specifies the operational state of the syslog service. The default is active.
<b>Header</b>	<p>Specifies the IP header in syslog packets to circuitlessIP or default.</p> <ul style="list-style-type: none"> <li>• If the value is default, the IP address of the VLAN is used for syslog packets that are transmitted in-band using input/output (I/O) ports.</li> <li>• If the value is circuitlessIP, the circuitless IP address is used in the IP header for all syslog messages (in-band or out-of-band). If you configure multiple circuitless IPs, the first circuitless IP configured is used.</li> </ul> <p>The default value is default.</p>

## Configuring the System Log Table

### About this task

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

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

### Procedure

1. In the navigation pane, expand the **Configuration > Edit > Diagnostics** folders.
2. Click **System Log**.
3. Click the **System Log Table** tab.
4. Click **Insert**.
5. Configure the parameters as required.
6. Click **Insert**.
7. To modify mappings, double-click a parameter to view a list of options.
8. Click **Apply**.

## System Log Table Field Descriptions

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

Name	Description
<b>Id</b>	Specifies the ID for the syslog host. The range is 1–10.

*Table continues...*

Name	Description
<b>AddressType</b>	Specifies if the address is an IPv4 or IPv6 address.
<b>Address</b>	Specifies the IP address of the syslog host. You can log system log messages to external system log hosts with both IPv4 and IPv6 addresses.
<b>UdpPort</b>	Specifies the UDP port to use to send messages to the syslog host (514–530). The default is 514.
<b>Enable</b>	Enables or disables the sending of messages to the syslog host. The default is disabled.
<b>HostFacility</b>	Specifies the syslog host facility used to identify messages (local0 to local7). The default is local7.
<b>Severity</b>	Specifies the message severity for which syslog messages are sent. The default is info warning error fatal.
<b>MapInfoSeverity</b>	Specifies the syslog severity to use for INFO messages. The default is info.
<b>MapWarningSeverity</b>	Specifies the syslog severity to use for WARNING messages. The default is warning.
<b>MapErrorSeverity</b>	Specifies the syslog severity to use for ERROR messages. The default is error.
<b>MapFatalSeverity</b>	Specifies the syslog severity to use for FATAL messages. The default is emergency.
<b>MapMfgSeverity</b>	Specifies the syslog severity to use for Accelar manufacturing messages. The default is notice.
<b>SecureForwardingTcpPort</b>	Specifies the TCP port to use for secure forwarding for a particular host. The default is 1025.
<b>SecureForwardingMode</b>	Enables or disables secure forwarding of syslog over remote port forwarding. The supported values are tls and none. The default is none, which means that secure forwarding is disabled.
<b>SecureForwardingServerCertName</b>	Specifies the server certificate name.  Certificate validation is done only if the server certificate name is configured.

## Configuring Email Notification

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

### About this task

The SMTP feature is disabled by default.

### Before you begin

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

## Procedure

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

**\* Note:**

You must configure at least one recipient.


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

## Globals Field Descriptions

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

Name	Description
<b>ServerAddressType</b>	Specifies the type of server address as either an IPv4 address or a hostname. If you use a hostname, you must configure the DNS client on the switch.
<b>ServerAddress</b>	Specifies the SMTP server address. You can use either a hostname or an IPv4 address. If you use a hostname, you must configure the DNS client on the switch.
<b>ReceiverEmailsList</b>	Specifies the recipient list. The recipients receive the email notification generated by the switch.  You must configure at least one email recipient and can create a maximum of five email recipients. You can specify multiple addresses in a single command by separating them with a comma.

*Table continues...*

Name	Description
	<p>You cannot use quotation marks (") or commas (,) in email addresses. Other restrictions for the format of the email address follow RFC5321.</p> <p>The maximum length for the address is 254 characters.</p>
<b>NumOfEmails</b>	Shows the total number of addresses in <b>ReceiverEmailsList</b> .
<b>SenderEmail</b>	Specifies the email address that appears in the From field of the message that the switch generates. By default, the switch uses <i>SystemName@default.com</i> .
<b>DomainName</b>	<p>Specifies the SMTP domain name.</p> <p>The maximum length is 254 characters.</p>
<b>Port</b>	<p>Specifies the TCP port on the switch that the SMTP client uses to communicate with the SMTP server. The default value is 25.</p> <p> <b>Note:</b></p> <p>You must disable the SMTP feature before you can change an existing SMTP port configuration.</p> <p>The port you specify must match the port that the SMTP server uses.</p>
<b>SystemStatusSendTimer</b>	Specifies the interval, in seconds, at which the switch sends status information. The default is 30 seconds. A value of 0 means the switch does not send status information.
<b>Enable</b>	Enables or disables the SMTP feature. By default, SMTP is disabled.
<b>LogEventIds</b>	<p>Specifies the list of events that generate email notification. You can specify multiple event IDs in a single command by separating them with a comma.</p> <p>The event ID can be up to 10 digits in hexadecimal format.</p>
<b>NumOfEventIds</b>	Shows the total number of IDs in <b>LogEventIds</b> .
<b>DefaultLogEventIds</b>	Shows the default list of event IDs that generate email notification.
<b>NumOfDefaultEventIds</b>	Shows the total number of IDs in <b>DefaultLogEventIds</b> .

---

## SNMP Trap Configuration Using CLI

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

For more information about how to configure SNMP community strings and related topics, see [Configuring Security for VOSS](#).

---

### Configuring an SNMP Host

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

#### Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure an SNMPv1 host:

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

3. Configure an SNMPv2c host:

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

4. Configure an SNMPv3 host:

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

5. Ensure that the configuration is correct:

```
show snmp-server host
```

#### Example

Configure the target table entry. Configure an SNMPv3 host.

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



## Variable Definitions

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

Variable	Value
inform [timeout <1-2147483647>] [retries <0-255>] [mms <0-2147483647>]	Sends SNMP notifications as inform (rather than trap). To use all three options in one command, you must use them in the following order: <ol style="list-style-type: none"> <li>1. timeout &lt;1-2147483647&gt; specifies the timeout value in seconds with a range of 1–214748364.</li> <li>2. retries &lt;0-255&gt; specifies the retry count value with a range of 0–255.</li> <li>3. mms &lt;0-2147483647&gt; specifies the maximum message size as an integer with a range of 0–2147483647.</li> </ol>
filter WORD<1-32>	Specifies the filter profile to use.
noAuthNoPriv authNoPriv AuthPriv	Specifies the security level.
port <1-65535>	Specifies the host server port number.
WORD<1-32>	Specifies the security name, which identifies the principal that generates SNMP messages.
WORD<1-256>	Specifies either an IPv4 or IPv6 address.

## Configuring an SNMP Notify Filter Table

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

### Before you begin

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

### Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Create a new notify filter table:

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

3. Ensure that the configuration is correct:

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

### Example

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

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

## Variable Definitions

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

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

## Configuring SNMP Interfaces

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

### Procedure

1. Enter Global Configuration mode:
 

```
enable
configure terminal
```
2. Configure the destination and source IP addresses for SNMP traps:
 

```
snmp-server sender-ip {A.B.C.D} {A.B.C.D}
```
3. If required, send the source address (sender IP) as the sender network in the notification message:
 

```
snmp-server force-trap-sender enable
```
4. If required, force the SNMP and IP sender flag to use the same value:

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

### Example

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

## Variable Definitions

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

Variable	Value
authentication-trap enable	Activates the generation of authentication traps.
force-iphdr-sender enable	Automatically configures the SNMP and IP sender to the same value. The default is disabled.
force-trap-sender enable	Sends the configured source address (sender IP) as the sender network in the notification message.
sender-ip <A.B.C.D> <A.B.C.D>	Configures the SNMP trap receiver and source IP addresses. Specify the IP address of the destination SNMP server that receives the SNMP trap notification in the first IP address.  Specify the source IP address of the SNMP trap notification packet that is transmitted in the second IP address. If this address is 0.0.0.0, the system uses the IP address of the local interface that is closest (from an IP routing table perspective) to the destination SNMP server.

## Enabling SNMP Trap Logging

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

### Before you begin

- You must configure and enable the syslog server.

### About this task

#### Note:

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

### Procedure

- Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

- Enable SNMP trap logging:

```
snmplog enable
```

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

```
no snmplog enable
```

4. View the contents of the SNMP log:

```
show logging file module snmplog
```

**Example**

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

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

**Variable Definitions**

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

Variable	Value
enable	Enables the logging of traps.  Use the command <code>no snmplog enable</code> to disable the logging of traps.

---

## SNMP Trap Configuration Using EDM

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

For information about how to configure SNMP community strings and related topics, see [Configuring Security for VOSS](#).

---

## Configuring an SNMP Host Target Address

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

## Procedure

1. In the navigation pane, expand the **Configuration > Edit > SnmpV3** folders.
2. Click **Target Table**.
3. In the **Target Table** tab, click **Insert**.
4. In the **Name** box, type a unique identifier.
5. In the **TDomain** box, select the transport type of the address. Select either **ipv4Tdomain** or **ipv6Tdomain**.
6. In the **TAddress** box, type the transport address and User Datagram Protocol (UDP) port.
7. In the **Timeout** box, type the maximum round trip time.
8. In the **RetryCount** box, type the number of retries to be attempted.
9. In the **TagList** box, type the list of tag values.
10. In the **Params** box, type the SnmpAdminString.
11. In the **TMask** box, type the mask.
12. In the **MMS** box, type the maximum message size.
13. Click **Insert**.

## Target Table Field Descriptions

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

Name	Description
<b>Name</b>	Specifies a unique identifier for this table. The name is a community string.
<b>TDomain</b>	Specifies the transport type of the address. <b>ipv4Tdomain</b> specifies the transport type of address is an IPv4 address. <b>ipv6Tdomain</b> specifies the transport type of address is IPv6. The default is ipv4Tdomain.
<b>TAddress</b>	Specifies the transport address in xx.xx.xx.xx:port format, for example: 192.1.2.12:162, where 162 is the trap listening port on the system 192.1.2.12.
<b>Timeout</b>	Specifies the maximum round trip time required to communicate with the transport address. The value is in 1/100 seconds from 0–2147483647. The default is 1500.  After the system sends a message to this address, if a response (if one is expected) is not received within this time period, you can assume that the response is not delivered.
<b>RetryCount</b>	Specifies the maximum number of retries if a response is not received for a generated message. The count can be in the range of 0–255. The default is 3.

*Table continues...*

Name	Description
<b>TagList</b>	Contains a list of tag values used to select target addresses for a particular operation. A tag refers to a class of targets to which the messages can be sent.
<b>Params</b>	Contains SNMP parameters used to generate messages to send to this transport address. For example, to receive SNMPv2C traps, use TparamV2.
<b>TMask</b>	Specifies the mask. The value can be empty or in six-byte hex string format. Tmask is an optional parameter that permits an entry in the TargetAddrTable to specify multiple addresses.
<b>MMS</b>	Specifies the maximum message size. The size can be zero, or 484–2147483647. The default is 484.  Although the maximum message size is 2147483647, the device supports the maximum SNMP packet size of 8192.

## Configuring Target Table Parameters

### About this task

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

### Procedure

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

## Target Params Table Field Descriptions

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

Name	Description
<b>Name</b>	Identifies the target table.

*Table continues...*

Name	Description
<b>MPModel</b>	Specifies the message processing model to use to generate messages: SNMPv1, SNMPv2c, or SNMPv3/USM.
<b>SecurityModel</b>	Specifies the security model to use to generate messages: SNMPv1, SNMPv2c, or USM. You can receive an inconsistentValue error if you try to configure this variable to a value for a security model that the implementation does not support.
<b>SecurityName</b>	Identifies the principal on whose behalf SNMP messages are generated.
<b>SecurityLevel</b>	Specifies the security level used to generate SNMP messages: noAuthNoPriv, authNoPriv, or authPriv.

## Configuring SNMP Notify Filter Profiles

### About this task

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

### Procedure

1. In the navigation pane, expand the **Configuration > Edit > Snpv3** folders.
2. Click **Notify Table**.
3. Click the **Notify Filter Table** tab.
4. Click **Insert**.
5. In the **NotifyFilterProfileName** box, type a name for the notify filter profile.
6. In the **Subtree** box, type subtree location information in x.x.x.x.x.x.x.x. format.
7. In the **Mask** box, type the mask location in hex string format.
8. From the **Type** options, select **included** or **excluded**.
9. Click **Insert**.

## Notify Filter Table Field Descriptions

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

Name	Description
<b>NotifyFilterProfileName</b>	Specifies the name of the filter profile used to generate notifications.
<b>Subtree</b>	Specifies the MIB subtree that, if you combine it with the mask, defines a family of subtrees, which are included in or excluded from the filter profile. For more information, see RFC 2573.

*Table continues...*

Name	Description
<b>Mask</b>	Specifies the bit mask (in hexadecimal format) that, in combination with the subtree, defines a family of subtrees, which are included in or excluded from the filter profile.
<b>Type</b>	Indicates whether the family of filter subtrees are included in or excluded from a filter. The default is included.

## Configuring SNMP Notify Filter Profile Table Parameters

### Before you begin

- The notify filter profile exists.

### About this task

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

### Procedure

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

## Notify Filter Profile Table Field Descriptions

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

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

## Enabling Authentication Traps

### About this task

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

### Procedure

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



2. Click **General**.
3. Click the **Error** tab.
4. Select **AuthenticationTraps**.
5. Click **Apply**.

## Error Field Descriptions

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

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

---

## Viewing the Trap Sender Table

### About this task

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

### Procedure

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

## Trap Sender Table Field Descriptions

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

Name	Description
<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 8: MACsec performance

## MACsec Statistics

This feature is not supported on all hardware platforms. For more information about feature support, see [Release Notes for VOSS](#).

MAC Security (MACsec) is an IEEE 802<sup>®</sup> standard that allows authorized systems in a network to transmit data confidentially and to take measures against data transmitted or modified by unauthorized devices.

The switch supports the following statistics that provide a measure of MACsec performance.



**Table 7: General MACsec statistics**

Statistics	Description
TxUntaggedPkts	Specifies the number of transmitted packets without the MAC security tag (SecTAG), with MACsec disabled on the interface.
TxTooLongPkts	Specifies the number of transmitted packets discarded because the packet length is greater than the Maximum Transmission Unit (MTU) of the Common Port interface.
RxUntaggedPkts	Specifies the number of received packets without the MAC security tag (SecTAG), with MACsec <i>not</i> operating in strict mode.
RxNoTagPkts	Specifies the number of received packets without the MAC security tag (SecTAG), with MACsec operating in strict mode.
RxBadTagPkts	Specifies the number of received packets discarded with an invalid SecTAG or with a zero value Packet Number (PN)/invalid Integrity Check Value (ICV).
RxUnknownSCIPkts	Specifies the number of packets received with an unknown Secure Channel Identifier (SCI) and with MACsec <i>not</i> operating in strict mode.
RxNoSCIPkts	Specifies the number of packets received with an unknown Secure Channel Identifier (SCI) and with MACsec operating in strict mode.
RxOverrunPkts	Specifies the number of packets discarded because the number of received packets exceeded the cryptographic performance capabilities.

**Table 8: Secure-channel inbound MACsec statistics**

Statistics	Description
UnusedSAPkts	Specifies the summation of received unencrypted packets on all SAs of this secure channel, with MACsec <i>not</i> in strict mode.

*Table continues...*

Statistics	Description
NoUsingSAPkts	Specifies the summation of received packets that were discarded along with either encrypted packets or packets that were received with MACsec operating in strict mode.
LatePkts	Specifies the number of packets received that have been discarded for this Secure Channel (SC) with Replay Protect enabled.   <b>Note:</b> The switch does not support Replay Protect.
NotValidPkts	Specifies the summation of packets that were discarded in all SAs of the SC because they were not valid with one of the following conditions: <ul style="list-style-type: none"> <li>• MACsec was operating in strict mode</li> <li>• The packets received were encrypted but contained erroneous fields.</li> </ul>
InvalidPkts	Specifies the summation of all packets received that were not valid for this SC, with MACsec operating in <i>check</i> mode.
DelayedPkts	Specifies the summation of packets for this SC, with the Packet Number (PN) of the packets lower than the lower bound replay protection PN.   <b>Note:</b> The switch does not support Replay Protect.
UncheckedPkts	The total number of packets for this SC that: <ul style="list-style-type: none"> <li>• were encrypted and had failed the integrity check</li> <li>• were <i>not</i> encrypted and had failed the integrity check</li> <li>• were received when MACsec validation was not enabled</li> </ul>
OKPkts	Specifies the total number of Integrity Check Validated (ICV) packets for all SAs of this Secure Channel. The number of octets of User Data recovered from received frames that were integrity protected but not encrypted.
OctetsValidated	Specifies the number of octets of plaintext recovered from received packets that were integrity protected but not encrypted.
OctetsDecrypted	Specifies the number of octets of plaintext recovered from received packets that were integrity protected and encrypted.

**Table 9: Secure-channel outbound MACsec statistics**

Statistics	Description
ProtectedPkts	Specifies the number of integrity protected but not encrypted packets for this transmitting SC.
EncryptedPkts	Specifies the number of integrity protected and encrypted packets for this transmitting SC.
OctetsProtected	Specifies the number of plain text octets that are integrity protected but not encrypted on the transmitting SC.
OctetsEncrypted	Specifies the number of plain text octets that are integrity protected and encrypted on the transmitting SC.

## Viewing MACsec Statistics Using the CLI

Use the following procedure to view MAC Security (MACsec) statistics using CLI.

### Viewing MACsec Statistics

Perform this procedure to view the MACsec statistics.

#### Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View the general MACsec statistics:

```
show macsec statistics [{slot/port[/sub-port]}[-slot/port[/sub-port]]
[,...]]
```

3. View the secure-channel inbound MACsec statistics:

```
show macsec statistics [{slot/port[/sub-port]}[-slot/port[/sub-port]]
[,...]] secure-channel inbound
```

4. View the secure-channel outbound MACsec statistics:

```
show macsec statistics [{slot/port[/sub-port]}[-slot/port[/sub-port]]
[,...]] secure-channel outbound
```

#### Example

Display general MACsec statistics, inbound MACsec statistics, and outbound MACsec statistics:

**\* Note:**

Slot and port information can differ depending on hardware platform. For more information, see the hardware documentation for your platform.

The switch supports MACsec on specific ports. For more information, see your hardware documentation.

```
Switch:1>enable
Switch:1#show macsec statistics 1/40
```

MACSEC Port Statistics				
PortId	TxUntagged Packets	TxTooLong Packets	RxUntagged Packets	RxNoTag Packets
1/40	0	0	0	0
PortId	RxBadTag Packets	RxUnknown SCIPackets	RxNoSCI Packets	RxOverrun Packets

```
1/40      0          0          0          0
```

```
Switch:1#show macsec statistics 1/40 secure-channel inbound
```

```
=====
MACSEC Port Inbound Secure Channel Statistics
=====
PortId    UnusedSA    NoUsingSA    Late    NotValid    Invalid
Packets   Packets     Packets      Pkts    Packets     Packets
-----
1/40      0           0            0       100037      0
PortId    Delayed     Unchecked    Ok       Octets      Octets
Packets   Packets     Packets      Pkts    Validated   Decrypted
-----
1/40      0           0            0       53528828    0
```

```
Switch:1#show macsec statistics 1/40 secure-channel outbound
```

```
=====
MACSEC Port Outbound Secure Channel Statistics
=====
PortId    Protected   Encrypted    Octets     Octets
Packets   Packets     Packets      Protected  Encrypted
-----
1/40      0           99946       0          53434154
```

---

## Viewing MACsec Statistics Using EDM

Use the following procedures to view MAC Security (MACsec) statistics using EDM.

---

### Viewing MACsec Interface Statistics

Use this procedure to view the MACsec interface statistics using EDM.

This feature is not supported on all hardware platforms. If you do not see this command in EDM, the feature is not supported on your hardware. For more information about feature support, see [Release Notes for VOSS](#).

#### Procedure

1. In the Device Physical View tab, select the port for which you need to view the MACsec interface statistics.

The switch supports MACsec on specific ports. For more information, see your hardware documentation.

2. In the navigation pane, expand the **Edit > Port > General** folders.
3. Click the **MacSec Interface Stats** tab.

**\* Note:**

Use the **Clear Stats** button to clear MACsec interface statistics. The **Clear Stats** button is available to clear single-port as well as multiple-port MACsec interface statistics.

## MacSec Interface Stats Field Descriptions

The following table describes the fields in the MacSec Interface Stats tab.

Field	Description
<b>TxUntaggedPkts</b>	Specifies the number of transmitted packets without the MAC security tag (SecTAG), with MACsec disabled on the interface.
<b>TxTooLongPkts</b>	Specifies the number of transmitted packets discarded because the packet length is greater than the maximum transmission unit (MTU) of the common port interface.
<b>RxUntaggedPkts</b>	Specifies the number of received packets without the MAC security tag (SecTAG), with MACsec <i>not</i> operating in strict mode.
<b>RxNoTagPkts</b>	Specifies the number of received packets without the MAC security tag (SecTAG), with MACsec operating in strict mode.
<b>RxBadTagPkts</b>	Specifies the number of received packets discarded with an invalid SecTAG, or with a zero value packet number (PN), or invalid Integrity Check Value (ICV).
<b>RxUnknownSCIPkts</b>	Specifies the number of packets received with an unknown secure channel identifier (SCI), and with MACsec <i>not</i> operating in strict mode.
<b>RxNoSCIPkts</b>	Specifies the number of packets received with an unknown secure channel identifier (SCI), and with MACsec operating in strict mode.
<b>RxOverrunPkts</b>	Specifies the number of packets discarded because the number of received packets exceeded the cryptographic performance capabilities.

---

## Viewing Secure Channel (SC) Inbound Statistics

Use this procedure to view the secure channel (SC) inbound statistics using EDM.

This feature is not supported on all hardware platforms. If you do not see this command in EDM, the feature is not supported on your hardware. For more information about feature support, see [Release Notes for VOSS](#).

## Procedure

1. In the Device Physical View tab, select the port for which you need to view the SC inbound statistics.

The switch supports MACsec on specific ports. For more information, see your hardware documentation.


2. In the navigation pane, expand the **Edit > Port > General** folders.
3. Click the **SC Inbound Stats** tab.

 **Note:**


Use the **Clear Stats** button to clear single-port secure channel inbound statistics. The **Clear Stats** button is not available to clear multiple-port secure channel inbound statistics.

## SC Inbound Stats Field Descriptions

The following table describes the fields in the **SC Inbound Stats** tab.

Field	Description
<b>UnusedSAPkts</b>	Specifies the summary of received unencrypted packets on all SAs of this secure channel, with MACsec <i>not</i> in strict mode.
<b>NoUsingSAPkts</b>	Specifies the summary of received packets that were discarded along with either encrypted packets or packets that were received with MACsec operating in strict mode.
<b>LatePkts</b>	Specifies the number of packets received that have been discarded for this secure channel (SC) with Replay Protect enabled.   <b>Note:</b> The switch does not support Replay Protect.
<b>NotValidPkts</b>	Specifies the summary of packets that were discarded in all SAs of the SC because they were not valid with one of the following conditions: <ul style="list-style-type: none"> <li>• MACsec was operating in strict mode.</li> <li>• The packets received were encrypted but contained erroneous fields.</li> </ul>
<b>InvalidPkts</b>	Specifies the summary of all packets received that were not valid for this SC, with MACsec operating in <i>check</i> mode.
<b>DelayedPkts</b>	Specifies the summary of packets for this SC, with the packet number (PN) of the packets lower than the lower bound replay protection PN.

*Table continues...*

Field	Description
	<p> <b>Note:</b></p> <p>The switch does not support Replay Protect.</p>
<b>UncheckedPkts</b>	<p>The total number of packets for this SC that:</p> <ul style="list-style-type: none"> <li>• Were encrypted and had failed the integrity check.</li> <li>• Were <i>not</i> encrypted and had failed the integrity check.</li> <li>• Were received when MACsec validation was not enabled.</li> </ul>
<b>AcceptedPkts</b>	<p>Specifies the total number of Integrity Check Validated (ICV) packets for all SAs of this Secure Channel. The number of octets of User Data recovered from received frames that were integrity protected but not encrypted.</p>
<b>OctetsValidated</b>	<p>Specifies the number of octets of plaintext recovered from received packets that were integrity protected but not encrypted.</p>
<b>OctetsDecrypted</b>	<p>Specifies the number of octets of plaintext recovered from received packets that were integrity protected and encrypted.</p>

## Viewing Secure Channel (SC) Outbound Statistics

Use this procedure to view the secure channel (SC) outbound statistics using EDM.

This feature is not supported on all hardware platforms. If you do not see this command in EDM, the feature is not supported on your hardware. For more information about feature support, see [Release Notes for VOSS](#).

### Procedure

1. In the Device Physical View tab, select the port for which you need to view the SC outbound statistics.  
The switch supports MACsec on specific ports. For more information, see your hardware documentation.
2. In the navigation pane, expand the **Edit > Port > General** folders.
3. Click the **SC Outbound Stats** tab.

 **Note:**

Use the **Clear Stats** button to clear single-port secure channel outbound statistics. The **Clear Stats** button is not available to clear multiple-port secure channel outbound statistics.



## SC Outbound Stats Field Descriptions

The following table describes the fields in the **SC Outbound Stats** tab.

Field	Description
<b>ProtectedPkts</b>	Specifies the number of integrity protected but not encrypted packets for this transmitting SC.
<b>EncryptedPkts</b>	Specifies the number of integrity protected and encrypted packets for this transmitting SC.
<b>OctetsProtected</b>	Specifies the number of plain text octets that are integrity protected but not encrypted on the transmitting SC.
<b>OctetsEncrypted</b>	Specifies the number of plain text octets that are integrity protected and encrypted on the transmitting SC.

# Chapter 9: Remote Monitoring

This chapter provides conceptual information and procedures to configure Remote Monitoring (RMON1) and (RMON2).

---

## Remote Monitoring

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). Use CLI, or EDM, to globally enable RMON on the system. After you globally enable RMON, you enable monitoring for individual devices on a port-by-port basis.

RMON1 is the original version of the protocol, which collects information for OSI Layer 1 and Layer 2 in Ethernet networks. RMON1 provides traffic statistics at the MAC layer, and provides statistics on Ethernet segments for packets and bytes received and transmitted.

You can use RMON1 to:

- Configure alarms for user-defined events.
- Collect Ethernet statistics.
- Log events.
- Send traps for events.

Within EDM, you can configure RMON1 alarms that relate to specific events or variables. You can also specify events associated with alarms to trap or log-and-trap. In turn, the system traps or logs tripped alarms.

You can view all RMON1 information using CLI or EDM. Alternatively, you can use any management application that supports SNMP traps to view RMON1 trap information.

This section describes RMON1 alarms, RMON1 history, RMON1 events, and RMON1 statistics.

### **RMON1 alarms**

You can configure alarms to alert you if the value of a variable goes out of range. You can define RMON1 alarms on any MIB variable that resolves to an integer value. You cannot use string variables (such as system description) as alarm variables.

You can use RMON1 alarm to monitor anything that has a MIB OID associated with it and a valid instance.

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.


After you activate alarms, you can:

- View the activity in a log and/or a trap.
- Create a script directing the system to sound an audible alert at a console.
- Create a script directing the system to send an e-mail.
- Create a script directing the system to call a pager.

The system polls the alarm variable and the system compares the result against upper and lower limit values you select when you create the alarm. If the system reaches or crosses the alarm variable during the polling period, 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 cause a console to beep, send an e-mail, or call a pager.

The upper limit of the alarm is the rising value, and the lower limit is the falling value. RMON1 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.

The following figure shows how alarms fire:



**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 system crosses the opposite threshold. Therefore, you must 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.

You can 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 you define the lower limit of exiting octets at 260 and you define the upper limit at 320 (or at any value greater than  $260 + 52 = 312$ ).

The rising alarm fires the first time outbound traffic, other than spanning tree Bridge Protocol Data Units (BPDUs), occurs. The falling alarm fires after outbound traffic, other than spanning tree, ceases. This process provides the time intervals of any nonbaseline outbound traffic.

If you define the alarm with a falling threshold of less than 260 and the alarm polling interval is at 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.

The following figure shows an example of the alarm threshold:



**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 the system logs an event or trap.

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 the value 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.

**\* Note:**

If you create an alarm that monitors a variable that does not exist, you will receive an error message and the creation will fail. Also, if the variable you are monitoring is no longer valid at the time of sampling, the switch removes the alarm automatically. For example, if you create an alarm that monitors some information about a VLAN, and that VLAN is later removed, then the switch silently removes the associated alarm at the next sampling interval.

### RMON1 history

The RMON1 history group records periodic statistical samples from a network. A sample is a history and the system gathers the sample in time intervals referred to as buckets.

You can use RMON1 history for the MAC layer in the network. You cannot use RMON1 history for application and network layer protocols.

You enable and create histories to establish a time-dependent method to gather RMON1 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 system reaches the last bucket, the system dumps bucket 1 and recycles the bucket to hold a new bucket of statistics. Then the system dumps bucket 2, and so forth.

### **RMON1 events**

RMON1 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 system records the activity.

You can use RMON1 events to monitor anything that has a MIB OID associated with it and a valid instance.

An event specifies whether a trap, a log, or both a trap and a log generates to view alarm activity.

You must create an event before associating it with an alarm, otherwise an error occurs. Also, you cannot delete an event as long as there are alarms associated with it. If you try to do so, an error message displays.

### **RMON1 statistics**

You can use EDM to gather and graph statistics in a variety of formats, or you can save the statistics to a file and export the statistics to a third-party presentation or graphing application.

### **RMON1 scaling limits**

The following tables shows the scaling limits for RMON1 elements.

#### **\* Note:**

When the log table reaches the maximum 500 log limit, the oldest third of the logs per event is removed to make room for new events. For all other elements, a message displays when you reach the maximum limit and no other element can be added.

Alarms	100
Events	100
History (entries in the history control table with 2000 buckets shared between them)	20
Logs	500
Statistics (entries in stats table)	100

## RMON 2

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).


Use CLI or EDM, to globally enable RMON on the system.

After you globally enable RMON, you enable monitoring for individual devices on a port-by-port basis.

RMON1 is the original version of the protocol, which collects information for OSI Layer 1 and Layer 2 in Ethernet networks. RMON1 provides traffic statistics at the MAC layer, and provides statistics on Ethernet segments for packets and bytes received and transmitted.

The RMON2 feature monitors network and application layer protocols on configured network hosts, either VLAN or port interfaces, that you enable for monitoring. The RMON2 feature expands the capacity of RMON1 to upper layer protocols in the OSI model.

The following figure shows which form of RMON monitors which layers in the OSI model:



**Figure 3: OSI model and RMON**

The RMON2 feature 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). The switch supports a partial implementation of RMON2. The RMON2 feature adds the following MIBs: protocol directory, protocol distribution, address map, network-layer host and application layer host for the traffic passing through the (Control Processor) CP for these MIB tables.

The system only collects statistics for IP packets that pass through the CP. RMON2 does not monitor packets on other interfaces processed on the switch that do not pass through the CP.

After you globally enable RMON2, enable monitoring for individual devices. Identify the network hosts for the system to monitor with a manual configuration on the interfaces you want to monitor.

The RMON2 feature monitors a list of predefined protocols. The system begins to collect protocol statistics immediately after you enable RMON.

The RMON2 feature collects statistics on:

- Protocols predefined by the system.
- Address mapping between physical and network address on particular network hosts that you configure for monitoring.
- Network host statistics for particular hosts on a network layer protocol (IP) that you configure for monitoring.
- Application host statistics for a particular host on an application layer protocol that you configure for monitoring.

## **RMON2 MIBs**

This section describes the following MIBs, on which RMON2 can collect statistics: protocol directory, protocol distribution, address map, network-layer host, and application layer host.

### **Protocol directory MIB**

The protocol directory is a master directory that lists all of the protocols RMON2 can monitor. The protocols include network layer, transport layer, and application layer protocols, under the OSI model. The system only monitors statistics for the predefined protocols. You cannot delete or add additional protocols to this table. The protocol directory MIB is enabled by default for the predefined protocols.

The predefined protocols include:

- Internet Protocol (IP)
- Secure Shell version 2 (SSHv2)
- Transmission Control Protocol (TCP)
- User Datagram Protocol (UDP)
- File Transfer Protocol (FTP)
- Hypertext Transfer Protocol (HTTP)
- Telnet
- Remote login (rlogin)
- Trivial File Transfer Protocol (TFTP)
- Simple Network Management Protocol (SNMP)

### **Protocol distribution MIB**

The protocol distribution MIB collects traffic statistics that each protocol generates by local area network (LAN) segment. The switch acts as the probe and the system collects protocol statistics for the entire switch as part of the group for all of the protocols predefined in the protocol directory table. The protocol distribution control table is part of this group. The protocol distribution control table is predefined with an entry for the management IP for the switch to represent the network segment where the system collects the statistics.

No CLI or EDM support exists to add or delete entries in this table.

### **Address map MIB**

The address map MIB maps the network layer IP to the MAC layer address.

The system populates the address map control table MIB with an entry for each host interface that you enable for monitoring on the switch.

### **Network layer host MIB**

The network layer host MIB monitors the Layer 3 traffic statistics for each host. The network layer host MIB monitors traffic packets in and out of hosts based on the network layer address. The network layer host controls the network and application layer host tables.

The system populates an entry for the management IP of the switch to represent the network segment where the system collects the statistics. You have to enable each host interface that you want to monitor on the switch.

The system only collects statistics for this group from packets that go to the CP.

### **Application layer host MIB**

The application layer host MIB monitors traffic statistics by application protocol for each host.

The system populates an entry for the management IP of the switch to represent the network segment where the system collects the statistics. You have to enable each host interface that you want to monitor on the switch.

The system only collects statistics for this group from packets that go to the CP.

---

## **RMON Configuration Using CLI**

This section contains procedures to configure RMON using Command Line Interface (CLI).

For information about RMON statistics, see the following sections in the Statistics chapter:

- [Displaying RMON statistics for specific ports](#) on page 206
- [Viewing RMON statistics](#) on page 224

---

## **Configuring RMON**

Enable RMON1 and RMON2 globally, and configure RMON1 alarms, events, history, statistics, and whether port utilization is calculated in half or full duplex. By default, RMON1 and RMON2 are disabled globally.

For RMON1, you enable RMON globally, and then you can use RMON1 alarm, history, events, and statistics for the MAC layer in the network. You cannot use RMON1 history or statistics for application and network layer protocols.

For RMON2, you enable RMON globally, and then you enable RMON on the host interfaces you want to monitor.



## Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enable RMON1 and RMON2 globally:

```
rmon
```

3. Configure an RMON1 alarm:

```
rmon alarm <1-65535> WORD <1-1536> <1-3600> {absolute|delta}
[falling-threshold <-2147483647-2147483647> event <1-65535>] [owner
WORD<1-127>] [rising-threshold <-2147483647-2147483647> event
<1-65535>]
```

4. Configure an RMON1 event:

```
rmon event <1-65535> [community WORD<1-127>] [description
WORD<0-127>] [log] [owner WORD<1-127>] [trap] [trap_dest
[{A.B.C.D}]] [trap_src [{A.B.C.D}]]
```

5. Configure RMON1 history:

```
rmon history <1-65535> {slot/port [/sub-port] [-slot/port [/sub-port]
[,...]} [buckets <1-65535>] [interval <1-3600>] [owner WORD<1-127>]
```

6. Configure RMON1 statistics:

```
rmon stats <1-65535> {slot/port [/sub-port] [-slot/port [/sub-port]
[,...]} [owner <1-127>]
```

7. Configure whether the system calculates port utilization in half or full duplex:

```
rmon util-method [half|full]
```

## Example

Configure RMON globally, an RMON1 alarm, and RMON1 event:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#rmon
Switch:1(config)#rmon event 60534 community public description "Rising Event" log trap
Switch:1(config)#rmon alarm 4 rcCliNumAccessViolations.0 10 absolute rising-threshold 2
event 60000
```

## Variable Definitions

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

Variable	Value
<p>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;]</p>	<p>Creates an alarm interface.</p> <ul style="list-style-type: none"> <li>• &lt;1-65535&gt;— Specifies the interface index number from 1 to 65535. Each entry defines a diagnostics sample at a particular interval for an object on the device. The default is 1.</li> <li>• WORD &lt;1-1536&gt;— Specifies the variable name or OID. The entry is case sensitive and can have a string length of 1 to 1536.</li> <li>• {absolute   delta} — Specifies the sample type.</li> <li>• rising-threshold &lt;-2147483648-2147483647&gt; [&lt;event: 1-65535&gt;] — Specifies the rising threshold from -2147483648 to 2147483647, which is 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, the system generates a single event. The system also generates a single event if the first sample after this entry that becomes valid is greater than or equal to the rising alarm, or the rising or falling alarm. After the system generates a rising event, the system does not generate another such event until the sampled value falls below this threshold and reaches the alarm falling threshold. You cannot modify this object if the associated alarm status is equal to valid.</li> </ul> <p>&lt;1-65535&gt;— Specifies the rising event index, which the system uses after the system crosses a rising threshold. The event entry identified by a particular value of this index is the same as identified by the same value of the event index object. If no corresponding entry exists in the event table, no association exists. In particular, if this value is zero, the system does not generate an associated event, as zero is not a valid event index. You cannot modify this object if the associated alarm status is equal to valid.</p> <ul style="list-style-type: none"> <li>• falling-threshold &lt;-2147483648-2147483647&gt; [&lt;event: 1-65535&gt;] — Specifies the falling threshold from -2147483648 to 2147483647, which 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, the system generates a single event. The system also generates a single event if the first sample after this entry that becomes valid is less than or equal to this threshold and the associated alarm startup alarm is equal to falling alarm or rising or falling alarm. After the system generates a falling event, the system does not generate another such event until the sampled value rises above this threshold, and reaches the alarm rising threshold. You cannot</li> </ul>

*Table continues...*

Variable	Value
	<p>modify this object if the associated alarm status is equal to valid.</p> <p>&lt;1-65535&gt; – Specifies the index of the event entry that the system uses after a falling threshold is crossed. The event entry identified by a particular value of this index is the same as identified by the same value of the event index object. If no corresponding entry in the event table exists, no association exists. In particular, if this value is zero, the system does not generate an event, as zero is not a valid event index. You cannot modify this object if the associated alarm status is equal to valid. The default is 60535.</p> <ul style="list-style-type: none"> <li>owner <i>WORD</i>&lt;1-127&gt; — Specifies the name of the owner, with a string length 1 to 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><b>* Note:</b></p> <p>When configuring from CLI, the default owner is <code>cli</code>; when configuring with SNMP, the default owner is <code>snmp</code>. The default command only sets the owner to default. No other parameters can be changed after you create the alarm.</p> <p>Use the no operator to disable RMON alarms: <code>no rmon alarm [&lt;1-65535&gt;]</code></p>
<p>event &lt;1-65535&gt; [community <i>WORD</i>&lt;1-127&gt;] [description <i>WORD</i>&lt;0-127&gt;] [log] [owner <i>WORD</i>&lt;1-127&gt;] [trap]</p>	<p>Create an event.</p> <ul style="list-style-type: none"> <li>&lt;1-65535&gt;— Specifies the event index number. Each entry defines one event that the system generates after the appropriate conditions occur. The default is 1.</li> <li>log — Specifies if this event stores a log when the event is triggered by the alarm.</li> <li>trap — Specifies if this event sends a trap when the event is triggered by the alarm. The trap will be sent to all the snmp-server hosts configured in the snmp table.</li> <li>description <i>WORD</i>&lt;0-127&gt;— Specifies the event description, with a string length of 0 to 127.</li> <li>owner <i>WORD</i>&lt;1-127&gt; — Specifies the name of the owner, with a string length of 1 to 127.</li> <li>community <i>WORD</i>&lt;1-127&gt; — Specifies the SNMP community where you can send SNMP traps, with a string length 1 to 127.</li> </ul> <p>You can set the community, but the trap is not filtered out. The trap is sent to all configured snmp-server hosts, regardless of the value of this field.</p>

Table continues...

Variable	Value
	Use the no operator to delete a RMON event: <code>no rmon event [&lt;1-65535&gt;] [log]</code>
<p><code>history &lt;1-65535&gt; {slot/port [/sub-port] [-slot/port [/sub-port] [, ...]] [buckets &lt;1-65535&gt;] [interval &lt;1-3600&gt;] [owner WORD&lt;1-127&gt;]</code></p>	<p>Configures RMON history.</p> <ul style="list-style-type: none"> <li>• <code>&lt;1-65535&gt;</code> — Specifies the history index number that uniquely identifies an entry in the history control table. Each entry defines a set of samples at a particular interval for an interface on the default. The default value is 1.</li> <li>• <code>{slot/port [/sub-port] [-slot/port [/sub-port] [, ...]]</code> — Specifies the single port interface. Identifies the source for which the system collects and places historical data in a media-specific table on behalf of this history control entry. The source is an interface on this device. The statistics in this group reflect all packets on the local network segment that attaches to the identified interface.</li> <li>• <code>buckets &lt;1-65535&gt;</code>— Specifies the requested number of discrete time intervals where the system saves data in the part of the media-specific table associated with this history control entry. The default value is 50.</li> <li>• <code>interval &lt;1-3600&gt;</code>— Specifies the time interval in seconds over which the system samples the data for each bucket in the part of the media-specific table associated with this history control entry. 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 the associated counters. Consider the minimum time in which a counter can overflow on a particular media type, and then set the history control interval to a value less than this interval, which is typically most important for the octets counter in a media-specific table. The default value is 1800.</li> <li>• <code>owner WORD&lt;1-127&gt;</code>— Specifies the name of the owner.</li> </ul>
<p><code>stats &lt;1-65535&gt; {slot/port [/sub-port] [-slot/port [/sub-port] [, ...]] owner WORD&lt;1-127&gt;</code></p>	<p>Configures RMON statistics.</p> <ul style="list-style-type: none"> <li>• <code>&lt;1-65535&gt;</code>— Specifies the control Ether statistics entry index number.</li> <li>• <code>{slot/port [/sub-port] [-slot/port [/sub-port] [, ...]]</code>— Specifies the single port interface.</li> <li>• <code>owner WORD&lt;1-127&gt;</code> — Specifies the name of the owner.</li> </ul> <p>Use the no operator to delete a RMON Ether stats control interface: <code>no rmon stats [&lt;1-65535&gt;]</code></p>
<p><code>util-method [half full]</code></p>	<p>Configures whether port utilization is calculated in half or full duplex to calculate port usage.</p> <ul style="list-style-type: none"> <li>• <code>half</code>—Configures the string to half duplex.</li> </ul>

*Table continues...*

Variable	Value
	<ul style="list-style-type: none"> <li>• full—Configures the string to full duplex.</li> </ul> <p>After you select half for half duplex, RMON uses InOctets and the speed of the port to calculate port usage (this is the standard RMON RFC 1271 convention). After you select full for full duplex, RMON uses InOctets and OutOctets, and 2X the speed of the port to calculate port usage. If you select full, but the port operates in half-duplex mode, the calculation defaults to the RFC1271 convention. The default is half.</p>

## Enabling Remote Monitoring on an Interface

Use the following procedure to enable Remote Monitoring (RMON) on an interface.

### Before you begin

- Enable RMON globally.

### Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enable RMON on a particular VLAN:

```
vlan rmon <1-4059>
```

3. Enter GigabitEthernet Interface Configuration mode:

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

#### \* Note:

If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

4. Enable RMON on a particular port:

```
rmon
```

### Example

Enable RMON on VLAN 2:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#vlan rmon 2
```

**Enable RMON on port 3/8:**

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface gigabitethernet 3/8
Switch:1(config-if)#rmon
```

**Variable Definitions**

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

Variable	Value
<1-4059>	Specifies the VLAN ID on which to configure RMON.

**Displaying RMON Information**

View RMON1 and RMON2 information on the switch. You can display information on RMON1 alarms, events, history, logs, and statistics. You can also display RMON2 information on application host statistics, control tables, network host statistics, and protocol distribution statistics.

**Procedure**

1. View RMON1 information:

```
show rmon {alarm|event|history|log|stats}
```

2. View RMON2 information:

```
show rmon {address-map|application-host-stats WORD<1-64>|application
protocols|ctl-table|protocol-dist-stats|network-host-stats}
```

**Example**

View RMON event, log, and statistics information:

```
Switch:1(config)#show rmon event
```

```
=====
                        Rmon Event
=====
INDEX DESCRIPTION      TYPE      COMMUNITY OWNER      LAST_TIME_SENT
-----
60534 Rising Event     log-and-trap public     192.0.2.155 none
60535 Falling Event    log-and-trap public     192.0.2.155 8 day(s), 19:14:32
```

```
Switch: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"

Switch:1(config)#show rmon stats

=====
Rmon Ether Stats
=====

INDEX  PORT   OWNER
-----
1      1/10   monitor

```

## Variable Definitions

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

Variable	Value
address-map	Displays the RMON2 address map. This RMON2 parameter expands RMON capacity to display information on network, transport, and application layers.
alarm	Displays the RMON1 alarm table.
application-host-stats <i>WORD&lt;1-64&gt;</i>	Displays RMON2 application host statistics from one of the following protocols: TCP, UDP, FTP, Telnet HTTP, rLogin, SSHv2, TFTP, SNMP, HTTPS. This RMON2 parameter expands RMON capacity to display network, transport, and application layers.
ctl-table	Displays the RMON2 control tables. This RMON2 parameter expands RMON capacity to display network, transport, and application layers.
event	Displays the RMON1 event table.
history	Displays the RMON1 history table. This RMON1 parameter displays and is limited to link layer information, including as MAC information.
log	Displays the RMON1 log table.
network-host-stats	Displays RMON2 network-host statistics. This RMON2 parameter expands RMON capacity to display network, transport, and application layers.
protocol-dist-stats	Displays RMON2 protocol distribution statistics. This RMON2 parameter expands RMON capacity to display network, transport, and application layers.
stats	Displays the RMON1 statistics table. This RMON1 parameter displays and is limited to link layer information, including as MAC information.

## Displaying RMON Status

View the current RMON status on the switch.

### Procedure

1. Log on to the switch to enter User EXEC mode.
2. View RMON status:

```
show rmon
```

**Example**

```
Switch:1# show rmon
RMON Info :
Status      : enable
```

**Displaying RMON Address Maps**

View the maps of network layer address to physical address to interface.

The probe adds entries based on the source MAC and network addresses in packets without MAC-level errors.

**Procedure**

1. Log on to the switch to enter User EXEC mode.
2. View RMON address maps:

```
show rmon address-map
```

**Example**

```
Switch:1# show rmon address-map
=====
                        Rmon Address Map Table
=====
PROTOIDX  HOSTADDR      SOURCE  PHYADDR      LASTCHANGE
-----
1          192.0.2.11    2060    b0:ad:aa:42:a5:03  10/09/15 17:30:41
```

**Job Aid**

The following table describes the fields in the output for the **show rmon address-map** command.

Parameter	Description
PROTOIDX	Shows a unique identifier for the entry in the table.
HOSTADDR	Shows the network address for this entry. The format of the value depends on the protocol portion of the local index.
SOURCE	Shows the interface or port on which the network address was most recently seen.
PHYADDR	Shows the physical address on which the network address was most recently seen.
LASTCHANGE	Shows when the entry was created or last changed. If this value changes frequently, it can indicate duplicate address problems.

**Displaying RMON Application Host Statistics**

View application host statistics to see traffic statistics by application protocol for each host.

**Procedure**

1. Log on to the switch to enter User EXEC mode.
2. View RMON application host statistics:



```
show rmon application-host-stats WORD<1-64>
```

**Example**

```
Switch:1# show rmon application-host-stats ?
WORD<1-64> Select one of these application protocols
           {TCP|UDP|FTP|TELNET|HTTP|RLOGIN|SSH|TFTP|SNMP|HTTPS}
Switch:1# show rmon application-host-stats FTP

=====
                        Rmon Application Host Stats
=====
HOSTADDR      INPKT      OUTPKT      INOCT      OUTOCT      CREATETIME
-----
192.0.2.10    0          0          0          0          10/09/15 17:29:54
```

**Job Aid**

The following table describes the fields in the output for the **show rmon application-host-stats** command.

Parameter	Description
HOSTADDR	Shows the network address for this entry. The format of the value depends on the protocol portion of the local index.
INPKT	Shows the number of packets for this protocol type, without errors, transmitted to this address. This value is the number of link-layer packets so a single, fragmented network-layer packet can increment the counter several times.
OUTPKT	Shows the number of packets for this protocol type, without errors, transmitted by this address. This value is the number of link-layer packets so a single, fragmented network-layer packet can increment the counter several times.
INOCT	Shows the number of octets transmitted to this address, excluding octets in packets that contained errors. This value counts octets in the entire packet that contained the protocol, not just the particular protocol frames.
OUTOCT	Shows the number of octets transmitted by this address, excluding octets in packets that contained errors. This value counts octets in the entire packet that contained the protocol, not just the particular protocol frames.
CREATETIME	Shows when the entry was last activated.

**Displaying RMON Control Tables**

View RMON control tables to see the data source for both network layer and application layer host statistics.

**Procedure**

1. Log on to the switch to enter User EXEC mode.
2. View RMON control tables:

```
show rmon ctl-table
```

**Example**

```
Switch:1# show rmon ctl-table
```

```
=====
                        Rmon Control Table
=====
                        Protocol Directory Table
=====
IDX  PROTOCOL  ADDRMAPCFG  HOSTCFG  MATRIXCFG  OWNER
-----
1    IP        SUPPORTED   SUPPORTED  NOT SUPPORTED  Switch-1
2    TCP        SUPPORTED   SUPPORTED  NOT SUPPORTED  Switch-1
3    UDP        SUPPORTED   SUPPORTED  NOT SUPPORTED  Switch-1
4    FTP        SUPPORTED   SUPPORTED  NOT SUPPORTED  Switch-1
5    SSH        SUPPORTED   SUPPORTED  NOT SUPPORTED  Switch-1
6    TELNET    SUPPORTED   SUPPORTED  NOT SUPPORTED  Switch-1
7    HTTP      SUPPORTED   SUPPORTED  NOT SUPPORTED  Switch-1
8    RLOGIN    SUPPORTED   SUPPORTED  NOT SUPPORTED  Switch-1
9    TFTP      SUPPORTED   SUPPORTED  NOT SUPPORTED  Switch-1
10   SNMP      SUPPORTED   SUPPORTED  NOT SUPPORTED  Switch-1
11   HTTPS     SUPPORTED   SUPPORTED  NOT SUPPORTED  Switch-1
=====
                        Protocol Distribution Control Table
=====
IDX  DATASOURCE  DROPPFRAMES  CREATETIME  OWNER
-----
1    0.0.0.0     0             09/22/15 19:29:13  Switch-1
=====
                        Address Map Control Table
=====
IDX  DATASOURCE  DROPPFRAMES  OWNER
-----
1    0.0.0.0     0             Switch-1
=====
                        Host Control Table
=====
IDX  DATASOURCE  NHDROPPFRAMES  AHDROPPFRAMES  OWNER
-----
1    0.0.0.0     0               0               Switch-1
=====
```

**Job Aid**

The following table describes the fields in the output for the `show rmon ctl-table` command.

Parameter	Description
ADDRMAPCFG	Describes and configures the probe support for the network layer and application layer host tables for this protocol. The value can be one of the following: <ul style="list-style-type: none"> <li>• <b>NOT SUPPORTED</b></li> <li>• <b>SUPPORTED OFF</b></li> <li>• <b>SUPPORTED ON</b></li> </ul>

*Table continues...*

Parameter	Description
	If the value is <b>SUPPORTED ON</b> , the probe adds entries to the address map table that maps the network layer address to the MAC layer address.
AHDROPFRAMES	Shows the total number of application layer host frames that the probe receives and drops. This value does not include packets that were not counted because they had MAC-layer errors.
CREATETIME	Shows when the entry was last activated.
DATASOURCE	Shows the source of data for the entry.
DROPFRAMES	Shows the total number of frames that the probe receives and drops. This value does not include packets that were not counted because they had MAC-layer errors.
HOSTCFG	Describes and configures the probe support for the network layer and application layer host tables for this protocol. The value can be one of the following: <ul style="list-style-type: none"> <li>• <b>NOT SUPPORTED</b></li> <li>• <b>SUPPORTED OFF</b></li> <li>• <b>SUPPORTED ON</b></li> </ul> If the value is <b>SUPPORTED ON</b> , the probe adds entries to the Host Control table to collect statistics for network layer and application layer hosts.
IDX	Shows a unique identifier for the entry in the table.
MATRIXCFG	Describes and configures the probe support for the network layer and application layer host tables for this protocol. The value can be one of the following: <ul style="list-style-type: none"> <li>• <b>NOT SUPPORTED</b></li> <li>• <b>SUPPORTED OFF</b></li> <li>• <b>SUPPORTED ON</b></li> </ul>
NHDROPFRAMES	Shows the total number of network host frames that the probe receives and drops. This value does not include packets that were not counted because they had MAC-layer errors.
OWNER	Shows the entity that configured this entry.
PROTOCOL	Shows the protocols RMON2 can monitor: <ul style="list-style-type: none"> <li>• Internet Protocol (IP)</li> <li>• Transmission Control Protocol (TCP)</li> <li>• User Datagram Protocol (UDP)</li> <li>• File Transfer Protocol (FTP)</li> <li>• Secure Shell version 2 (SSHv2)</li> <li>• Telnet</li> <li>• Hypertext Transfer Protocol (HTTP)</li> </ul>

*Table continues...*

Parameter	Description
	<ul style="list-style-type: none"> <li>• Remote login (RLOGIN)</li> <li>• Trivial File Transfer Protocol (TFTP)</li> <li>• Simple Networking Management Protocol (SNMP)</li> <li>• Hypertext Transfer Protocol Secure (HTTPS)</li> </ul>

## Displaying RMON Network Host Statistics

View network host statistics to see Layer 3 traffic statistics for each host. The network layer host MIB monitors traffic packets in and out of hosts based on the network layer address.

### Procedure

1. Log on to the switch to enter User EXEC mode.
2. View RMON network host statistics:

```
show rmon network-host-stats
```

### Job Aid

The following table describes the fields in the output for the `show rmon network-host-stats` command.

Parameter	Description
HOSTADDR	Shows the host address for this entry.
INPKT	Shows the number of packets without errors transmitted to this address. This value is the number of link-layer packets so a single, fragmented network-layer packet can increment the counter several times.
OUTPKT	Shows the number of packets without errors transmitted by this address. This value is the number of link-layer packets so a single, fragmented network-layer packet can increment the counter several times.
INOCT	Shows the number of octets transmitted to this address, excluding octets in packets that contained errors. This value counts octets in the entire packet that contained the protocol, not just the particular protocol frames.
OUTOCT	Shows the number of octets transmitted by this address, excluding octets in packets that contained errors. This value counts octets in the entire packet that contained the protocol, not just the particular protocol frames.
CREATETIME	Shows when the entry was last activated.

## Displaying RMON Protocol Distribution Statistics

View protocol distribution statistics to see traffic statistics that each protocol generates by local area network (LAN) segment.

### Procedure

1. Log on to the switch to enter User EXEC mode.
2. View RMON protocol distribution statistics:

```
show rmon protocol-dist-stats
```

### Example

```
Switch:1# show rmon protocol-dist-stats
```

```
=====
                        Rmon Protocol Dist Stats
=====
```

PROTOCOL	PKTS	OCTETS
IP	0	0
TCP	0	0
UDP	0	0
FTP	0	0
SSH	0	0
TELNET	0	0
HTTP	0	0
RLOGIN	0	0
TFTP	0	0
SNMP	0	0
HTTPS	0	0

```
=====
```

## RMON Configuration Using EDM

This section contains procedures to configure RMON using Enterprise Device Manager (EDM).

For information about RMON statistics, see the following sections in the Statistics chapter:

- [Enabling RMON statistics](#) on page 304
- [Viewing RMON statistics](#) on page 305

## Enabling RMON Globally

### About this task

You must globally enable RMON before you can use RMON2 functions. If you attempt to enable an RMON2 function before the global flag is disabled, EDM informs you that the flag is disabled and prompts you to enable the flag. You can configure RMON1 while RMON is globally disabled.

If you want to use nondefault RMON parameter values, you can configure them before you enable RMON, or as you configure the RMON functions.

### Procedure

1. In the navigation pane, expand the **Configuration > Serviceability > RMON** folders.
2. Click **Options**.
3. Click the **Options** tab.
4. Select the **Enable** check box.
5. In the **UtilizationMethod** option, select a utilization method.

- 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 <b>Enable</b> check box, the RMON agent starts immediately. To disable RMON, clear the <b>Enable</b> check box and click <b>Apply</b> 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.

---

## Enabling RMON on a Port or VLAN

Use the following procedure to enable RMON on an interface.

### Before you begin

- Enable RMON globally.

### Procedure

- Enable RMON on a VLAN:
  - In the navigation pane, expand the **Configuration > VLAN** folders.
  - Click **VLANs**.
  - Click the **Advanced** tab.
  - In the row for the VLAN, double-click the **RmonEnable** field, and then select **enable**.
  - Click **Apply**.
- Enable RMON on a port:
  - In the Device Physical View, select a port.
  - In the navigation pane, expand the **Configuration > Edit > Port** folders.
  - Click **General**.
  - Click the **Interface** tab.
  - For the **RmonEnable** field, select **enable**.
  - Click **Apply**.

## Enabling RMON1 History

### About this task

Use RMON1 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 the history characteristics, you cannot modify them; you must delete the history and create another one.

### Procedure

1. In the navigation pane, expand the **Configuration > Serviceability > RMON** folders.
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 the system collects and places historical data 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 the system save data in the part of the media-specific table associated with

*Table continues...*

Name	Description
	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 the system save data 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, the system adds a new bucket to the media-specific table. After the number of buckets reaches the value of this object and the system is going to add a new bucket to the media-specific table, the agent deletes the oldest bucket associated with this entry so the system can added the new bucket. 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 system allows the number of associated media-specific entries to grow.
<b>Interval</b>	Specifies the interval in seconds over which the system samples data 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.
<b>Owner</b>	Specifies the entity that configured this entry and uses the assigned resources.

---

## Disabling RMON1 History

### About this task

Disable RMON1 history on a port if you do not want to record a statistical sample from that port.



**Procedure**

1. In the navigation pane, expand the **Configuration > Serviceability > RMON** folders.
2. Click **Control**.
3. In the **History** tab, select the row that contains the port ID to delete.
4. Click **Delete**.

## Viewing RMON1 History Statistics

View RMON1 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 pane, expand the **Configuration > Graph** folders.
3. Click **Port**.
4. Click the **RMON History** tab.
5. Select the statistics you want to graph.
6. Click the button for the type of graph you require (bar, pie, chart, or line).

## RMON History Field Descriptions

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

**Table 10: Variable definitions**

Parameter	Description
<b>SampleIndex</b>	Identifies the particular sample this entry represents among all samples associated with the same history control entry. This index starts at one and increases by one as each new sample is taken.
<b>Utilization</b>	Specifies 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>	Specifies 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>	Specifies the number of packets (including bad packets) received during this sampling interval.
<b>BroadcastPkts</b>	Specifies the number of good packets received during this sampling interval that were directed to the broadcast address.
<b>MulticastPkts</b>	Specifies the number of good packets received during this sampling interval that the system directs to a multicast address. This number does not include packets addressed to the broadcast address.

*Table continues...*

Parameter	Description
<b>DropEvents</b>	Specifies the total number of events in which the probe dropped packets 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 the system detects this condition.
<b>CRCAlignErrors</b>	The number of packets the system receives 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>	Specifies the number of packets the system receives 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>	Specifies the number of packets the system receives 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>	<p>Specifies 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).</p> <p>It is entirely normal for Fragments to increment because it counts both runts (which are normal occurrences due to collisions) and noise hits.</p>
<b>Collisions</b>	<p>Specifies 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>An RMON probe inside a repeater can ideally report collisions between the repeater and one or more other hosts (transmit collisions as defined by IEEE 802.3k) plus receiver collisions observed on any coax segments to which the repeater is connected.</p>

## Creating an RMON1 Alarm

After you enable RMON1 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 entry.

### Before you begin

- You must globally enable RMON.

### Procedure

- In the navigation pane, expand the **Configuration > Serviceability > RMON** folders.
- Click **Alarms**.
- Click the **Alarms** tab.
- Click **Insert**.
- 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.
- In the **SampleType** option, select a sample type.
- In the **Interval** box, type a sample interval in seconds.
- In the **Index** box, type an index number.
- In the **RisingThreshold** box, type a rising threshold value.
- In the **RisingEventIndex** box, type a rising threshold event index.
- In the **FallingThreshold** box, type a falling threshold value.
- In the **FallingEventIndex** box, type a falling threshold event index.
- In the **Owner** box, type the owner of the alarm.
- 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—Configures 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.

*Table continues...*

Name	Description
<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.</p> <p>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), 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 the system articulates SNMP access control 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 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, the system returns a badValue error. After the variable name of an established alarmEntry is no longer available in the selected MIB view, the probe changes 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>	<p>Specifies the method of sampling the selected variable and calculating the value to be compared against the thresholds. If the value of this object is absoluteValue, the value of the system compares the selected variable directly with the thresholds at the end of the sampling interval. If the value of this object is deltaValue, the system subtracts the value of the selected variable at the last sample from the current value, and the system compares the difference with the thresholds. You cannot modify this object if the associated alarmStatus object is equal to valid. The default is deltaValue.</p>
<b>Value</b>	<p>Specifies the value of the statistic during the last sampling period. For example, if the sample type is deltaValue, this value is the difference between the samples at the beginning and end of the period. If the sample type is absoluteValue, this value is the sampled value at the end of the period. This system compares the value with the rising and falling thresholds. The value during the current sampling period is not made available until the period is completed and remains available until the next period is complete.</p>
<b>StartupAlarm</b>	<p>Specifies the alarm that is sent after this entry is first set to valid. If the first sample after this entry becomes valid is greater than or equal to the risingThreshold and alarmStartupAlarm is equal to the risingAlarm or the risingOrFallingAlarm, then the system generates a single rising alarm. If the first sample after this entry becomes valid is less than or equal to the</p>

*Table continues...*

Name	Description
	fallingThreshold and alarmStartupAlarm is equal to the fallingAlarm or the risingOrFallingAlarm, then the system generates a single falling alarm. You cannot modify this object if the associated alarmStatus object is equal to valid.
<b>RisingThreshold</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, the system generates a single event. The system also generates a single event 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>	<p>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 no corresponding entry exists in the eventTable, no association exists. In particular, if this value is zero, the system generates no associated event, as zero is not a valid event index. You cannot modify this object if the associated alarmStatus object is equal to valid.</p> <p> <b>Note:</b> You must create the event prior to associating it to an alarm.</p>
<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, the system generates a single event. The system also generates a single event if the first sample after this entry becomes valid is less than or equal to this threshold and the associated alarmStartupAlarm is equal to fallingAlarm or risingOrFallingAlarm. After the system generates a falling event, the system does not generate another similar event 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>	<p>Specifies the index of the eventEntry that the system uses 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, the system generates no associated event, as zero is not a valid event index. You cannot modify this object if the associated alarmStatus object is equal to valid.</p> <p> <b>Note:</b> You must create the event prior to associating it to an alarm.</p>
<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.

---

## Viewing RMON1 Alarms

View the RMON1 alarm information to see alarm activity.

### Procedure

1. In the navigation pane, expand the **Configuration > Serviceability > RMON** folders.
2. Click **Alarms**.
3. Click the **Alarm** tab.

---

## Deleting an RMON1 Alarm

Delete an RMON1 alarm if you no longer want it to appear in the log.

### Procedure

1. In the navigation pane, expand the **Configuration > Serviceability > RMON** folders.
2. Click **Alarms**.
3. Select the alarm you must delete.
4. Click **Delete**.

---

## Creating an RMON1 Event

Create a custom rising and falling RMON1 event to specify if alarm information is sent to a trap, a log, or both.

### Procedure

1. In the navigation pane, expand the **Configuration > Serviceability > RMON** folders.
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 entry defines one event that the system generates after the appropriate conditions occur. The default is 1.
<b>Description</b>	Specifies a comment that describes this event entry.
<b>Type</b>	Specifies the type of notification that the probe makes about this event. In the case of a log, the system makes an entry in the log table for each event. In the case of SNMP traps, the system sends an SNMP trap to one or more management stations.
<b>Community</b>	Specifies the SNMP community 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.

---

## Viewing RMON1 Events

View RMON1 events to see how many events occurred.

### Procedure

1. In the navigation pane, expand the **Configuration > Serviceability > RMON** folders.
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 entry defines one event that the system generates after the appropriate conditions occur. The default is 1.
<b>Description</b>	Specifies a comment that describes this event entry.
<b>Type</b>	Specifies the type of notification that the probe makes about this event. In the case of a log, the system makes an entry in the log table for each event. In the case of SNMP traps, the system sends an SNMP trap to one or more management stations.

*Table continues...*

Name	Description
<b>Community</b>	Specifies the SNMP community 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.

---

## Deleting an Event

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

### Procedure

1. In the navigation pane, expand the **Configuration > Serviceability > RMON** folders.
2. Click **Alarms**.
3. Click the **Events** tab.
4. Select the event you must delete.
5. Click **Delete**.

---

## Viewing the RMON Log

### About this task

View the trap log to see which activity occurred.

### Procedure

1. In the navigation pane, expand the **Configuration > Serviceability > RMON** folders.
2. Click **Alarms**.
3. Click the **Log** tab.

## Log Field Descriptions

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

Name	Description
<b>EventIndex</b>	Specifies an index that uniquely identifies an entry in the event table. Each entry defines one event that is generated under appropriate conditions.

*Table continues...*



Name	Description
<b>Index</b>	Specifies an index that uniquely identifies an entry in the log table generated by the same event entries.
<b>Time</b>	Specifies the creation time for this log entry.
<b>Description</b>	Specifies an implementation dependent description of the event that activated this log entry.

## Viewing the Protocol Directory

View the protocol directory to see the list of protocols that RMON2 can monitor. You cannot change the list of protocols.

### About this task

The protocol directory MIB is enabled by default for the predefined protocols.

### Procedure

1. In the navigation pane, expand the **Configuration > Serviceability > RMON** folders.
2. Click **Protocol Directory**.
3. Click the **Protocol Directories** tab.

## Protocol Directories Field Descriptions

Use the data in the following table to use the **Protocol Directories** tab.

Name	Description
<b>Index</b>	Shows a unique identifier for the entry in the table.
<b>Protocol</b>	Shows the protocols RMON2 can monitor: <ul style="list-style-type: none"> <li>• Internet Protocol (IP)</li> <li>• Secure Shell version 2 (SSHv2)</li> <li>• Transmission Control Protocol (TCP)</li> <li>• User Datagram Protocol (UDP)</li> <li>• File Transfer Protocol (FTP)</li> <li>• Hypertext Transfer Protocol (HTTP)</li> <li>• Telnet</li> <li>• Remote login (rlogin)</li> <li>• Trivial File Transfer Protocol (TFTP)</li> <li>• Simple Networking Management Protocol (SNMP)</li> </ul>

*Table continues...*

Name	Description
<b>AddressMapConfig</b>	<p>Describes and configures the probe support for the network layer and application layer host tables for this protocol. The value can be one of the following:</p> <ul style="list-style-type: none"> <li>• notSupported</li> <li>• supportedOff</li> <li>• supportedOn</li> </ul> <p>If the value is supportedOn, the probe adds entries to the Address Map tab that maps the network layer address to the MAC layer address.</p>
<b>HostConfig</b>	<p>Describes and configures the probe support for the network layer and application layer host tables for this protocol. The value can be one of the following:</p> <ul style="list-style-type: none"> <li>• notSupported</li> <li>• supportedOff</li> <li>• supportedOn</li> </ul> <p>If the value is supportedOn, the probe adds entries to the Host Control tab to collect statistics for network layer and application layer hosts.</p>
<b>MatrixConfig</b>	<p>Describes and configures the probe support for the network layer and application layer host tables for this protocol. The value can be one of the following:</p> <ul style="list-style-type: none"> <li>• notSupported</li> <li>• supportedOff</li> <li>• supportedOn</li> </ul>
<b>Owner</b>	Shows the entity that configured this entry.

---

## Viewing the Data Source for Protocol Distribution Statistics

View the Distribution Control tab to see the network segment data source on which the protocol distribution statistics are measured. The management IP mentioned as a data source represents the IP that the SNMP agent uses to access the switch.

### Procedure

1. In the navigation pane, expand the **Configuration > Serviceability > RMON** folders.
2. Click **Protocol Distribution**.
3. Click the **Distribution Control** tab.

## Distribution Control Field Descriptions

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

Name	Description
<b>Index</b>	Shows a unique identifier for the entry in the table.
<b>DataSource</b>	Specifies the source of data for this protocol distribution.
<b>DroppedFrames</b>	Shows the total number of frames that the probe receives and drops but does not include in the StatsDropEvents value. This event can occur if the probe is out of resources and sheds the load from this collection. This value does not include packets that were not counted because they had MAC-layer errors.
<b>CreateTime</b>	Shows the value of the sysUpTime when the entry was last activated.
<b>Owner</b>	Shows the entity that configured this entry.

---

## Viewing Protocol Distribution Statistics

View protocol distribution statistics to see traffic statistics that each protocol generates by local area network (LAN) segment.

### Procedure

1. In the navigation pane, expand the **Configuration > Serviceability > RMON** folders.
2. Click **Protocol Distribution**.
3. Click the **Distribution Stats** tab.

## Distribution Stats Field Descriptions

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

Name	Description
<b>LocalIndex</b>	Identifies the protocol distribution an entry is part of, as well as the particular protocol that it represents.
<b>Pkts</b>	Shows the number of packets without errors received for this protocol type. This value is the number of link-layer packets so a single, fragmented network-layer packet can increment the counter several times.
<b>Octets</b>	Shows the number of octets in packets received for this protocol type since it was added to the table.

*Table continues...*

Name	Description
	This value does not include octets in packets that contained errors. This value counts octets in the entire packet that contained the protocol, not just the particular protocol frames.

---

## Viewing the Host Interfaces Enabled for Monitoring

View the entries in the address map control tab to see which host interfaces are enabled for monitoring on the switch. Each entry in this table enables the discovery of addresses on a new interface.

### Procedure

1. In the navigation pane, expand the **Configuration > Serviceability > RMON** folders.
2. Click **Address Map**.
3. Click the **Address Map Control** tab.

## Address Map Control Field Descriptions

Use the data in the following table to use the **Address Map Control** tab.

Name	Description
<b>Index</b>	Shows a unique identifier for the entry in the table.
<b>DataSource</b>	Shows the source of data for the entry.
<b>DroppedFrames</b>	Shows the total number of frames that the probe receives and drops but does not include in the StatsDropEvents value. This event can occur if the probe is out of resources and sheds the load from this collection. This value does not include packets that were not counted because they had MAC-layer errors.
<b>Owner</b>	Shows the entity that configured this entry.

---

## Viewing Address Mappings

View the mappings of network layer address to physical address to interface.

### About this task

The probe adds entries on this tab based on the source MAC and network addresses in packets without MAC-level errors.

The probe populates this table for all protocols on the **Protocol Directories** tab with a value of **AddressMapConfig** equal to **supportedOn**.

## Procedure

1. In the navigation pane, expand the **Configuration > Serviceability > RMON** folders.
2. Click **Address Map**.
3. Click the **Address Map** tab.

## Address Map Field Descriptions

Use the data in the following table to use the **Address Map** tab.

Name	Description
<b>LocalIndex</b>	Shows a unique identifier for the entry in the table.
<b>HostAddress</b>	Shows the network address for this entry. The format of the value depends on the protocol portion of the local index.
<b>Source</b>	Shows the interface or port on which the network address was most recently seen.
<b>PhysicalAddress</b>	Shows the physical address on which the network address was most recently seen.
<b>LastChange</b>	Shows the value of the sysUpTime when the entry was created or last changed. If this value changes frequently, it can indicate duplicate address problems.

---

## Viewing the Data Source for Host Statistics

View the Host Control tab to see the data source for both network layer and application layer host statistics.

### Procedure

1. In the navigation pane, expand the **Configuration > Serviceability > RMON** folders.
2. Click **Network Layer Host**.
3. Click the **Host Control** tab.

## Host Control Field Descriptions

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

Name	Description
<b>Index</b>	Shows a unique identifier for the entry in the table.
<b>DataSource</b>	Shows the source of data for the associated host table. The statistics in this group reflect all packets

*Table continues...*

Name	Description
	on the local network segment that attaches to the identified interface.
<b>NHDropFrames</b>	Shows the total number of frames that the probe receives and drops but does not include in the StatsDropEvents value. This event can occur if the probe is out of resources and sheds the load from this collection. This value does not include packets that were not counted because they had MAC-layer errors.
<b>AHDropFrames</b>	Shows the total number of frames that the probe receives and drops but does not include in the StatsDropEvents value. This event can occur if the probe is out of resources and sheds the load from this collection. This value does not include packets that were not counted because they had MAC-layer errors.
<b>Owner</b>	Shows the entity that configured this entry.

## Viewing Network Host Statistics

View network host statistics to see Layer 3 traffic statistics for each host. The network layer host MIB monitors traffic packets in and out of hosts based on the network layer address.

### Procedure

1. In the navigation pane, expand the **Configuration > Serviceability > RMON** folders.
2. Click **Network Layer Host**.
3. Click the **Network Host Stats** tab.

## Network Host Stats Field Descriptions

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

Name	Description
<b>LocalIndex</b>	Shows a unique identifier for the entry in the table.
<b>HostAddress</b>	Shows the host address for this entry.
<b>InPkts</b>	Shows the number of packets without errors transmitted to this address. This value is the number of link-layer packets so a single, fragmented network-layer packet can increment the counter several times.
<b>OutPkts</b>	Shows the number of packets without errors transmitted by this address. This value is the number

*Table continues...*

Name	Description
	of link-layer packets so a single, fragmented network-layer packet can increment the counter several times.
<b>InOctets</b>	Shows the number of octets transmitted to this address, excluding octets in packets that contained errors. This value counts octets in the entire packet that contained the protocol, not just the particular protocol frames.
<b>OutOctets</b>	Shows the number of octets transmitted by this address, excluding octets in packets that contained errors. This value counts octets in the entire packet that contained the protocol, not just the particular protocol frames.
<b>CreateTime</b>	Shows the value of the sysUpTime when the entry was last activated.

## Viewing Application Host Statistics

View application host statistics to see traffic statistics by application protocol for each host.

### Procedure

1. In the navigation pane, expand the **Configuration > Serviceability > RMON** folders.
2. Click **Application Layer Host**.
3. Click the **Application Host Stats** tab.

## Application Host Stats Field Descriptions

Use the data in the following table to use the Application Host Stats tab.

Name	Description
<b>Index</b>	Shows a unique identifier for the entry in the table.
<b>HostAddress</b>	Identifies the network layer address of this entry.
<b>LocalIndex</b>	Identifies the network layer protocol of the address.
<b>InPkts</b>	Shows the number of packets for this protocol type, without errors, transmitted to this address. This value is the number of link-layer packets so a single, fragmented network-layer packet can increment the counter several times.
<b>OutPkts</b>	Shows the number of packets for this protocol type, without errors, transmitted by this address. This value is the number of link-layer packets so a single,

*Table continues...*

Name	Description
	fragmented network-layer packet can increment the counter several times.
<b>InOctets</b>	Shows the number of octets transmitted to this address, excluding octets in packets that contained errors. This value counts octets in the entire packet that contained the protocol, not just the particular protocol frames.
<b>OutOctets</b>	Shows the number of octets transmitted by this address, excluding octets in packets that contained errors. This value counts octets in the entire packet that contained the protocol, not just the particular protocol frames.
<b>CreateTime</b>	Shows the value of the sysUpTime when the entry was last activated.

## 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 11: RMON alarm variables**

Category	Subcategory	Variable	Definition
Security		rcCliNumAccessViolations.0	The number of CLI access violations detected by the system.
		rcWebNumAccessBlocks.0	The number of accesses the Web server blocked.
		snmpInBadCommunityNames.0	The total number of SNMP messages delivered to the SNMP protocol entity that represented an SNMP operation not allowed 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.

*Table continues...*



Category	Subcategory	Variable	Definition
		ifInErrors	For packet-oriented interfaces, the number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol. For character-oriented or fixed-length interfaces, the number of inbound transmission units that contained errors, preventing them from being deliverable to a higher-layer protocol.
		ifOutDiscards	The number of outbound packets discarded even though no errors were detected to prevent the packets being transmitted. One possible reason for discarding such a packet is to free buffer space.
		ifOutErrors	For packet-oriented interfaces, the number of outbound packets that were not transmitted because of errors. For character-oriented or fixed-length interfaces, the number of outbound transmission units that were not transmitted because of errors.
	Ethernet	dot3StatsAlignmentErrors	A count of frames received on a particular interface that are not an integral number of octets in length and do not pass the Frame Check Sequence (FCS) check. The count represented by an instance of this object increments when the alignmentError status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions exist are, according to the conventions of IEEE 802.3 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.

*Table continues...*

Category	Subcategory	Variable	Definition
			<p>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.</p>
		dot3StatsSingleCollisionFrames	<p>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.</p>
		dot3StatsMultipleCollisionFrames	<p>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.</p>
		dot3StatsSQETestErrors	<p>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</p>

*Table continues...*

Category	Subcategory	Variable	Definition
			defined in section 7.2.2.2.4 of ANSI/IEEE 802.3-1985 and its generation is described in section 7.2.4.6 of the same document.
		dot3StatsDeferredTransmissions	A count of frames for which the first transmission attempt on a particular interface is delayed because the medium is busy. The count represented by an instance of this object does not include frames involved in collisions.
		dot3StatsLateCollisions	The number of times that a collision is detected on a particular interface later than 512 bit-times into the transmission of a packet; 512 bit-times corresponds to 51.2 microseconds on a 10 Mb/s system. A (late) collision included in a count represented by an instance of this object is also considered as a (generic) collision for purposes of other collision-related statistics.
		dot3StatsExcessiveCollisions	A count of frames where the transmission on a particular interface fails due to excessive collisions.
		dot3StatsInternalMacTransmitErrors	<p>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.</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</p>

*Table continues...*

Category	Subcategory	Variable	Definition
			transmission errors on a particular interface that are not otherwise counted.
		dot3StatsCarrierSenseErrors	The number of times the carrier sense condition was lost or never asserted when the switch attempted to transmit a frame on a particular interface. The count represented by an instance of this object increments at most once for each transmission attempt, even if the carrier sense condition fluctuates during a transmission attempt.
		dot3StatsFrameTooLongs	A count of frames received on a particular interface that exceeds the maximum permitted frame size. The count represented by an instance of this object increments when the frameTooLong status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions obtained are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.
		dot3StatsInternalMacReceiveErrors	<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</p>

*Table continues...*

Category	Subcategory	Variable	Definition
			object can represent a count of transmission errors on a particular interface that are not otherwise counted.
	IP	ipInHdrErrors.0	The number of input datagrams discarded due to errors in the datagram IP headers, including bad checksums, version number mismatch, other format errors, time-to-live exceeded, and errors discovered in processing IP options.
		ipInDiscards.0	The number of discarded input IP datagrams where no problems were encountered to prevent continued processing. An example of why they were discarded can be lack of buffer space. This counter does not include any datagrams discarded while awaiting reassembly.
		ipOutDiscards.0	The number of output IP datagrams where no problems were encountered to prevent transmission to the destination, but that were discarded (for example, for lack of buffer space). This counter includes datagrams counted in ipForwDatagrams if packets meet this (discretionary) discard criterion.
		ipFragFails.0	The number of IP datagrams discarded because they needed to be fragmented at this entity but were not, for example, because the Don't Fragment flag was set.
		ipReasmFails.0	The number of failures detected by the IP reassembly algorithm (for whatever reason: timed out, errors, and so forth). This is not necessarily a count of discarded IP fragments because some algorithms (notably the algorithm in RFC 815) can lose track of the number of fragments by

*Table continues...*

Category	Subcategory	Variable	Definition
			combining them as they are received.
		icmpInParmProbs.0	The number of ICMP In parameter problem messages received.
		icmpOutParmProbs.0	The number of ICMP Out parameter problem messages received.
	MLT	rcStatMltEtherAlignmentErrors	The number of frames received on an MLT that are not an integral number of octets in length, but do not pass the FCS check.
		rcStatMltEtherFCSErrors	The number of frames received on an MLT that are an integral number of octets in length, but do not pass the FCS check.
		rcStatMltEtherSingleCollFrames	The number of successfully transmitted frames on a particular MLT where transmission is inhibited by exactly one collision.
		rcStatMltEtherMultipleCollFrames	The number of successfully transmitted frames on a particular MLT where transmission is inhibited by more than one collision.
		rcStatMltEtherSQETestError	A count of times that the SQE TEST ERROR message is generated by the PLS sublayer for a particular MLT.
		rcStatMltEtherDeferredTransmiss	A count of frames where the first transmission attempt on a particular MLT is delayed because the medium is busy. The count represented by an instance of this object.
		rcStatMltEtherLateCollisions	The number of times that a late collision is detected on a particular MLT later than 512 bit-times into the transmission of a packet; 512-bit-times corresponds to 51.2-microseconds on a 10 Mb/s system.
		rcStatMltEtherExcessiveCollis	The number of times that excessive collisions are detected

*Table continues...*

Category	Subcategory	Variable	Definition
			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 object, the AlignmentErrors object, or the FCSErrors object.

*Table continues...*

Category	Subcategory	Variable	Definition
	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.
		rcStgPortInBadBpdus	The number of bad BPDUs received by this port.
		dot1dTpPortInDiscards	Count of valid frames received that were discarded (that is, filtered) by the forwarding process.
Traffic	Interface	ifInOctets	The total number of octets received on the interface, including framing characters.
		ifInMulticastPkts	The number of packets, delivered by this sublayer to a higher sublayer, that are addressed to a multicast address at this sublayer. For a MAC layer protocol, this number includes both Group and Functional addresses.
		ifInBroadcastPkts	The number of packets, delivered by this sublayer to a higher (sub) layer, that are addressed to a broadcast address at this sublayer.
		ifInUnkownProtos	For packet-oriented interfaces, the number of packets received through the interface that are discarded because of an unknown or unsupported protocol. For character-oriented or fixed-length interfaces that support protocol multiplexing, the number of transmission units received 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.

*Table continues...*



Category	Subcategory	Variable	Definition
		ifOutOctets	The total number of octets transmitted from the interface, including framing characters.
		ifOutMulticastPkts	The total number of packets that higher-level protocols requested be transmitted, and that are addressed to a multicast address at this sublayer, including those that are discarded or not sent. For a MAC layer protocol, this includes both Group and Functional addresses.
		ifoutBroadcastPkts	The total number of packets that higher level protocols requested transmitted, and that were addressed to a broadcast address at this sublayer, including those discarded or not sent.
		ifLastChange	The value of sysUpTime at the time the interface entered its current operational state. If the current state was entered prior to the last reinitialization of the local network management subsystem, this object contains a value of zero.
	RmonEther Stats	etherStatsOctets	The total number of octets of data (including those in bad packets) received on the network (excluding framing bits but including FCS octets). Use this object as a reasonable estimate of Ethernet utilization. For greater precision, sample the etherStatsPkts and etherStatsOctets objects before and after a common interval.
		etherStatsPkts	The total number of packets (including bad packets, broadcast packets, and multicast packets) received.
		etherStatsBroadcastPkts	The total number of good packets received that are directed to the broadcast address. This number does not include multicast packets.

*Table continues...*

Category	Subcategory	Variable	Definition
		etherStatsMulticastPkts	The total number of good packets received that are directed to a multicast address. This number does not include packets directed to the broadcast address.
		etherStatsCRCAAlignErrors	The total number of packets received that had a length (excluding framing bits, but including FCS octets) of 64 to 1518 octets, inclusive, but had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).
		etherStatsUndersizePkts	The total number of packets received that are less than 64 octets long (excluding framing bits, but including FCS octets) and were otherwise well formed.
		etherStatsOversizePkts	The total number of packets received that are longer than 1518 octets (excluding framing bits, but including FCS octets) and were otherwise well formed.
		etherStatsFragments	<p>The total number of packets received that are less than 64 octets in length (excluding framing bits but including FCS octets) and had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).</p> <p>It is entirely normal for etherStatsFragments to increment because it counts both runts (which are normal occurrences due to collisions) and noise hits.</p>
		etherStatsCollisions	The best estimate of the total number of collisions on this Ethernet segment.
	IP	ipInReceives.0	All incoming IP packets.

*Table continues...*

Category	Subcategory	Variable	Definition
		ipInAddrErrors.0	The number of bad IP destination addresses.
		ipForwDatagrams.0	IP packets forwarded.
		ipInUnknownProtos.0	Number of unsupported IP protocols.
		ipInDelivers.0	The number of IP In packets delivered.
		ipOutRequests.0	The total number of IP datagrams that local IP user protocols supplied to IP in request for transmission.
		ipOutNoRoutes.0	The number of IP datagrams discarded because no route was found to transmit to the destination.
		ipFragOKs.0	The number of IP datagrams successfully fragmented.
		ipFragCreates.0	The number of IP datagram fragments generated as a result of fragmentation.
		ipReasmReqds.0	The number of requests to reassemble fragments.
		ipReasmOKs.0	The number of fragments reassembled successfully.
	ICMP	icmplnSrcQuenchs.0	The number of ICMP Source Quench messages received.
		icmplnRedirects.0	The number of ICMP redirect messages.
		icmplnEchos.0	The number of ICMP Echo requests messages received.
		icmplnEchosReps.0	The number of ICMP Echo reply messages received.
		icmplnTimeStamps.0	The number of ICMP timestamp request messages received.
		icmplnTimeStampsReps.0	The number of ICMP timestamp reply messages received.
		icmplnAddrMasks.0	The number of ICMP mask request messages reviewed.
		icmplnAddrMasksReps.0	The number of ICMP mask reply messages reviewed.

*Table continues...*

Category	Subcategory	Variable	Definition
		icmpInDestUnreachs.0	The number of ICMP destinations unreachable messages received.
		icmpInTimeExcds.0	The number of ICMP Time Exceeded messages received.
		icmpOutSrcQuenchs.0	The number of ICMP Source Quench messages sent.
		icmpOutRedirects.0	The number of ICMP redirect messages sent.
		icmpOutEchos.0	The number of ICMP Echo request messages sent.
		icmpOutEchosReps.0	The number of ICMP Echo reply messages sent.
		icmpOutTimeStamps.0	The number of ICMP Timestamp request messages sent.
		icmpOutTimeStampsReps.0	The number of ICMP Timestamp reply messages sent.
		icmpOutAddrMasks.0	The number of ICMP Address mask messages sent.
		icmpOutAddrMasksReps.0	The number of ICMP Address mask reply messages sent.
		icmpOutDestUnreachs.0	The number of ICMP destination unreachable messages sent.
		icmpOutTimeExcds.0	The number of ICMP time exceeded messages sent.
	Snmp	snmpInPkts.0	The total number of messages delivered to the SNMP entity from the transport service.
		snmpOutPkts.0	The total number of SNMP messages passed from the SNMP protocol entity to the transport service.
		snmpInBadVersions.0	The total number of SNMP messages delivered to the SNMP protocol entity that were intended for an unsupported SNMP version.
		snmpInBadCommunityUses.0	The total number of SNMP messages delivered to the SNMP protocol entity that represented an SNMP operation that was not allowed by the SNMP community named in the message.

*Table continues...*

Category	Subcategory	Variable	Definition
		snmpInTooBigs.0	The total number of SNMP PDUs delivered to the SNMP protocol entity and for which the value of the error-status field is tooBig.
		snmpInNoSuchNames.0	The total number of SNMP PDUs delivered to the SNMP protocol entity and for which the value of the error-status field is noSuchName.
		snmpInBadValues. 0	The total number of SNMP PDUs received that were generated by the SNMP protocol entity and for which the value of the error-status field is badValue.
		snmpInReadOnlys.0	The total number of valid SNMP PDUs delivered to the SNMP protocol entity and for which the value of the error-status field is readOnly. It is a protocol error to generate an SNMP PDU that contains the value readOnly in the error-status field; as such, this object is provided as a means of detecting incorrect implementations of the SNMP.
		snmpInGenErrs.0	The total number of SNMP PDUs delivered to the SNMP protocol entity and for which the value of the error-status field is genErr.
		snmpInTotalReqVars.0	The total number of MIB objects retrieved successfully by the SNMP protocol entity as the result of receiving valid SNMP Get-Request and Get-Next PDUs.
		snmpInTotalSetVars.0	The total number of MIB objects altered successfully by the SNMP protocol entity as the result of receiving valid SNMP Set-Request PDUs.
		snmpInGetRequests.0	The total number of SNMP Get-Request PDUs accepted and processed by the SNMP protocol entity.
		snmpInGetNexts.0	The total number of SNMP Get-Next PDUs accepted and

*Table continues...*

Category	Subcategory	Variable	Definition
			processed by the SNMP protocol entity.
		snmpInSetRequests.0	The total number of SNMP Set-Request PDUs accepted and processed by the SNMP protocol entity.
		snmpInGetResponses.0	The total number of SNMP Get-Response PDUs accepted and processed by the SNMP protocol entity.
		snmpInTraps.0	The total number of SNMP Trap PDUs accepted and processed by the SNMP protocol entity.
		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.

*Table continues...*

Category	Subcategory	Variable	Definition
		snmpOutTraps.0	The total number of SNMP Trap PDUs generated by the SNMP protocol entity.
	Bridge	rcStgTimeSinceTopologyChange	The time (in hundredths of a second) since the last topology change was detected by the bridge entity.
		rcStgTopChanges	The total number of topology changes detected by this bridge since the management entity was last reset or initialized.
		rcStgMaxAge	The maximum age of Spanning Tree Protocol information learned from the network on any port before it is discarded, in hundredths of a second. This is the actual value that this bridge is currently using.
		rcStgPortForwardTransitions	The number of times this port transitioned from the Learning state to the Forwarding state.
		rcStgPortInConfigBpdus	The number of Config BPDUs received by this port.
		rcStgPortInTcnBpdus	The number of Topology Change Notification BPDUs received by this port.
		rcStgPortOutConfigBpdus	The number of Config BPDUs transmitted by this port.
		rcStgPortOutTcnBpdus	The number of Topology Change Notification BPDUs transmitted by this port.
		dot1dTpPortInFrames	The number of frames received by this port from its segment. A frame received on the interface corresponding to this port is counted by this object only if it is for a protocol being processed by the local bridging function, including bridge management frames.
		dot1dTpPortOutFrames	The number of frames transmitted by this port to its segment. A frame transmitted on the interface corresponding to this port is

*Table continues...*

Category	Subcategory	Variable	Definition
			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 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	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.
	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.

*Table continues...*



Category	Subcategory	Variable	Definition
		rcStatMltIfExtnIfHCInOctets	The total number of octets received on this MLT interface including framing characters detected by the high-count (64-bit) register.
		rcStatMltIfExtnIfHCInUcastPkts	The number of packets delivered by this MLT interface to a higher MLT that were not addressed to a multicast or broadcast address as detected by the high-count (64-bit) register.
		rcStatMltIfExtnIfHCInMulticastPkt	The total number of multicast packets delivered to this MLT interface detected by the high-count (64-bit) register.
		rcStatMltIfExtnIfHCInBroadcastPkt	The total number of broadcast packets delivered to this MLT interface detected by the high-count (64-bit) register.
		rcStatMltIfExtnIfHCOutOctets	The total number of octets transmitted from the MLT interface, including framing characters.
		rcStatMltIfExtnIfHCOutUcastPkts	The number of packets transmitted by this MLT interface to a higher MLT that were not addressed to a multicast or broadcast address as detected by the high-count (64-bit) register.
		rcStatMltIfExtnIfHCOutMulticast	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.
		rcStatMltIfExtnIfHCOutBroadcast	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.

 **Note:**

In addition to these elements that are offered in a graphical way by EDM, you can manually set any valid OID in the variable field to be monitored by an alarm. For these cases, the name of the variable cannot be translated automatically in OID, the exact OID must be set as a sequence of numbers.

# Chapter 10: sFlow

---

## sFlow Fundamentals

sFlow monitors traffic in a data network. Use sFlow to monitor routers and switches in the network, and capture traffic statistics about those devices. sFlow uses sampling to provide scalability for network-wide monitoring, and therefore applies to high speed networks. The switch sends the sampled data as a User Datagram Protocol (UDP) packet to the specified host and port.

sFlow consists of the following:

- sFlow agent—Performs two types of sampling:
  - Flow samples: Flow sampling randomly samples an average of 1 out of n packets for each operation.
  - Counter samples: Counter sampling periodically polls and exports counters for a configured interface. This type of sampling uses a counter to determine if the packet is sampled. Each packet that an interface receives, and that a filter does not drop, reduces the counter by one. After the counter reaches zero, the sFlow agent takes a sample.

 **Note:**

Only generic interface counters and Ethernet interface counters are supported.

- sFlow datagrams—Supports both flow samples and counter samples. Datagrams can be sent from the front panel port or an out-of-band (OOB) port. Each datagram provides information about the sFlow version, the originating IP address of the device, a sequence number, the number of samples it contains, and one or more flow and/or counter samples.
- sFlow collector—Located on a central server and runs software that analyzes and reports on network traffic. Two sFlow collectors can be configured to be reachable over a management network or Shortest Path Bridging (SPB). The preferred network is SPB.

### Limitations

- Application-specific integrated circuit (ASIC) or Software Development Kit (SDK) limitation—To avoid wobbling, the recommended counter interval for sFlow is 20 seconds. Minor wobbling can still occur even after configuring the recommended counter interval due to the interaction between the sFlow agent counter export schedule and the frequency with which the switch ASIC SDK copies and caches counters from the ASIC.
- sFlow supports a maximum of two collectors.
- UDP datagram size and the collector buffer are restricted to 1400 bytes. sFlow sends datagrams to the collector when the buffer reaches the 1400-byte capacity or after a timeout of one second is triggered. The collector buffer size cannot be modified.

- The switch supports IPv4 collector IP addresses.
- VLAN counters/statistics are not supported.
- sFlow can be enabled only on the front panel ports.
- You cannot configure the sampling limit. The sampling limit applies system-wide rather than on a per port basis. Sampling rates differ depending on the hardware platform so any sampled packets beyond the limit are dropped. For more information about feature support, see [Release Notes for VOSS](#).
- The switch does not support egress sampling. The switch supports only ingress sampling.
- The switch does not support enabling sFlow on a link aggregation group (LAG) interface. However, you can enable sFlow on the member interfaces of a LAG.
- The sFlow collector can be reachable through the Management VRF, the Global Routing Table (GRT) or if your switch supports doing so, through a user created VRF (virtual routing and forwarding). If the sFlow collector is hosted in either the GRT or a user created VRF, SPB reachability only supports using Layer 2 VSN or IP shortcuts to access the collector. Layer 3 VSNs are not supported in accessing the collector when it is hosted in the GRT or a User created VRF.
- **\* Note:**  
This restriction applies to the VSP 8600 only. Other platforms mirror copies to both destinations.

A packet can have only one mirror destination so you cannot configure sFlow and Port Mirroring on the same port.

### Configuration considerations

- If the sFlow collector has two network interface controller (NIC) cards, to avoid dropped sFlow datagrams that are a result of reverse path checks, you can add a route to the agent-ip address for the NIC card on which the sFlow datagrams are received.
- First preference is always given to either the GRT or management VRF to where the sFlow agent IP address is configured. For example, if you configure the sFlow agent IP address as part of GRT, the GRT route to the collector is given preference over the management VRF. If the management network hosts a collector with a collector IP address that is reachable over SPB as a result of redistributing direct routes on a peer Backbone Edge Bridge (BEB) or in situations where the GRT has a default route (0.0.0.0) and the collector route is in the local management VRF, first preference is given to the VRF where you have configured the sFlow agent IP address.

### Example

After you configure the sFlow agent on the network device that you want to monitor, the system collects flow samples or counter samples, and exports these traffic statistics as sFlow datagrams to the sFlow collector on a server or appliance.

For example, after the buffers reach capacity or a timeout is triggered, an sFlow datagram, which is a UDP packet, sends the measurement information to the sFlow collector buffers. The UDP payload contains the sFlow datagram.

The following figure shows the sFlow agent on various routers and switches with sFlow datagrams being sent to the sFlow collector.




Table 12: sFlow legend

Number	Description
1	sFlow collector
2	sFlow datagrams
3	sFlow agents

As a general rule, drop action occurs after sampling completes. However, in situations related to Layer 1 errors such as, MTU exceeded packets, the drop action occurs before sampling begins. For errors such as, frame too long, packets are dropped due to the size of the frame being greater than the interface MTU. In this situation, the packets are dropped before sampling begins so only counter polling occurs. To enable trace, use `line-card 1 trace level 232 <0-4>`.

**! Important:**

The defined sampling rate, an average of 1 out of n packets/operations does not provide a 100% accurate result, but it does provide a result with quantifiable accuracy.

## sFlow Configuration Using CLI

Use sFlow to capture traffic statistics to monitor traffic in a data network. This section provides procedures to view and configure sFlow using CLI.

### Configuring the agent-ip and Enabling sFlow Globally

Configure the sFlow agent IPv4 address, and then enable sFlow before the system can monitor and capture traffic statistics to send to an sFlow collector. By default, sFlow is globally disabled.

#### Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enable the agent IPv4 address:

```
sflow agent-ip {A.B.C.D}
```

3. Enable sFlow:

```
sflow enable
```

4. Verify the global configuration:

```
show sflow
```

#### Example

Globally enable sFlow, and then verify the configuration.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#sflow agent-ip 192.0.2.27
INFO: Agent IP address must exist in MGMT or GRT vrf
Switch:1(config)#sflow enable
Switch:1(config)#show sflow
=====
                        sFlow Global Configuration
=====
Global State           : Enabled
Agent IP               : 192.0.2.27
```

#### Next steps

After you configure the agent-ip and globally enable sFlow, proceed to configuring the sFlow collector.

### Variable Definitions

Use the data in the following table to use the `sflow agent-ip` command.

Variable	Definition
{A.B.C.D.}	Specifies the agent-ip address (IPv4).

## Configuring an sFlow Collector

Configure an sFlow collector to determine the device to which the sFlow agent sends sFlow datagrams. You can configure up to two collectors for each interface slot in the chassis.

### Before you begin

- You must globally enable sFlow.

### About this task

The sFlow datagrams that the agent sends to the collector are not encrypted. Use a VLAN to create a secure measurement network to route sFlow datagrams.

To further protect the sFlow collector, configure it to accept only sFlow datagrams, or to check sequence numbers and verify source addresses.

### Procedure

- Enter Global Configuration mode:

```
enable
configure terminal
```

- Configure the collector information:

```
sflow collector <1-2> address {A.B.C.D} [Owner WORD <1-20>] [port
<1-65535>] [timeout <1-65535>]
```

- Verify the collector configuration:

```
show sflow collector <1-2>
```

### Example

Configure collector ID, and then verify the configuration.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#sflow collector 1 address 192.0.2.26 owner sflow1 port 6343 timeout 497
Switch:1(config)#sflow collector 2 address 192.0.2.27 owner sflow2 port 6343 timeout 531
Switch:1(config)#show sflow collector
```


#### sFlow Collector Configuration Info

Id	Owner	Collector-IP	Port	Timeout(secs)	Reachable via
1	sflow1	192.0.2.26	6343	497	192.0.2.15
2	sflow2	192.0.2.27	6343	531	192.0.2.16

All 2 out of 2 Total Num of sflow collector entries displayed

## Variable Definitions

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

Variable	Value
collector <1-2>	<p>Specifies the id to export sFlow datagrams to the collector id.</p> <p>Use the no operator to remove an sflow collector id and a collector name. <code>no sflow collector &lt;1-2&gt; owner WORD&lt;1-20&gt;</code></p> <p>To configure the default value, enter <code>default sflow collector &lt;1-2&gt;</code></p>
address {A.B.C.D.}	<p>Specifies the collector IP address.</p> <p>Use the no operator to remove an sflow collector address. <code>no sflow collector &lt;1-2&gt; address {A.B.C.D.}</code></p>
owner WORD<1-20>	<p>Specifies the sFlow collector name.</p>
port <1-65535>	<p>Specifies the destination UDP port. The default port is 6343.</p> <p>To configure the default value, enter <code>default sflow collector &lt;1-2&gt; port</code></p>
timeout <1-65535>	<p>Specifies the time remaining (in seconds) before the collector is released.</p> <p>The default timeout is 0, which means the timeout is not used and the collector sends data forever.</p> <p>To configure the default value, enter <code>default sflow collector &lt;1-2&gt; timeout</code></p>
vrf WORD<1-16>	<p>Specifies the name of the VRF used to reach the collector.</p> <p> <b>Note:</b> This parameter is not supported on all hardware platforms.</p>

## Configuring the Packet Sampling Rate

Configure the packet sampling rate at port level to determine how many packets the system counts before it takes a sample.

### Before you begin

- You must globally enable sFlow.



## About this task

If you configure a conservative sampling rate to prevent overloading the sFlow agent, the result will reflect high values that do not reflect typical traffic levels.

## Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable

configure terminal

interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}
```

**\* Note:**

If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the collector id:

```
sflow collector <1-2>
```

3. Configure the sampling rate:

```
sflow sampling-rate <8192-1000000>
```

4. Verify the configuration:

```
show sflow interface {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}
```

## Example

Configure sampling rates for ports 2/1, 2/2, 2/3, and 2/4.

```
Switch:1(config-if)#interface gigabitethernet 2/1,2/2
Switch:1(config-if)#sflow collector 1
Switch:1(config-if)#sflow sampling-rate 10000
Switch:1(config-if)#interface gigabitethernet 2/3
Switch:1(config-if)#sflow collector 2
Switch:1(config-if)#sflow sampling-rate 8192
Switch:1(config-if)#interface gigabitethernet 2/4
Switch:1(config-if)#sflow collector 2
Switch:1(config-if)#sflow sampling-rate 12001
Switch:1(config-if)#show sflow interface enabled
```

```
=====
                          sFlow Port Configuration Info
=====
```

Port	Packet-Sample-Rate	Max-Header-Size	Counter-interval (in secs)	Collector-list
2/1	10000	128	0	1
2/2	10000	128	0	1
2/3	8192	128	0	2
2/4	12001	128	0	2

```
-----
All 4 out of 4 Total Num of sflow port entries displayed
```

## Variable Definitions

Use the data in the following table to use the `sflow sampling-rate` and `show sflow interface` commands.

Variable	Value
<8192–1000000>	Configures the packet sampling rate on a port. The default value is 0 (disabled). To configure the default value, enter <code>default sflow sampling-rate</code> .
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

---

## Configuring sFlow Maximum Header Size

Configure the maximum header size on a single port or multiple ports.

### Before you begin

- You must globally enable sFlow.

### Procedure

1. Enter GigabitEthernet Interface Configuration mode:

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

**\* Note:**

If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the maximum-header size:

```
sflow max-header-size <64-256>
```

3. Verify the configuration:

```
show sflow interface {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}
```

## Example

For ports 1/1 to 1/10, configure the maximum header size, and then verify the configuration.

```
Switch:1(config-if)#interface gigabitethernet 1/1-1/10
Switch:1(config-if)#sflow max-header-size 255
Switch:1(config-if)#show sflow interface 1/1-1/10
```

```
=====
                          sFlow Port Configuration Info
=====
Port      Packet-Sample-Rate  Max-Header-Size  Counter-interval  Collector-list
              (in secs)
-----
1/1       0                   255              525               1,2
1/2       0                   255              525               1,2
1/3       0                   255              525               1,2
1/4       0                   255              525               1,2
1/5       0                   255              525               1,2
1/6       0                   255              525               1,2
1/7       0                   255              525               1,2
1/8       0                   255              525               1,2
1/9       0                   255              525               1,2
1/10      0                   255              525               1,2
```

## Variable Definitions

Use the data in the following table to use the `max-header-size` command.

Variable	Value
<64–256>	Identifies the maximum number of bytes to be copied from the sampled packet. Default 128 bytes.

## Configuring the Counter Sampling Interval

Configure the counter sampling interval values at port level to determine how often the sFlow agent polls and exports counters for a configured interface.

### Before you begin

- You must globally enable sFlow.

### Procedure

- Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-
port]][,...]}
```

**\* Note:**

If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the counter sampling interval:

```
sflow counter-interval <1-3600>
```

3. Verify the configuration:

```
show sflow interface {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...]}
```

### Example

Verify all slots use the default polling-interval configuration.

```
Switch:1(config-if)#sflow counter-interval 525
Switch:1(config-if)#show sflow interface 1/1-1/10
```

```
=====
                        sFlow Port Configuration Info
=====
Port      Packet-Sample-Rate  Max-Header-Size  Counter-interval  Collector-list
                        (in secs)
-----
1/1       0                   128              525               1,2
1/2       0                   128              525               1,2
1/3       0                   128              525               1,2
1/4       0                   128              525               1,2
1/5       0                   128              525               1,2
1/6       0                   128              525               1,2
1/7       0                   128              525               1,2
1/8       0                   128              525               1,2
1/9       0                   128              525               1,2
1/10      0                   128              525               1,2
-----
```

All 10 out of 10 Total Num of sflow port entries displayed

## Variable Definitions

Use the data in the following table to use the **sflow counter-interval** and **show sflow interface** commands.

Variable	Value
<1-3600>	Specifies the polling interval for a slot. Default value is 0 (disabled).
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

## Viewing sFlow Statistics

Display statistics for sFlow datagrams.

### Before you begin

- You must globally enable sFlow.

### Procedure

- Log on to the switch to enter User EXEC mode.
- View sFlow statistics:

```
show sflow statistics [collector <1-2>]
```

### Example

```
Switch:1>show sflow statistics
```

```
=====
sFlow Statistics Info
=====
Collector-id      sFlow-Datagrams
-----
1                  1001
2                   0
-----
```

```
All 2 out of 2 Total Num of sflow statistics entries displayed
```

## Clearing sFlow Statistics

Use this procedure to clear the statistics for each collector.

### Before you begin

- You must globally enable sFlow.

### Procedure

- Enter Privileged EXEC mode:

```
enable
```

- Clear sFlow statistics:

```
clear sflow statistics [collector <1-2>]
```

- Verify the collector information:

```
show sflow statistics [collector <1-2>]
```

### Example

Clear the statistics for collector ID 1.

```
Switch:1>enable
Switch:1#clear sflow statistics collector 1
Switch:1#show sflow statistics collector 1
```

```

=====
                          sFlow Statistics Info
=====
Collector-id                sFlow-Datagrams
-----
1                            0
-----

All 1 out of 1 Total Num of sflow statistics entries displayed

```

## sFlow Configuration Using EDM

Use sFlow to capture traffic statistics to monitor traffic in a data network. This section provides procedures to view and configure sFlow using EDM.

### Enabling sFlow and Configuring the Agent IP Address

Use this procedure to enable sFlow and configure the sFlow Agent IP address so the system can send packets to an sFlow collector.

#### About this task

Application Telemetry and sFlow both use the sFlow Globals tab.

#### Before you begin

You *must* enable sFlow before you enable Application Telemetry.

#### Procedure

1. In the navigation pane, expand the **Configuration > Serviceability** folders.
2. Click **Sflow**.
3. Click the **Globals** tab.
4. Check **AdminEnable** to enable sFlow.
5. In the **AgentAddress** field, enter the agent IPv4 address.
6. Click **Apply**.

### Globals Field Descriptions

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

Name	Description
<b>AdminEnable</b>	Shows whether sFlow is enabled. By default, the check box is not enabled.

*Table continues...*

Name	Description
<b>AgentAddressType</b>	Specifies the collector IP address type. Only IPv4 collector addresses are supported.
<b>AgentAddress</b>	Specifies the agent IP address of an interface that exists in the local management VRF or GRT.

## Configuring an sFlow Collector

Use this procedure to configure the device used as either an sFlow Collector or an Application Telemetry Analytics Engine. This device is where the agent sends sFlow datagrams and Application Telemetry packets for analysis.

sFlow supports up to two collectors for each interface slot in the chassis. However, Application Telemetry supports Collector 1 only.

### \* Note:

- You can configure two Collectors, but Application Telemetry uses Collector 1 only. You *must* configure Collector 1 before you enable Application Telemetry.
- Before you change or remove Collector 1, you must disable Application Telemetry.
- By default, Application Telemetry is globally disabled.

### About this task

#### + Tip:

You can configure the Collector tab to select only the columns you are interested in seeing. By default, the AddressType option does not appear. To make the AddressType column visible, click the down arrow on one of the menu headings, navigate to Columns, and select the AddressType check box.

### Procedure

1. In the navigation pane, expand the **Configuration > Serviceability** folders.
2. Click **Sflow**.
3. Click the **Collector** tab.
4. Configure the fields for Collector 1.
5. Click **Apply**.

## Collector Field Descriptions

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

Name	Description
<b>Index</b>	Shows collector 1 and collector 2. The switch exports sFlow and Application Telemetry traffic to the collector.
<b>Owner</b>	Specifies the sFlow collector name. The string length is 1 to 20 characters.
<b>Timeout</b>	Specifies the time remaining (in seconds) before the collector is released and stops sampling.  The default timeout is 0, which means the timeout is not used and the collector sends data forever.
<b>Address</b>	Specifies the collector IP address. If the default address is set to 0.0.0.0, no traffic is sent.
<b>Port</b>	Specifies the destination port. The default port is 6343.
<b>IsReachable</b>	Shows whether the sFlow collector is reachable.
<b>NextHop</b>	If the collector is reachable, shows the name or address of the next hop through which the collector is reachable.

## Configuring the Packet Samples and Counter Samples

Configure the packet sampling rate to determine how many packets the system counts before it takes a sample and configure the counter sampling interval to determine how often the sFlow agent polls and exports counters for a configured interface. You can also configure the maximum header size on a single port or multiple ports.

### Before you begin

- You must globally enable sFlow.

### About this task

#### Tip:

You can configure the Interfaces tab to select only the columns you are interested in seeing. By default, the Instances option does not appear. To make the **Instances** column visible, click the down arrow on one of the menu headings, navigate to Columns, and select the Instances check box.

### Procedure


- In the navigation pane, expand the **Configuration > Serviceability** folders.
- Click **Sflow**.
- Click the **Interfaces** tab.



4. In the **DataSource** column, navigate to the slot and port where you want to configure sFlow, and configure the following:
  - a. **PacketSamplingRate**—Double-click the field, and enter a sampling rate value. The range is 8192 to 1000000.
  - b. **MaximumHeaderSize**—Double-click the field, and enter a maximum header size value. The range is 64 to 256.
  - c. **Interval**—Double-click the field, and enter the counter sampling interval value in seconds. The range is 1 to 3600.
5. Click **Apply**.

## Interfaces Field Descriptions

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

Name	Description
<b>DataSource</b>	Shows the slot and port for which traffic statistics are collected.  The field name is different on different hardware platforms.
<b>Instance</b>	Shows the number of sFlow samplers associated with a specific datasource. The value for each sFlow sampler instance can range from 1–65535.   <b>Note:</b> You must select this field for it to display on the Interfaces tab.
<b>Collectors</b>	Shows the collectors that have been configured for the sFlow agent to send sFlow datagrams. Two collectors are supported.  The field name is different on different hardware platforms.
<b>PacketSamplingRate</b>	Specifies the packet sampling rate to determine how many packets the system counts before it take a sample.  Range 8192–1000000. Default 0.
<b>MaximumHeaderSize</b>	Specifies the maximum header size on a single port or multiple ports.  Range 64–256 bytes. Default 128 bytes
<b>Interval</b>	Specifies the counter sampling interval to determine how often the sFlow agent polls and exports counters for a configured interface.  Range 1–3600 seconds. Default 0.

## Displaying sFlow Statistics

Use the following procedure to display (true) sFlow statistics. Statistics for sFlow are cleared (false), by default.

### About this task

#### + Tip:

You can configure the Stats tab to select only the columns you are interested in seeing. All the options appear, by default. To hide a column, click the down arrow on one of the menu headings, navigate to Columns, and select the check box for the column you want to hide.

### Procedure

1. In the navigation pane, expand the **Configuration > Serviceability** folders.
2. Click **sFlow**.
3. Click the **Stats** tab.
4. In the **ClearStats** column, double-click the field, and select true or false from the list.
5. Click **Apply**.

## Stats Field Descriptions

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

Name	Description
<b>Index</b>	Shows sFlow collector ID 1 and 2
<b>DatagramCount</b>	Shows the number of datagrams that have been sent to the collector.
<b>ClearStats</b>	Shows whether the sFlow statistics are displayed (true) or cleared (false). The default is false.

## Configuring counter polling

### About this task

Use this procedure to configure counter polling using the EDM.

This tab is not supported on all hardware platforms.

### Procedure

1. In the navigation pane, expand the **Configuration > Serviceability** folders.
2. Click **sFlow**.
3. Click the **Counter Polling** tab.
4. In the **Index** column, navigate to the slot and port where you want to configure sFlow.

5. Double-click the **Interval** field and enter a value in seconds.
6. Click **Apply**.

## Counter Polling Field Descriptions

Use the data in the following table to use the Counter Polling tab.

Name	Description
<b>Index</b>	Identifies the source of the data for counter polling.
<b>Receiver</b>	The sFlow receiver associated with the counter polling.
<b>Interval</b>	The maximum number of seconds between successive samples of the counters associated with this data source. A sampling interval of 0 disables counter sampling.

# Chapter 11: Application Telemetry Fundamentals

Extreme Networks offers two Analytics solutions that monitor traffic on your network:

- sFlow
- Application Telemetry

## Important:

You can use either sFlow, or sFlow with Application Telemetry or both at the same time as they can coexist on a switch. Note that, in order to enable Application Telemetry, you *must* enable sFlow first.

In both solutions, the switch collects flow information and sends it to a central server that processes the information and provides statistical data in the form of reports. Then you can use Extreme Management Center to analyze the reports to give you a full understanding of the applications on your network and learn who is using those applications. Extreme Management Center also provides information such as DoS tracking, security monitoring, and statistics for protocols, ports, and applications.

This chapter describes how Application Telemetry works and how to configure it. Because there is some commonality between the two features, this chapter also describes some sFlow features.

For further information about sFlow, see [sFlow Fundamentals](#) on page 163.

For more information about Extreme Management Center, see the documentation on the Extreme Networks Documentation portal ([www.extremenetworks.com/documentation/](http://www.extremenetworks.com/documentation/)) with special attention to the *Application Analytics User Guide*.

---

## How Application Telemetry Works

Both sFlow and Application Telemetry mirror packets to a server for deep packet inspection, but they collect streams in very different ways:

- sFlow samples 1 out of n packets to create flow streams. This methodology achieves scalability and applies to high speed networks, but it provides limited application visibility.
- Application Telemetry does not sample some packets like sFlow; it monitors all traffic and uses policy rules to filter packets for analysis. This pattern matching methodology enables

Application Telemetry to monitor *all* application-level traffic flows at wire speed on *all* interfaces simultaneously.

The policy rules that Application Telemetry uses are ACL and ACE filters that are pre-configured in a policy configuration file called `sflow.pol`. This policy file is not user configurable. These rules enable the switch to recognize several signatures that represent a combination of the following:

- IP protocol type (TCP/UDP)
- TCP flags
- Layer 4 port numbers
- data patterns (defined as offset/data/mask triplets)

*Pattern matching* enables Application Telemetry to target very specific, well-defined packets in each flow and not full streams of traffic. Thus, the switch mirrors only a relatively few packets to the Analytics Engine. It is the Analytics Engine that performs deep packet inspection to create reports of statistical data.

### **Important:**

When you enable Application Telemetry, the switch loads the filter rules based on the logic below:

- Application Telemetry uses the `apptelemetry.pol` or the `sflow.pol` file because the filter rules can exist in either file. The `sflow.pol` file is the default file and is included with the image that is loaded on the switch. This file contains the default filter rules. The `apptelemetry.pol` file is the user-defined file, which can be updated by the Extreme Management Center. To use this file, configure Application Telemetry using the Extreme Management Center. When you enable Application Telemetry from Extreme Management Center, the updated `apptelemetry.pol` file is placed in `/intflash/`.
- When you enable Application Telemetry, the feature uses the files in the following order:
  - If the user-defined file (`apptelemetry.pol`) exists, then the switch loads the rules from this file.
  - If the `apptelemetry.pol` file does not exist or if there is a problem reading this file, then the switch uses the default `sflow.pol` file.

---

## Common Elements Between sFlow and Application Telemetry

sFlow and Application Telemetry send mirrored packets from a common *source* to a common *destination*. sFlow sends samples directly to the destination, while Application Telemetry sends mirrored packets through a GRE tunnel, to the same destination.

The tunnel source is the switch that you want to monitor:

- sFlow sends sampled flows.
- Application Telemetry sends packets that match its policy rules.

Both sFlow and Application Telemetry use an agent to package either the sFlow streams or the Application Telemetry packets. To configure the agent, they both use the `sflow agent-ip` command.

**\* Note:**

The switch sends only one mirrored copy. For information on which mirrored copies take precedence, see [Configuration considerations](#) on page 182.

The tunnel destination for the mirrored traffic is a server where software performs a deep packet inspection of the mirrored traffic.

- sFlow sends flow and counter samples as datagrams to the sFlow Collector.
- Application Telemetry sends packets that match the policy rules over a GRE tunnel to the Analytics Engine.

To configure the tunnel destination, they both use the `sflow collector <1-2>` command.

**! Important:**

You can configure two Collectors, but Application Telemetry uses Collector 1 only. You *must* configure Collector 1 before you enable Application Telemetry.

---

## Operational Considerations and Restrictions

The following section describes operational considerations for deploying Application Telemetry, including general considerations that apply to all platforms, followed by a summary of platform-specific considerations.

### General Considerations

The following section describes general Application Telemetry operational considerations for all platforms.

- When you enable Application Telemetry, it is globally enabled on all slots. You cannot disable the feature on a per slot basis.
- Application Telemetry supports IPv4 and IPv6 packets.
- Application Telemetry filter rules are not user configurable. However, an updated `app-telemetry.pol` file can be installed through the Extreme Management Center.
- If a user-created filter rule (ACL) conflicts with an Application Telemetry defined filter, the user-created rule always takes precedence.
- There are two configurable sFlow collectors (Collector 1 and Collector 2). However, Application Telemetry uses Collector 1 only and you must configure it before enabling Application Telemetry.

- In a Fabric Extend deployment, Application Telemetry does not mirror ingressing NNI to UNI IpShortcut traffic. Application Telemetry sends (and counts) two packets for each UNI to NNI packet that is sent with IpShortcuts over Fabric Extend.

### Platform-specific Considerations

The following table provides a summary of operational considerations for different VOSS switches.

**Table 13: Summary of Application Telemetry Operational Considerations**

Category	Attribute	VSP 4000, 7200, 7400, 8200, 8400	VSP 8600
Supported flow types	Flows that ingress standard VLAN ports	Supported	Supported
	Flows that ingress UNI ports	Supported	Supported
	Flows that ingress NNI ports and egress UNI ports (Layer 2 VSN)	Supported	Not supported
	Flows that ingress NNI ports and egress UNI ports (Layer 3 VSN)	Supported Exception: Not supported on VSP 7400	Not supported
	Flows that ingress NNI ports and terminate locally	Supported	Not supported
	Flow that ingress NNI ports and egress NNI ports	Not supported	Not supported
	Flows on DvR Controllers	Supported	Not applicable
Application Telemetry collector/server reachability	GRT	Yes	Yes
	VRF	No	Yes Exception: management VRF
	Fabric Connect – Layer 2 VSNs	Yes When the Analytics Engine is reachable over a Layer 2 VSN, the GRE packets are encapsulated with MAC-in-MAC (IEEE 802.1ah) at the originating BEB. The MAC-in-MAC header is removed at the terminating BEB and the	Yes When the Analytics Engine is reachable over a Layer 2 VSN, the GRE packets are encapsulated with MAC-in-MAC (IEEE 802.1ah) at the originating BEB. The MAC-in-MAC header is removed at the terminating BEB and the

*Table continues...*

Category	Attribute	VSP 4000, 7200, 7400, 8200, 8400	VSP 8600
		original GRE packet is sent to the collector. Note also that the MAC-in-MAC encapsulation plus the GRE encapsulation adds 60 bytes to the original packet. Therefore, if the original packet is close to the maximum transmission unit (MTU), the mirrored copy may exceed the MTU and be dropped.	original GRE packet is sent to the collector. Note also that the MAC-in-MAC encapsulation plus the GRE encapsulation adds 60 bytes to the original packet. Therefore, if the original packet is close to the maximum transmission unit (MTU), the mirrored copy may exceed the MTU and be dropped.
	Fabric Connect – IP Shortcut Routing	Yes	Yes
	Fabric Connect – Layer 3 VSNs	No	No
Coexistence with sFlow	If you enable sFlow and Application Telemetry simultaneously on the same port	The switch sends the sFlow datagrams and Application Telemetry packets to the collector.	If the packet matches the Application Telemetry rules, the switch mirrors the packet to the GRE tunnel and sends it to the Analytics Engine and it cannot be sampled by sFlow.  If the packet does not match the Application Telemetry rules and the packet gets sampled, the switch sends it as an sFlow datagram to the sFlow Collector.
Coexistence with security filters	IPv6 security filters or IPv6 source guard	Not supported (consistency checks in place)  Exception: Allowed on VSP 7400	Allowed
Coexistence with mirroring	Mirroring resources	Only 3 mirror ports can be configured for general port mirroring	No impact to number of mirror ports
	If rx port mirroring is enabled on a port, and Application Telemetry is enabled, when a packet that	The switch generates the remote mirrored packet, and the port-based mirroring copy.	The switch generates the remote mirrored packet only. The switch does not generate the port-based mirroring copy.

*Table continues...*



Category	Attribute	VSP 4000, 7200, 7400, 8200, 8400	VSP 8600
	matches one Application Telemetry entry criterion comes to this port		If a packet does not match an Application Telemetry rule, the switch generates the port-based mirroring copy.
Coexistence with Unicast Reverse Path Forwarding (uRPF)	If you enable uRPF mode on the switch	The MTU values for both IPv4 and IPv6 packets on the same VLAN are always matched. Different Layer 3 MTU sizes on the same VLAN are not allowed in uRPF mode.	The URPF boot config flag is not applicable. Even when uRPF is enabled, IPv6 MTU can be different from IPv4 MTU; both need not be the same.
High Availability	Application Telemetry deployed in a High Availability environment	Not applicable	Supported
Counters	If packets match both user defined filters (ACLs) and Application Telemetry rules, and if both rules have counters	Both counters incremented	ACL counters incremented only
Match off-set	smb, kerberosasreq2 and kerberostgsreq packet types	kerberosasreq2 and kerberostgsreq packet types supported. Smb – not available	kerberosasreq2 and kerberostgsreq packet types supported with an off-set of 24 bytes only; an off-set of 40 bytes is not supported

## Configuration Overview

After the optional step of uploading the `apptelemetry.pol` file to flash memory using Extreme Management Center, activate Application Telemetry by configuring the following:

1. Configure the IP address of the egress interface for the GRE tunnel with the `sFlow agent-ip` command.
2. Enable sFlow with the `sflow enable` command.
3. Configure the IP address of the Analytics Engine with the `sFlow collector 1` command.
4. Enable Application Telemetry with the `app-telemetry enable` command.

The following figure shows the Application Telemetry agent on various routers and switches with packets being sent to the Analytics Engine.




Figure 4: Application Telemetry Overview

Table 14: Application Telemetry Legend

Number	Description
1	Analytics Engine
2	GRE tunnels
3	Application Telemetry agents

## Application Telemetry Configuration Using CLI

Use Application Telemetry to capture traffic statistics to monitor traffic in a data network. This section provides procedures to view and configure this feature using CLI.

### Configuring the Agent IP Address

Use this procedure to configure the source of the Application Telemetry packets.

## Procedure

1. Enter Global Configuration mode:  

```
enable
```

```
configure terminal
```
2. Enable the agent IPv4 address:  

```
sflow agent-ip {A.B.C.D}
```

## Example

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch:1(config)#sflow agent-ip 192.0.2.27
```

## Variable Definitions

Use the data in the following table to use the `sflow agent-ip` command.

Variable	Definition
{A.B.C.D.}	Specifies the agent-ip address (IPv4).

## Configuring an Analytics Engine and Enabling Application Telemetry Globally

Use this procedure to enable Application Telemetry and configure the device used as either an sFlow Collector or an Application Telemetry Analytics Engine. This device is where the agent sends sFlow datagrams and Application Telemetry packets for analysis.

sFlow supports up to two collectors for each interface slot in the chassis. However, Application Telemetry supports Collector 1 only.

### \* Note:

- You can configure two Collectors, but Application Telemetry uses Collector 1 only. You *must* configure Collector 1 before you enable Application Telemetry.
- Before you change or remove Collector 1, you must disable Application Telemetry.
- By default, Application Telemetry is globally disabled.

### Before you begin

- You must configure the agent IP address.
- You must enable sFlow before you can enable Application Telemetry.

## Procedure

1. Enter Global Configuration mode:  

```
enable
```

```
configure terminal
```

2. Configure the Analytics Engine information using Collector 1:

```
sflow collector 1 address {A.B.C.D} [owner WORD<1-20>] [vrf WORD<1-16>]
```

3. Verify the Analytics Engine configuration:

```
show sflow collector 1
```

4. Enable Application Telemetry:

```
app-telemetry enable
```

5. Verify the global configuration:

```
show app-telemetry status
```

**\* Note:**

The output of this command shows whether Application Telemetry is enabled or not and if the collector is reachable.

**Example**

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#sflow collector 1 address 192.0.2.26 owner sflow1 port 6343 timeout 497
Switch:1(config)#show sflow collector 1
```

```
=====
                        sFlow Collector Configuration Info
=====
```

Id	Owner	Collector-IP	Port	Timeout(secs)	Reachable via
1	sflow1	192.0.2.26	6343	497	192.0.2.15

```
=====
```

```
All 1 out of 1 Total Num of sflow collector entries displayed
Switch:1(config)#app-telemetry enable
Switch:1(config)#show app-telemetry status
Application Telemetry is enabled
Collector is reachable via 192.0.2.26
```

**Variable Definitions**

Use the data in the following table to use the **sflow collector** command.

Variable	Value
<1-2>	Specifies the ID of the collector where you want to send packets for analysis. Application Telemetry uses Collector 1 only.
owner WORD<1-20>	Specifies the name of the collector.
Collector-IP {A.B.C.D.}	Specifies the IP address of the collector.

*Table continues...*

Variable	Value
port <1–65535>	Specifies the destination port. The default port is 6343.  * <b>Note:</b> Application Telemetry does not use this parameter.
timeout <1–65535>	Specifies the time remaining (in seconds) before the collector is released.  The default timeout is 0, which means the timeout is not used and the collector sends data forever.  * <b>Note:</b> Application Telemetry does not use this parameter.
vrf WORD<1–16>	Specifies the name of the VRF used to reach the collector.  * <b>Note:</b> This parameter is not supported on all hardware platforms.

## Viewing Application Telemetry Counters

Use the following procedure to view the Application Telemetry status counters. The switch assigns an ID to each counter and displays information about each filter rule by name. The information includes how many packets were transmitted to the Analytics Engine that matched the specified pattern in the rule and the total number of bytes in the packets.

### Procedure

1. Log on to the switch to enter User EXEC mode.
2. View Application Telemetry counters:

```
show app-telemetry counter [name <WORD<1-32> | id <1-2000>]
```

### Example

```
Switch:1>show app-telemetry counter
```

```
=====
Application Telemetry Counters
=====
EntryId      Name           Packets      Bytes
-----
1            ssh            1258         72145
2            sslclient     457          27000
=====
```

```
All 2 out of 2 Total Num of Application Telemetry counters entries displayed
```

## Clearing Application Telemetry Counters

Use this procedure to clear the Application Telemetry status counters. You can clear all of the counters or specify just the counters you want to clear by name or ID.

### Procedure

1. Log on to the switch to enter User EXEC mode.
2. Clear Application Telemetry counters:

```
clear app-telemetry counter [name <WORD<1-32> | id <1-2000>]
```

3. Verify that the counters were cleared:

```
show app-telemetry counter [name <WORD<1-32> | id <1-2000>]
```

### Example

Clear the counters.

```
Switch:1>enable
Switch:1#clear app-telemetry counter
Switch:1>show app-telemetry counter
```

```
=====
Application Telemetry Counters
=====
EntryId      Name           Packets      Bytes
-----
1            ssh            0            0
2            sslclient     0            0
-----
```

```
All 2 out of 2 Total Num of Application Telemetry counters entries displayed
```

## Application Telemetry Configuration Using EDM

Use Application Telemetry to capture traffic statistics to monitor traffic in a data network. This section provides procedures to view and configure this feature using EDM.

## Enabling sFlow and Configuring the Agent IP Address

Use this procedure to enable sFlow and configure the sFlow Agent IP address so the system can send packets to an sFlow collector.

### About this task

Application Telemetry and sFlow both use the sFlow Globals tab.

## Before you begin

You *must* enable sFlow before you enable Application Telemetry.

### Procedure

1. In the navigation pane, expand the **Configuration > Serviceability** folders.
2. Click **Sflow**.
3. Click the **Globals** tab.
4. Check **AdminEnable** to enable sFlow.
5. In the **AgentAddress** field, enter the agent IPv4 address.
6. Click **Apply**.

## Globals Field Descriptions

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

Name	Description
<b>AdminEnable</b>	Shows whether sFlow is enabled. By default, the check box is not enabled.
<b>AgentAddressType</b>	Specifies the collector IP address type. Only IPv4 collector addresses are supported.
<b>AgentAddress</b>	Specifies the agent IP address of an interface that exists in the local management VRF or GRT.

## Configuring an sFlow Collector

Use this procedure to configure the device used as either an sFlow Collector or an Application Telemetry Analytics Engine. This device is where the agent sends sFlow datagrams and Application Telemetry packets for analysis.

sFlow supports up to two collectors for each interface slot in the chassis. However, Application Telemetry supports Collector 1 only.

### \* Note:

- You can configure two Collectors, but Application Telemetry uses Collector 1 only. You *must* configure Collector 1 before you enable Application Telemetry.
- Before you change or remove Collector 1, you must disable Application Telemetry.
- By default, Application Telemetry is globally disabled.

### About this task

### + Tip:

You can configure the Collector tab to select only the columns you are interested in seeing. By default, the AddressType option does not appear. To make the AddressType column visible,

click the down arrow on one of the menu headings, navigate to Columns, and select the AddressType check box.

### Procedure

1. In the navigation pane, expand the **Configuration > Serviceability** folders.
2. Click **Sflow**.
3. Click the **Collector** tab.
4. Configure the fields for Collector 1.
5. Click **Apply**.

## Collector Field Descriptions

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

Name	Description
<b>Index</b>	Shows collector 1 and collector 2. The switch exports sFlow and Application Telemetry traffic to the collector.
<b>Owner</b>	Specifies the sFlow collector name. The string length is 1 to 20 characters.
<b>Timeout</b>	Specifies the time remaining (in seconds) before the collector is released and stops sampling.  The default timeout is 0, which means the timeout is not used and the collector sends data forever.
<b>Address</b>	Specifies the collector IP address. If the default address is set to 0.0.0.0, no traffic is sent.
<b>Port</b>	Specifies the destination port. The default port is 6343.
<b>IsReachable</b>	Shows whether the sFlow collector is reachable.
<b>NextHop</b>	If the collector is reachable, shows the name or address of the next hop through which the collector is reachable.

---

## Enabling Application Telemetry Globally

Use this procedure to globally enable Application Telemetry so it can send packets to an Analytics Engine. By default, Application Telemetry is globally disabled.

### Before you begin

You must complete the following:

- Configure an agent IP address.
- Enable sFlow.



- Configure Collector 1.

### Procedure

1. In the navigation pane, expand the **Configuration > Serviceability** folders.
2. Click **Application Telemetry**.
3. Click the **Globals** tab.
4. Select the **AdminEnable** check box.
5. Click **Apply**.

## Globals Field Descriptions

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

Name	Description
<b>AdminEnable</b>	Shows whether Application Telemetry is enabled. By default, the check box is not enabled.
<b>ClearCounterStats</b>	Clears the Application Telemetry status counters.

## Viewing Application Telemetry Counters

Use the following procedure to view the Application Telemetry status counters. The switch assigns an ID to each counter and displays information about each filter rule by name. The information includes how many packets were transmitted to the Analytics Engine that matched the specified pattern in the rule and the total number of bytes in the packets.

### Procedure

1. In the navigation pane, expand the **Configuration > Serviceability** folders.
2. Click **Application Telemetry**.
3. Click the **Counter** tab.

## Clearing Application Telemetry Counters

Use this procedure to clear the Application Telemetry status counters. You can clear all of the counters or specify just the counters you want to clear by name or ID.

### Procedure

1. In the navigation pane, expand the **Configuration > Serviceability** folders.
2. Click **Application Telemetry**.
3. Click the **Counter** tab.
  - To clear all the counters, click the **Globals** tab and select **ClearCounterStats**.

- To clear specific counters, click the **Counter** tab, select the counters you want to clear, and then click **ClearStats**.
4. Click **Apply**.

---

## Viewing Application Telemetry Status

### About this task

Use this procedure to view the status of the Application Telemetry collector.

### Procedure

1. In the navigation pane, expand the **Configuration > Serviceability** folders.
2. Click **Application Telemetry**.
3. Click the **Status** tab.

### Status field descriptions

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

Name	Description
<b>Collector IP Address</b>	Shows the address of the Application Telemetry collector.
<b>IsReachable</b>	Shows whether the Application Telemetry collector is reachable.
<b>NextHop</b>	If the collector is reachable, shows the name or address of the next hop through which the collector is reachable.

# Chapter 12: Statistics

This chapter provides the procedures for using statistics to help monitor the performance of the switch using Enterprise Device Manager (EDM) and command line interface (CLI).

---

## Viewing Statistics Using CLI

This section contains procedures to view statistics in the CLI.

---

### Viewing TCP Statistics

View TCP statistics to manage network performance.

#### Procedure

View TCP statistics:

```
show ip tcp statistics
```

#### Example

```
Switch: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.
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.

**\* Note:**

This command is not available on all hardware platforms.

### Procedure

View port routing statistics:

```
show routing statistics interface [gigabitethernet] [{slot/port[-slot/
port] [, ...]]
```

**Example**

```
Switch:1#show routing statistics interface gigabitethernet 1/7-1/9
```

Port Stats Routing					
PORT NUM	IN_FRAME UNICAST	IN_FRAME MULTICAST	IN DISCARD	OUT_FRAME UNICAST	OUT_FRAME MULTICAST
1/7	1386	0	0	1344	0
1/8	1302	0	0	1344	0
1/9	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[/sub-port][/-slot/port[/sub-port]][,....]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

**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****About this task**

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

**\* Note:**

This command is not available on all hardware platforms.

**Procedure**

1. Enter Privileged EXEC mode:

```
enable
```

2. View bridging statistics for a specific port:

```
show interfaces GigabitEthernet statistics bridging [{slot/port[-slot/port] [,...]}]
```

**Example**

```
Switch: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  _FRAME    xSTP BPDU  xSTP BPDU
=====
1/1    179325     0          0          119310     179325     0          0
1/2    187951     26078     42         689486     179324     0          25617
1/3    0          0          0          0          0          0          0
1/4    0          0          0          0          0          0          0
1/5    0          0          0          0          0          0          0
1/6    394       0          0          948942     360        0          0
1/7    4689      0          0          863403     360        0          0
1/8    4369      3206     116        958752     360        0          3995
1/9    0          0          0          0          0          0          0
1/10   0          0          0          0          0          0          0
1/11   0          0          0          0          0          0          0
1/12   0          0          0          0          0          0          0
1/13   179325     0          0          42040      179325     0          0
1/14   187864     0          0          50437      179324     0          0
1/15   0          0          0          0          0          0          0
1/16   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[/sub-port] [-slot/port[/sub-port]][, ...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

## 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

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

**\* Note:**

Slot and port information can differ depending on hardware platform.

### Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View DHCP-relay statistics for a specific port or VRF.

```
show interfaces GigabitEthernet statistics dhcp-relay [vrf
WORD<1-16>] [vrfs WORD<0-255>] [{slot/port[/sub-port] [-slot/port[/
sub-port]] [, ...]]
```

### Example

View DHCP-relay statistics:

```
Switch:1>enable
Switch:1#show interfaces gigabitEthernet statistics dhcp-relay
```

```
=====
                        Port Stats Dhcp
=====
PORT_NUM VRF NAME          NUMREQUEST NUMREPLY
-----
1/12     GlobalRouter             0           2
1/13     GlobalRouter             3           2
2/3      GlobalRouter             0           2
=====
```

## 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>	Specifies a VRF instance by VRF name.
vrfids <i>WORD</i> <0-255>	Specifies the ID of the VRF.
{slot/port[/sub-port][/-slot/port[/sub-port]][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (1/1).  Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

## 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.

 **Note:**

Slot and port information can differ depending on hardware platform.

### Procedure

1. Show the number of requests and replies 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

```
Switch: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
-----
Port 1/12      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 WORD<0-16>	Specifies a VRF instance by the VRF name.
vrfids WORD<0-512>	Specifies the ID of the VRF.

## Job Aid

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

**Table 19: show ip dhcp-relay counters option82 command**

Heading	Description
INTERFACE	Shows the VLAN or port associated with the respective relay interface.
IP ADDR	Shows the IP address of the respective relay interface.
FOUND OPT82	Shows the number of packets received that included option82. This number increases every time a valid DHCP packet that contains option82 arrives on the respective relay interface.
DROP PKT	Shows the number of packets the interface did not forward.  This number increases every time a DHCP packet that has option82 arrives on a relay interface but is not forwarded on the interface towards the server; the path towards the relay can include additional DHCP relays.  To determine the cause of the drop, you must enable trace on level 170.

*Table continues...*

Heading	Description
CIRC ID	Show the circuit ID associated with the respective interface.
ADD CIRC	Shows on how many packets the circuit ID was inserted for that interface.  This number increases every time the relay adds a circuit id sub-option in a generated option82 packet to send on an interface towards the server.  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.
DEL CIRC	Shows on how many packets the circuit id was removed for that interface.  This number increases every time the relay removes a circuit id sub-option from an option82 packet received on a interface towards the server.
REMOTE ID	Shows the remote ID associated with the respective interface. The value is the MAC address of the interface.
ADD REMID	Shows on how many packets the remote ID was inserted for that interface.  This number increases every time the relay adds a remote id sub-option in a generated option82 packet to send through an interface towards a server.  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.
DEL REMID	Shows on how many packets the remote ID was removed for that interface.  This number increases every time the relay removes a remote id sub-option from an option82 packet received on an interface towards a server.

---

## Displaying LACP Statistics for Specific Ports

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

 **Note:**

Slot and port information can differ depending on hardware platform.

## Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View statistics for specific ports:

```
show interfaces GigabitEthernet statistics lacp [{slot/port[/sub-
port]}[-slot/port[/sub-port]][,...]]
```

## Example

View LACP statistics:

```
Switch:1>enable
Switch:1#show interfaces gigabitethernet statistics lacp

=====
Port Stats Lacp
=====
PORT  TX      RX      TX      RX      TX      RX      RX      RX
NUM   LACPDU LACPDU MARKERPDU MARKERPDU MARKERRESPPDU MARKERRESPPDU UNKNOWN ILLEGAL
-----
1/39  0        0        0        0        0        0        0        0
1/40  0        0        0        0        0        0        0        0
2/37  0        0        0        0        0        0        0        0
2/38  0        0        0        0        0        0        0        0
```

## Variable Definitions

Use the data in the following table to use the **show interfaces GigabitEthernet statistics lacp** command.

Variable	Value
{slot/port[/sub-port]}[-slot/port[/sub-port]][,...]	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

## Job Aid

The following table describes parameters for the **show interfaces GigabitEthernet statistics lacp** command.

**Table 20: 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.

*Table continues...*

Parameter	Description
TX MARKERPDU	The count of transmitted marker protocol data units.
RX MARKERPDU	The count of received marker protocol data units.
TX MARKERRESPPDU	The count of transmitted marker protocol response data units.
RX MARKERRESPPDU	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 VLACP Statistics for Specific Ports

Display VLACP statistics for specific ports to manage network performance.

**\* Note:**

Slot and port information can differ depending on hardware platform.

### About this task

You can enable sequence numbers for each VLACPDU to assist in monitoring performance. The switch counts mismatched PDU sequence numbers to determine packet loss information. By default, sequence numbers are enabled.

You can use the show commands from Privileged EXEC mode but must enter Global Configuration mode to enable or disable the sequence numbers.

### Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Confirm sequence numbers are enabled:

```
show vlacp
```

3. **(Optional)** Enable sequence numbers for VLACPDUs:

```
vlacp sequence-num
```

4. View VLACP statistics:

```
show interfaces gigabitEthernet statistics vlacp [{slot/port[/sub-port]}[-slot/port[/sub-port]][,...]] ]
```

5. **(Optional)** View VLACP statistics history:

```
show interfaces gigabitEthernet statistics vlacp history [{slot/port[/sub-port]}[-slot/port[/sub-port]][,...]] ]
```

**6. (Optional) Clear VLACP statistics:**

```
clear vlacp stats [port {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...]]}
```

**7. (Optional) Disable sequence numbers for VLACPDUs:**

```
no vlacp sequence-num
```

**Example**

Determine if sequence numbers are enabled, and then view port statistics. Port numbering may differ depending on your product and configuration.

```
Switch:1(config)#show vlacp
=====
                        Vlacp Global Information
=====
                SystemId: 32:11:9f:20:00:00
                  Vlacp: enable
    Vlacp Sequence Number: enable

Switch:1(config)#show interfaces gigabitEthernet statistics vlacp
=====
                        Port Stats Vlacp
=====
PORT      TX      RX      SEQNUM
NUM      VLACPDU VLACPDU MISMATCH
-----
8/1      106058  105554    0
12/11    15      12        0
12/23    0        0         0
```

**Variable Definitions**

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

Variable	Value
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

**Job Aid**

The following table describes fields in the output for the **show interfaces gigabitEthernet statistics vlacp** command.

Field	Description
PORT NUM	Shows the slot and port number.

*Table continues...*

Field	Description
TX VLACPDU	Shows the number of VLACPDU transmitted on the port.
RX VLACPDU	Shows the number of valid VLACPDU received on the port.
SEQNUM MISMATCH	Shows the number of mismatched VLACPDU in terms of received sequence numbers on the port.

## Displaying RMON Statistics for Specific Ports

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

**\* Note:**

Slot and port information can differ depending on hardware platform.

### Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View statistics for specific ports:

```
show interfaces GigabitEthernet statistics rmon {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...]}
```

### Example

View RMON statistics:

```
Switch:1>enable
Switch:1#show interfaces gigabitEthernet statistics rmon 1/13
```

```

=====
                               Port Stats Rmon
=====
PORT  OCTETS   PKTS   MULTI  BROAD  CRC   UNDER  OVER  FRAG  COLLI
NUM                                     SIZE  SIZE  MENT  SION
-----
1/13  1943     21     8      13     0     0       0     0     0

```

## Variable Definitions

Use the data in the following table to use the **show interfaces GigabitEthernet statistics rmon** command.

Variable	Value
{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is

Variable	Value
	channelized, you must also specify the sub-port in the format slot/port/sub-port.

## Job Aid

The following table describes parameters for the `show interfaces GigabitEthernet statistics rmon` command output.

**Table 21: 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 (FCS) with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).

*Table continues...*

Parameter	Description
COLLISION	An estimated value for the total number of collisions on this Ethernet segment.

## Displaying Detailed Statistics for Ports

Display detailed statistics for specific ports to manage network performance.

**\* Note:**

Slot and port information can differ depending on hardware platform.

### Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View statistics for specific ports:

```
show interfaces GigabitEthernet statistics verbose {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...]}
```

### Example

View statistics for various ports:

```
Switch:1>enable
Switch:1#show interfaces gigabitethernet statistics verbose

Please widen the terminal for optimal viewing of data.

=====
Port Stats Interface Extended
=====
PORT_NUM IN_UNICST  OUT_UNICST  IN_MULTICST  OUT_MULTICST  IN_BRDCST  OUT_BRDCST  IN_LSM  OUT_LSM
-----
2/1      0            0            0            0            0            0            0        0
2/2      0            0            0            0            0            0            0        0
2/3      0            0            0            0            0            0            0        0
2/4      0            0            0            0            0            0            0        0
2/5      0            0            0            0            0            0            0        0
2/6      0            0            0            0            0            0            0        0
3/1      0            0            0            0            0            0            0        0
3/2      0            0            0            0            0            0            0        0
3/3      0            0            8702         34805         0            0            0        0
3/4      0            0            0            0            0            0            0        0
3/5      0            0            0            0            0            0            0        0
3/6      0            0            0            0            0            0            0        0
3/7      0            0            0            0            0            0            0        0
3/8      0            0            0            0            0            0            0        0
3/9      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[/sub-port][/-slot/port[/sub-port]][,....]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

## Job Aid

The following table describes parameters for the `show interfaces GigabitEthernet statistics verbose` command.

**Table 22: how interfaces GigabitEthernet statistics verbose field descriptions**

Parameter	Description
PORT_NUM	Indicates the port number.
IN_UNICAST	The count of inbound Unicast packets.
OUT_UNICAST	The count of outbound Unicast packets.
IN_MULTICAST	The count of inbound Multicast packets.
OUT_MULTICAST	The count of outbound Multicast packets.
IN_BRDCST	The count of inbound broadcast packets.
OUT_BRDCST	The count of outbound broadcast packets.

---

## Displaying IS-IS Statistics and Counters

Use the following procedure to display the IS-IS statistics and counters.

### Procedure

1. Display IS-IS system statistics:  
`show isis statistics`
2. Display IS-IS interface counters:  
`show isis int-counters`
3. Display IS-IS level 1 control packet counters:  
`show isis int-l1-ctrl-pkts`

**\* Note:**

The switch uses level 1 IS-IS. The switch does not support level 2 IS-IS. The command `show isis int-l2-ctrl-pkts` is not supported because the IEEE 802.1aq standard currently only defines the use of one hierarchy, Level 1.

4. Clear IS-IS statistics:

```
clear isis stats [error-counters] [packet-counters]
```

**Example**

```
Switch:1# show isis statistics
=====
                        ISIS System Stats
=====
LEVEL      CORR   AUTH   AREA   MAX SEQ   SEQ NUM   OWN LSP   BAD ID   PART   LSP   DB
          LSPs  FAILS  DROP   EXCEEDED SKIPS     PURGE   LEN     CHANGES OLOAD
-----
Level-1    0     0     0     0         1         0       0       0       0     0

Switch:1# show isis int-counters
=====
                        ISIS Interface Counters
=====
IFIDX      LEVEL  AUTH   ADJ      INIT      REJ      ID LEN  MAX AREA LAN  DIS
          L1-2  FAILS  CHANGES  FAILS     ADJ      ADJ     CHANGES CHANGES
-----
Mlt2       Level 1-2  0     1         0         0         0       0     0     0
Port1/21   Level 1-2  0     1         0         0         0       0     0     0

Switch:1# show isis int-l1-ctrl-pkts
=====
                        ISIS L1 Control Packet counters
=====
IFIDX      DIRECTION  HELLO      LSP      CSNP      PSNP
-----
Mlt2       Transmitted 13346      231      2         229
Mlt2       Received    13329      230      1         230
Port1/21   Transmitted 13340      227      2         226
Port1/21   Received    13335      226      1         227
```

**Variable Definitions**

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

Variable	Value
error-counters	Clears IS-IS stats error-counters.
packet-counters	Clears IS-IS stats packet-counters.

**Job Aid**

**show isis statistics**

The following table describes the fields in the output for the `show isis statistics` command.

Parameter	Description
LEVEL	Shows the level of the IS-IS interface.
CORR LSPs	Shows the number of corrupted LSPs detected.

*Table continues...*

Parameter	Description
AUTH FAILS	Shows the number of times authentication has failed on the global level.
AREA DROP	Shows the number of manual addresses dropped from the area.
MAX SEQ EXCEEDED	Shows the number of attempts to exceed the maximum sequence number.
SEQ NUM SKIPS	Shows the number of times the sequence number was skipped.
OWN LSP PURGE	Shows how many times the local LSP was purged.
BAD ID LEN	Shows the number of ID field length mismatches.
PART CHANGES	Shows the number of partition link changes.
LSP DB OLOAD	Show the number of times the switch was in the overload state.

### show isis int-counters

The following table describes the fields in the output for the **show isis int-counters** command.

Parameter	Description
IFIDX	Shows the interface index for the Ethernet or MLT interface.
LEVEL	Shows the level of the IS-IS interface.
AUTH FAILS	Shows the number of times authentication has failed per interface.
ADJ CHANGES	Shows the number of times the adjacencies have changed.
INIT FAILS	Shows the number of times the adjacency has failed to establish.
REJ ADJ	Shows the number of times the adjacency was rejected by another router.
ID LEN	Shows the ID field length mismatches.
MAX AREA	Shows the maximum area address mismatches.
LAN DIS CHANGES	Shows the number of times the DIS has changed.

### show isis int-l1-ctrl-pkts

The following table describes the fields in the output for the **show isis int-l1-ctrl-pkts** command.

Parameter	Description
IFIDX	Shows the interface index for the Ethernet or MLT interface.
DIRECTION	Shows the packet flow (Transmitted or Received).
HELLO	Shows the amount of interface-level Hello packets.
LSP	Shows the amount of LSP packets.
CSNP	Shows the amount of CSNPs.
PSNP	Shows the amount of PSNPs.

---

## Clearing ACL Statistics

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

**\* Note:**

The ACL statistics do not support security action on some hardware platforms. For more information, see [Configuring QoS and ACL-Based Traffic Filtering for VOSS](#).

**Procedure**

1. Enter Privileged EXEC mode:

```
enable
```

2. Enter the following command to clear default ACL statistics:

```
clear filter acl statistics default <acl-id>
```

3. Enter the following command to clear global ACL statistics:

```
clear filter acl statistics global <acl-id>
```

4. Enter the following command to clear all ACL statistics:

```
clear filter acl statistics all
```

5. Enter the following command to clear statistics associated with a particular ACL, ACE, or ACE type:

```
clear filter acl statistics <acl-id> <ace-id> [qos] [security]
```

## Variable Definitions

Use the information in the following table to use the `clear filter acl statistics` command.

Variable	Value
<acl-id>	Specifies the ACL ID. Use the CLI Help to see the available range for your switch.
<ace-id>	Specifies the ACE ID. Different hardware platforms support different ACE ID ranges. Use the CLI Help to see the available range for your switch.

## Viewing ACE Statistics

View ACE statistics to ensure that the filter operates correctly.

**\* Note:**

The ACL statistics do not support security action on some hardware platforms. For more information, see [Configuring QoS and ACL-Based Traffic Filtering for VOSS](#).

**Procedure**

1. Enter Privileged EXEC mode:

```
enable
```

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

```
show filter acl statistics <acl-id> <ace-id> [qos] [security]
```

## 3. View all ACE statistics:

```
show filter acl statistics all
```

## 4. View default ACE statistics:

```
show filter acl statistics default [<acl-id>]
```

## 5. View global statistics for ACEs:

```
show filter acl statistics global [<acl-id>]
```

**Example**

View ACE statistics:

**\* Note:**

Based on your hardware platform, the output may display all the ACL packets or segregate them as QoS and security ACL packets.

```
Switch:1>enable
Switch:1#show filter acl statistics all

=====
                        Acl Global Statistics Table
=====
Acl Id  Acl Name    Acl Type  Acl Sec  Acl Sec  Acl QOS  Acl QOS
        Packets Bytes    Packets Bytes
-----
1       ACL-1       inVlan    0         0         0         0
2       ACL-2       inVlan    0         0         0         0

Displayed 2 of 2 entries

=====
                        Acl Default Statistics Table
=====
Acl Id  Acl Name    Acl Type  Acl Sec  Acl Sec  Acl QOS  Acl QOS
        Packets Bytes    Packets Bytes
-----
1       ACL-1       inVlan    0         0         0         0
2       ACL-2       inVlan    0         0         0         0

Displayed 2 of 2 entries

--More-- (q = quit)

Switch:1#show filter acl statistics default

=====
                        Acl Default Statistics Table
=====
Acl Id  Acl Name    Acl Type  Acl Sec  Acl Sec  Acl QOS  Acl QOS
        Packets Bytes    Packets Bytes
-----
1       ACL-1       inVlan    0         0         0         0
2       ACL-2       inVlan    0         0         0         0
```

```

Displayed 2 of 2 entries
Switch:1#show filter acl statistics global 2
=====
                        Acl Global Statistics Table
=====
Acl Id  Acl Name    Acl Type  Acl Sec  Acl Sec  Acl QoS  Acl QoS
      Packets  Bytes    Packets  Bytes
-----
2      ACL-2      inVlan    0         0         0         0
=====
Displayed 1 of 1 entries

```

## Variable Definitions

Use the data in the following table to use the `show filter acl statistics` command.

Variable	Value
<code>&lt;acl-id&gt;</code>	Specifies the ACL ID. Use the CLI Help to see the available range for your switch.
<code>&lt;ace-id&gt;</code>	Specifies the ACE ID. Different hardware platforms support different ACE ID ranges. Use the CLI Help to see the available range for your switch.

## Job Aid

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

**Table 23: show filter acl statistics default 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.
Acl Sec Packets	Specifies the ACL secondary packets.
Acl Sec Bytes	Specifies the ACL secondary bytes.
Acl QoS Packets	Specifies the ACL QoS packets.
Acl QoS Bytes	Specifies the ACL QoS bytes.

## 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

```
Switch: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 24: show spanning-tree mstp statistics field descriptions**

Parameter	Description
MSTP Up Count	The number of times the MSTP port has been enabled. A Trap is generated on the occurrence of this event.
MSTP Down Count	The number of times the MSTP port 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 25: show spanning-tree rstp statistics field descriptions**

Parameter	Description
RSTP Up Count	The number of times RSTP port has been enabled. A Trap is generated on the occurrence of this event.
RSTP Down Count	The number of times RSTP port 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.

 **Note:**

Slot and port information can differ depending on hardware platform.

**Procedure**

View RSTP statistics on a port:

```
show spanning-tree rstp port statistics [{slot/port[/sub-port]}[-slot/port[/sub-port]][,...]]
```

**Example**

View RSTP statistics:

```
Switch:1#show spanning-tree rstp port statistics
=====
RSTP Port Statistics
```



```

=====
Port Number                : 4/1
Number of Fwd Transitions  : 0
Rx RST BPDUs Count        : 0
Rx Config BPDUs Count     : 0
Rx TCN BPDUs Count        : 0
Tx RST BPDUs Count        : 0
Tx Config BPDUs Count     : 0
Tx TCN BPDUs Count        : 0
Invalid RST BPDUs Rx Count : 0
Invalid Config BPDUs Rx Count : 0
Invalid TCN BPDUs Rx Count : 0
Protocol Migration Count   : 0
Port Number                : 4/2
Number of Fwd Transitions  : 0
Rx RST BPDUs Count        : 0
Rx Config BPDUs Count     : 0
Rx TCN BPDUs Count        : 0
Tx RST BPDUs Count        : 0
Tx Config BPDUs Count     : 0
--More-- (q = quit)

```

## Variable Definitions

Use the data in the following table to use the `show spanning-tree rstp port statistics` command.

Variable	Value
{slot/port[/sub-port][/-slot/port[/sub-port]][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

## Job Aid

The following table describes output for the `show spanning-tree rstp port statistics` command.

**Table 26: 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.

*Table continues...*

Parameter	Description
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

```
Switch: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 27: 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.

---

## Viewing vIST Statistics

View virtual IST (vIST) statistics for the switch.

### Procedure

1. Enter Privileged EXEC mode:  
`enable`
2. Display the vIST statistics:  
`show virtual-ist stat`
3. To clear the vIST statistics:  
`clear virtual-ist stats`

**Example**

```
Switch:1#show virtual-ist stat
=====
                        IST Message Statistics
=====
PROTOCOL MESSAGE          COUNT
-----
Ist Down                  : 0
Hello Sent                 : 0
Hello Recv                 : 0
Learn MAC Address Sent    : 0
Learn MAC Address Recv    : 0
MAC Address AgeOut Sent   : 0
MAC Address AgeOut Recv   : 0
MAC Address Expired Sent  : 0
MAC Address Expired Recv  : 0
Delete Mac Address Sent   : 0
Delete Mac Address Recv   : 0
Smlt Down Sent            : 0
Smlt Down Recv            : 0
Smlt Up Sent              : 0
Smlt Up Recv              : 0
Send MAC Address Sent     : 0
Send MAC Address Recv     : 0
IGMP Sent                 : 0
IGMP Recv                 : 0
Port Down Sent            : 0
Port Down Recv            : 0
Request MAC Table Sent    : 0
Request MAC Table Recv    : 0
Unknown Msg Type Recv    : 0
Mlt Table Sync Req Sent   : 0
Mlt Table Sync Req Recv   : 0
Mlt Table Sync Sent       : 0
Mlt Table Sync Recv       : 0
Port Update Sent          : 0
Port Update Recv          : 0
Entry Update Sent         : 0
Entry Update Recv         : 0
Dialect Negotiate Sent    : 0
Dialect Negotiate Recv    : 0
Update Response Sent      : 0
Update Response Recv      : 0
Transaction Que HiWaterM  : 0
Poll Count Hi Water Mark  : 0
```

**Job Aid**

The following table describes the output for the `show virtual-ist stat` command.

**Table 28: show virtual-ist stat field descriptions**

Parameter	Description
Ist Down	The count of how many sessions between the two peering switches went down since last boot.
Hello Sent	The count of transmitted hello messages.
Hello Recv	The count of received hello messages.

*Table continues...*

Parameter	Description
Learn MAC Address Sent	The count of transmitted learned MAC address messages.
Learn MAC Address Recv	The count of received learned MAC address messages.
MAC Address AgeOut Sent	The count of transmitted aging out MAC address messages.
MAC Address AgeOut Recv	The count of received aging out MAC address messages.
MAC Address Expired Sent	The count of transmitted MAC address age expired messages.
MAC Address Expired Recv	The count of received MAC address age expired messages.
Delete Mac Address Sent	The count of transmitted MAC address deleted messages.
Delete Mac Address Recv	The count of received MAC address deleted messages.
Smlt Down Sent	The count of transmitted SMLT down messages.
Smlt Down Recv	The count of received SMLT down messages.
Smlt Up Sent	The count of transmitted SMLT up messages.
Smlt Up Recv	The count of received SMLT up messages.
Send MAC Address Sent	The count of transmitted send MAC table messages.
Send MAC Address Recv	The count of received send MAC table messages.
IGMP Sent	The count of transmitted IGMP messages.
IGMP Recv	The count of received IGMP messages.
Port Down Sent	The count of transmitted port down messages.
Port Down Recv	The count of received port down messages.
Request MAC Table Sent	The count of transmitted MAC table request messages.
Request MAC Table Recv	The count of received MAC table request messages.
Unknown Msg Type Recv	The count of received unknown message type messages.
Mlt Table Sync Req Sent	The count of transmitted MLT table sync request messages.
Mlt Table Sync Req Recv	The count of received MLT table sync request messages.
Mlt Table Sync Sent	The count of transmitted MLT table sync messages.
Mlt Table Sync Recv	The count of received MLT table sync messages.
Port Update Sent	The count of transmitted port update messages.
Port Update Recv	The count of received port update messages.

*Table continues...*

Parameter	Description
Entry Update Sent	The count of transmitted entry update messages.
Entry Update Recv	The count of received entry update messages.
Dialect Negotiate Sent	The count of transmitted protocol ID messages.
Dialect Negotiate Recv	The count of received protocol ID messages.
Update Response Sent	The count of transmitted update response messages.
Update Response Recv	The count of received update response messages.
Transaction Que HiWaterM	The count of transaction queue high watermark messages.
Poll Count Hi Water Mark	The count of poll count high watermark messages.

## Showing RADIUS Server Statistics

### 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. Enter Privileged EXEC mode:  
`enable`
2. Display RADIUS server statistics:  
`show radius-server statistics`
3. Clear server statistics:  
`clear radius statistics`

### Example

```
Switch:1#show radius-server statistics
Responses with invalid server address: 0
  Radius Server(UsedBy) : 192.0.2.58(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 : 192.0.2.32

Radius Server(UsedBy) : 192.0.2.58 (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 : 192.0.2.32

--More-- (q = quit)

```

## Job Aid

The following table shows the field descriptions for the `show radius-server statistics` command output.

**Table 29: show radius-server statistics command fields**

Parameter	Description
RADIUS Server	The IP address of the RADIUS server.
AccessRequests	Number of access-response packets sent to the server; does not include retransmissions.
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.

*Table continues...*

Parameter	Description
	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

```
Switch:1(config)#show rmon stats
```

```
=====
                                     Rmon Ether Stats
=====
INDEX  PORT   OWNER
-----
1      1/10   monitor
```

## Job Aid

The following table describes parameters in the output for the `show rmon stats` command.

**Table 30: show rmon stats field descriptions**

Parameter	Description
Index	An index that uniquely identifies an entry in the Ethernet statistics table.
Port	Identifies the source of the data that this entry analyzes.
Owner	The entity that configured this entry and is therefore using the assign resources.

## Showing OSPF Error Statistics on a Port

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



## Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display extended information about OSPF errors for the specified port or for all ports:

```
show interfaces GigabitEthernet error ospf [{slot/port[/sub-port]}[-slot/port[/sub-port]][,...]]
```

## Variable Definitions

Use the following table to help you use the **show interfaces GigabitEthernet error ospf** command.

Variable	Value
{slot/port[/sub-port]}[-slot/port[/sub-port]][,...]	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

## Job Aid

The following table describes the output for the **show interfaces GigabitEthernet error ospf** command.

**Table 31: 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

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

## Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display OSPF interface statistics:

```
show ip ospf ifstats [detail vrf WORD<0-16> vrfids WORD<0-512>]
[mismatch vrf WORD<0-16> vrfids WORD<0-512>] [vlan <1-4059>] [vrf
WORD<0-16>] [vrfids WORD<0-512>]
```

## Example

```
Switch:1#show ip ospf ifstats
```

```
=====
                        OSPF Interface Statistics - GlobalRouter
=====
INTERFACE          ---HELLOS---  ---DBS---  -LS REQ--  --LS UPD---  --LS ACK---
                   RX      TX      RX      TX      RX      TX      RX      TX      RX      Tx
-----
192.0.2.3           76035  76355  33      32      4       9      2483  2551  2525  1247
192.0.2.8           76038  76349  0        0        0        0        0      0      0      0
```

## Variable Definitions

Use this table to help you use the `show ip ospf ifstats` command.

Variable	Value
detail	Shows detailed information.
mismatch	Shows the number of times the area ID is not matched.
vlan <1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
vrf WORD<1-16>	Specifies 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 32: 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.

*Table continues...*

Field	Description
HELLOS TX	Indicates the number of hello packets transmitted by this interface.
DBS RX	Indicates the number of database descriptor packets received by this interface.
DBS TX	Indicates the number of database descriptor packets transmitted by this interface.
LS REQ	Indicates the number of link state request packets received by this interface.
LS TX	Indicates the number of link state request packets transmitted by this interface.
LS UDP RX	Indicates the number of link state update packets received by this interface.
LS UDP TX	Indicates the number of link state update packets transmitted by this interface.
LS ACK RX	Indicates the number of link state acknowledge packets received by this interface.
LS ACK TX	Indicates the number of link state acknowledge packets transmitted by this interface.
VERSION	Indicates the OSPF version.
AREA	Indicates the OSPF area.
AUTHTYPE	Indicates the OSPF authentication type.
AUTHFAIL	The count of authentication fail messages.
NETMASK	Indicates the net mask.
HELLO	The count of Hello messages.
DEADTRR OPTION	The dead TRR option.

## Viewing OSPF Range Statistics

Use statistics to help you monitor Open Shortest Path First (OSPF) performance. You can also use statistics in troubleshooting procedures. OSPF range statistics include area ID, range network address, range subnet mask, range flag, and LSDB type.

### Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display the OSPF range statistics:

```
show ip ospf stats [vrf WORD<0-16>] [vrfs WORD<0-512>]
```

### Example

```
Switch:1#show ip ospf stats
```

```
=====
                        OSPF Statistics - GlobalRouter
=====
NumBufAlloc: 239603
NumBufFree: 239603
NumBufAllocFail: 0
```

## Statistics

```
NumBufFreeFail: 0
  NumTxPkt: 239655
  NumRxPkt: 317562
NumTxDropPkt: 0
NumRxDropPkt: 0
  NumRxBadPkt: 0
  NumSpfRun: 47
  LastSpfRun: 2 day(s), 04:18:58
  LsdbTblSize: 16
NumAllocBdDDP: 24
NumFreeBdDDP: 24
  NumBadLsReq: 0
NumSeqMismatch: 3
  NumOspfRoutes: 4
  NumOspfAreas: 1
NumOspfAdjacencies: 3

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

## Variable Definitions

Use the data in the following table to use the `show ip ospf stats` command.

Variable	Value
vrf <i>WORD</i> <0-16>	Specifies a VRF instance by VRF name.
vrfids <i>WORD</i> <0-16>	Specifies a VRF or range of VRFs by ID.

## Job Aid

The following table describes the show command output.

**Table 33: 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.

*Table continues...*

Parameter	Description
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.

---

## Clearing IP OSPF Statistics

Use the following procedure to clear all IPv4 OSPF statistics.

### Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Clear IPv4 OSPF statistics:

```
clear ip ospf stats [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

## Variable Definitions

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

Variable	Value
vrf WORD<0-16>	Specifies a VRF instance by VRF name.
vrfids WORD<0-512>	Specifies the ID of the VRF.

---

## Viewing Basic OSPF Statistics for a Port

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

### Procedure

1. Log on to the switch to enter User EXEC mode.
2. View basic OSPF statistics:

```
show ports statistics ospf main [{slot/port[/sub-port]}[-slot/port[/sub-port]][,...]
```

**Example**

View basic OSPF statistics:

```
Switch:1>enable
Switch:1#show ports statistics ospf main

=====
Port Stats Ospf
=====
PORT_NUM RX_HELLO TX_HELLO RXDB_DESCR TXDB_DESCR RXLS_UPDATE TXLS_UPDATE
-----
1/3      0          0          0          0          0          0
```

**Variable Definitions**

Use the data in the following table to use the `show ports statistics ospf main` command.

Variable	Value
{slot/port[/sub-port][/-slot/port[/sub-port]][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

**Job Aid**

The following table describes the output for the `show ports statistics ospf main` command.

**Table 34: show ports statistics ospf main output description**

Field	Description
PORT_NUM	Indicates the port number.
RX_HELLO	Indicates the number of hello packets this interface receives.
TX_HELLO	Indicates the number of hello packets this interface transmitted.
RXDB_DESCR	Indicates the number of database descriptor packets this interface receives.
TXDB_DESCR	Indicates the number of database descriptor packets this interface transmitted.
RXLS_UPDATE	Indicates the number of link state update packets this interface receives.
TXLS_UPDATE	Indicates the number of link state update packets this interface transmitted.

**Showing Extended OSPF Statistics**

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

## Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display extended OSPF information about the specified port or for all ports:

```
show ports statistics ospf extended [{slot/port[/sub-port]}[-slot/
port[/sub-port]][, ...]]
```

## Example

Display extended OSPF information:

```
Switch:1>enable
Switch:1#show ports statistics ospf extended

=====
Port Stats Ospf Extended
=====
PORT_NUM  RXLS_REQS  TXLS_REQS  RXLS_ACKS  TXLS_ACKS
-----
1/3        0           0           0           0
```

## Variable Definitions

Use the data in the following table to use the **show ports statistics ospf extended** command.

Variable	Value
{slot/port[/sub-port]}[-slot/port[/sub-port]][, ...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

## Job Aid

The following table describes the output for the **show ports statistics ospf extended** command.

**Table 35: 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 Ingress Port-rate Limit Statistics

Use this procedure to view the ingress port-rate limit statistics. The system displays the statistics of the dropped packets and bytes.

**\* Note:**

This command is not available on all hardware platforms.

**Procedure**

1. Enter Privileged EXEC mode:

```
enable
```

2. View the ingress port-rate limit statistics:

```
show interfaces gigabitethernet statistics rate-limiting [port
{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...]]
```

**Example**

```
Switch:1# show interfaces gigabitethernet statistics rate-limiting 1/1
```

QOS Interface Ingress Rate-Limiting Stats				
PORT	DROPPING PKTS RATE	DROPPING BYTES RATE	DROPPING PKTS	DROPPING BYTES
1/1	9224		9260507	
1430758		1436481032		

## Variable Definitions

Use the data in the following table to use the **show interfaces gigabitethernet statistics rate-limiting** command.

Variable	Value
{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

## Viewing Ingress Policer Statistics

Use this procedure to view the ingress policer statistics. The system displays individual policer statistics for specific ports to manage network performance.



**\* Note:**

This command is not available on all hardware platforms.

**Procedure**

1. Enter Privileged EXEC mode:

```
enable
```

2. View the policer statistics:

```
show interfaces gigabitethernet statistics policer {slot/port[/sub-
port] [-slot/port[/sub-port]] [,...]}
```

**Example**

```
Switch:1# show interfaces gigabitethernet statistics policer 1/3
```

QoS Ingress Port Policer Stats					
PORT NUM	TOTAL PKTS	TOTAL BYTES	YELLOW BYTES	RED BYTES	
1/3	420	31628	0	0	

**Variable Definitions**

Use the data in the following table to use the **show interfaces gigabitethernet statistics policer** command.

Variable	Value
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

**Viewing the Management Port Statistics**

Use this procedure to view the management port statistics.

**\* Note:**

This procedure only applies to hardware with a dedicated, physical management interface.

**Procedure**

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. View the management port statistics:

```
show interfaces mgmtethernet statistics
```

**Example**

View management port statistics:

```
Switch:1#show interfaces mgmtethernet statistics
=====
Port Stats Interface
=====
PORT   IN      OUT      IN      OUT
NUM   OCTETS OCTETS   PACKET  PACKET
-----
mgmt  7222116 44282   81789   586
PORT   IN      OUT      IN      OUT      OUTLOSS
NUM   FLOWCTRL FLOWCTRL PFC     PFC     PACKETS
-----
mgmt   0        0        0       0       0
```

## Viewing IP VRRPv3 Statistics

Use the following procedure to view IP VRRPv3 statistics to monitor network performance.

**Procedure**

1. Enter Privileged EXEC mode:

```
enable
```

2. Enter the following command to view VRRP statistics:

```
show ip vrrp statistics version <2-3>
```

3. Enter the following command to view VRRP statistics for the specified VRF:

```
show ip vrrp statistics vrf WORD<1-16> version <2-3>
```

4. Enter the following command to view VRRP statistics for the specified virtual router:

```
show ip vrrp statistics vrfids WORD<0-512> version <2-3>
```

**Example**

View IP VRRPv3 statistics:

```
Switch:1#show ip vrrp statistics
=====
VRRP Global Stats - GlobalRouter
=====
CHK_SUM_ERR  VERSION_ERR  VRID_ERR  VRRP_VERSION
-----
0             0            0         2
0             0            0         3
=====
```

```

VRRP Interface Stats - GlobalRouter
=====
VRRP ID  P/V      BECOME_MASTER ADVERTISE_RCV VERSION
-----
3         3         1             0             2
2         1/1       1             0             3

VRRP ID  P/V      ADVERTISE_INT_ERR TTL_ERR      PRIO_0_RCV  VERSION
-----
3         3         0             0             0           2
2         1/1       0             0             0           3

VRRP ID  P/V      PRIO_0_SENT  INVALID_TYPE_ERR ADDRESS_LIST_ERR UNKNOWN_AUTHTYPE  VERSION
-----
3         3         0             0             0           0           2
2         1/1       0             0             0           0           3

VRRP ID  P/V      AUTHTYPE_ERR  PACKLEN_ERR  VERSION
-----
3         3         0             0             2
2         1/1       0             0             3

```

## Variable Definitions

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

Variable	Value
<code>version</code>	Configures the VRRP version on the specified interface.
<code>&lt;2-3&gt;</code>	Specifies the version of VRRP (2 or 3) to be configured on the specified interface.
<code>vrf WORD&lt;1-16&gt;</code>	Specifies the name of the VRF.
<code>vrfids WORD&lt;0-512&gt;</code>	Specifies the ID of the VRF, and is an integer in the range of 0-512.

## Clearing IPv4 MSDP Statistics

Use the following procedure to clear all IPv4 Multicast Source Discovery Protocol (MSDP) statistics for all peers or a specific peer.

### About this task

The switch supports this command for local management VRF or global routing table (GRT). If you do not specify a VRF or VRF ID, the switch defaults to GRT.

### Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Clear IPv4 MSDP statistics for all peers:

```
clear ip msdp statistics [vrf WORD<0-16>] [vrfids WORD<0-512>]
```

3. Clear IPv4 MSDP statistics for a specific peer:

```
clear ip msdp statistics {A.B.C.D.} [vrf WORD<0-16>] [vrfids WORD<0-512>]
```

## Variable Definitions

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

Variable	Value
vrf WORD<0-16>	Specifies a VRF instance by VRF name.
vrfids WORD<0-512>	Specifies the ID of the VRF.
{A.B.C.D.}	Specifies the IPv4 MSDP address for a specific peer.

## Viewing IPv4 ICMP Statistics

View the collective IPv4 ICMP statistics for all VRF instances.

### Procedure

1. Log on to the switch to enter User EXEC mode.
2. View IPv4 ICMP statistics:

```
show ip icmp statistics
```

### Example

View IPv4 ICMP statistics:

```
Switch:1#show ip icmp statistics

Icmp Statistics:
-----
ICMP IN STATISTICS (includes GRT and all VRF instances) :
IcmpInMsgs           = 0
IcmpInErrors         = 0
IcmpInDestUnreachs  = 0
IcmpInTimeExcds     = 0
IcmpInParmProbs     = 0
IcmpInSrcQuenchs    = 0
IcmpInRedirects     = 0
IcmpInEchos         = 0
IcmpInEchoReps      = 0
IcmpInTimestamps    = 0
IcmpInTimestampReps = 0
IcmpInAddrMasks     = 0
IcmpInAddrMaskReps  = 0
ICMP OUT STATISTICS (includes GRT and all VRF instances) :
IcmpOutMsgs         = 0
IcmpOutErrors       = 0
IcmpOutDestUnreachs = 0
IcmpOutTimeExcds   = 0
IcmpOutParmProbs    = 0

IcmpOutSrcQuenchs   = 0
IcmpOutRedirects    = 0
```

```

IcmpOutEchos          = 0
IcmpOutEchoReps      = 0
IcmpOutTimestamps    = 0
IcmpOutTimestampReps = 0
IcmpOutAddrMasks     = 0
IcmpOutAddrMaskReps  = 0
0 messages dropped due to rate limiting

```

## Clearing IPv6 Statistics

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

### Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Clear all the IPv6 statistics:

```
clear ipv6 statistics all [vrf WORD<1-16> | vrfids WORD<0-512>]
```

3. Clear interface statistics:

```
clear ipv6 statistics interface [general|icmp] [gigabitethernet
{slot/port[/sub-port]}>| loopback <1-256> | mgmtethernet mgmt |
tunnel <1-2000> | vlan <1-4059>] [vrf WORD<1-16> | vrfids WORD<0-
512>]
```

4. Clear TCP statistics:

```
clear ipv6 statistics tcp [vrf WORD<1-16> | vrfids WORD<0-512>]
```

5. Enter the following command to clear UDP statistics:

```
clear ipv6 statistics udp [vrf WORD<1-16> | vrfids WORD<0-512>]
```

## Variable Definitions

Use the information in the following table to use the `clear ipv6 statistics` commands.

Variable	Value
gigabitethernet {slot/port[/sub-port]}	Identifies a single slot and port. If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
loopback <1-256>>	Identifies a loopback interface.

*Table continues...*

Variable	Value
mgmtethernet mgmt	Identifies the management interface. This parameter only applies to hardware with a dedicated, physical management interface.
tunnel <1-2000>	Identifies a 6in4 tunnel ID.
vlan<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
vrf WORD<1-16>	Specifies the VRF name.
vrfids WORD<0-512>	Specifies the VRF ID.

## Viewing IPv6 ICMP Statistics

View IPv6 ICMP statistics on an interface for ICMP messages sent over a particular interface.

### Procedure

1. Log on to the switch to enter User EXEC mode.
2. View IPv6 ICMP statistics:

```
show ipv6 interface icmpstatistics [gigabitethernet <slot/port[/sub-
port]> | loopback <1-256> | mgmtethernet mgmt|tunnel <1-2000> | vlan
<1-4059>] [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

### Example

View ICMP statistics:

```
Switch:1>show ipv6 interface icmpstatistics
=====
Icmp Stats
=====

Icmp stats for IfIndex = 192

IcmpInMsgs: 0
IcmpInErrors: 0
IcmpInDestUnreachs : 0
IcmpInAdminProhibs : 0
IcmpInTimeExcds : 0
IcmpInParmProblems : 0
IcmpInPktTooBigs : 0
IcmpInEchos : 0
IcmpInEchoReplies : 0
IcmpInRouterSolicits : 0
IcmpInRouterAdverts : 0
InNeighborSolicits : 0
InNbrAdverts : 0
```

```
IcmpInRedirects : 0
IcmpInGroupMembQueries : 0
IcmpInGroupMembResponses : 0
```

## Variable Definitions

Use the data in the following table to use the `show ipv6 interface icmpstatistics` command

Variable	Value
<1-4059>	Shows ICMP statistics for the specific interface index. If you do not specify an interface index, the command output includes all IPv6 ICMP interfaces.  Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
gigabitethernet {slot/port[/sub-port]}	Identifies a single slot and port. If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
loopback <1-256>>	Identifies a loopback interface.
mgmtethernet mgmt	Identifies the management interface. This parameter only applies to hardware with a dedicated, physical management interface.
tunnel <1-2000>	Identifies a 6in4 tunnel ID.
vrf WORD<1-16>	Specifies the VRF name.
vrfids WORD<0-512>	Specifies the VRF ID.

## Viewing IPv6 DHCP Relay Statistics

Display individual IPv6 DHCP Relay statistics for specific interfaces to manage network performance.

### Procedure

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

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

**\* Note:**

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

**Example**

```
Switch:1#show ipv6 dhcp-relay counters
=====
                                DHCPv6 Counters
=====
INTERFACE                        REQUESTS    REPLIES
-----
1111:0:0:0:0:0:1111              1          1
```

**Job Aid**

The following table explains the output of the `show ipv6 dhcp-relay counters` command.

**Table 36: 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.

**Viewing IPv6 OSPF Statistics**

View OSPF statistics to analyze trends.

**Procedure**

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

```
show ipv6 ospf statistics [vrf WORD <1-16>] [vrfids WORD <0-512>]
```

**Example**

View IPv6 OSPF statistics:

```
Switch:1>enable
Switch:1#show ipv6 ospf statistics
=====
                                OSPFv3 Statistics - GlobalRouter
=====
      NumTxPkt: 9958
      NumRxPkt: 8982
      NumTxDropPkt: 33
      NumRxDropPkt: 0
      NumRxBadPkt: 0
      NumSpfRun: 42
      LastSpfRun: 0 day(s), 02:44:32
      LsdbTblSize: 45
      NumBadLsReq: 0
      NumSeqMismatch: 0
      NumOspfAdjacencies: 7
```



## Job Aid

The following table explains the output of the `show ipv6 ospf statistics` command.

Field	Description
NumTxPkt	Shows the count of sent packets.
NumRxPkt	Shows the count of received packets.
NumTxDropPkt	Shows the count of sent, dropped packets.
NumRxDropPkt	Shows the count of received, dropped packets.
NumRxBadPkt	Shows the count of received, bad packets.
NumSpfRun	Shows the count of intra-area route table updates with calculations using this area link-state 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.

## Variable Definitions

Use the data in the following table to use the `show ipv6 ospf stats` command.

Variable	Value
vrf <i>WORD</i> <1-16>	Specifies the VRF name.
vrfids <i>WORD</i> <0-512>	Specifies VRF IDs.

---

## Clearing IPv6 OSPF Statistics

### About this task

Use the following procedure to clear all IPv6 OSPF statistics.

### Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Clear IPv6 statistics:

```
clear ipv6 ospf stats [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

## Variable Definitions

Use the data in the following table to use the `clear ipv6 ospf stats` command.

Variable	Value
vrf <i>WORD</i> <1-16>	Specifies the VRF name.
vrfids <i>WORD</i> <0-512>	Specifies VRF IDs.

## Viewing IPv6 Statistics on an Interface

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

### Procedure

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

```
show ipv6 interface statistics [gigabitethernet <slot/port[/sub-
port]> | loopback <1-256> | mgmtethernet mgmt | tunnel <1-2000> |
vlan <1-4059>] [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

### Example

View IPv6 statistics on an interface:

```
Switch:1>enable
Switch:1#show ipv6 interface statistics

=====
                          Interface Stats
=====

If Stats for mgmt, IfIndex = 64

InReceives: 404
InHdrErrors: 0
InTooBigErrors : 0
InNoRoutes : 0
InAddrErrors : 0
InUnknownProtos : 0
InTruncatedPkts : 0
InDiscards : 0
InDelivers : 404
OutForwDatagrams : 0
OutRequests : 417
OutDiscards : 0
OutFragOKs : 0
OutFragFails : 0
OutFragCreates : 0

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

## Variable Definitions

Use the data in the following table to use the `show ipv6 interface statistics` command

Variable	Value
<code>gigabitethernet {slot/port[/sub-port]}</code>	Identifies a single slot and port. If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
<code>loopback &lt;1-256&gt;&gt;</code>	Identifies a loopback interface.
<code>mgmtethernet mgmt</code>	Identifies the management interface. This parameter only applies to hardware with a dedicated, physical management interface.
<code>tunnel &lt;1-2000&gt;</code>	Identifies a 6in4 tunnel ID.
<code>vlan &lt;1-4059&gt;</code>	Shows statistics for the specific interface index. If you do not specify an interface index, the command output includes all IPv6 interfaces.  Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
<code>vrf WORD&lt;1-16&gt;</code>	Specifies the VRF name.
<code>vrfids WORD&lt;0-512&gt;</code>	Specifies the VRF ID.

## Displaying IPsec Statistics

Use the following procedure to clear Internet Protocol Security (IPsec) system statistics counters and display IPsec statistics on an interface. The device only clears system statistics counters on system reboot.

### Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display IPsec statistics for the system:  
`show ipsec statistics system`
3. Display IPsec statistics for an Ethernet interface:  
`show ipsec statistics gigabitethernet {slot/port[/sub-port]}[-slot/port[/sub-port]][,...]`
4. Display IPsec statistics for a VLAN interface:  
`show ipsec statistics vlan <1-4059>`
5. Display statistics for IPsec on the management interface:

**\* Note:**

This step only applies to hardware with a dedicated, physical management interface.

You can view the statistics on both the primary and standby management interface, if supported by your hardware platform.

```
show ipsec statistics mgmtethernet <mgmt | mgmt2>
```

6. Display statistics for IPsec on the loopback interface:

```
show ipsec statistics loopback <1-256>
```

7. Clear IPsec system statistics counters:

```
clear ipsec stats all
```

**Example**

Display IPsec statistics. Output is partial due to length.

```
Switch:1>show ipsec statistics system
```

```
=====
IPSEC Global Statistics
=====
InSuccesses           = 0
InSPViolations        = 0
InNotEnoughMemories   = 0
InAHESPReplays        = 0
InAHFailures          = 0
InESPFailures         = 0
OutSuccesses          = 0
OutSPViolations       = 0
OutNotEnoughMemories  = 0
generalError          = 0
InAHSuccesses         = 0
InESPSuccesses        = 0
OutAHSuccesses        = 0
OutESPSuccesses       = 0
OutKBytes             = 0
OutBytes              = 0
InKBytes              = 0
InBytes               = 0
--More-- (q = quit)
```

```
Switch:1>show ipsec statistics gigabitethernet 1/13
```

```
=====
Ipsec Port Stats
=====
Iindex                = 204
InSuccesses           = 0
InSPViolations        = 0
InNotEnoughMemories   = 0
InAHESPReplays        = 0
InAHFailures          = 0
InESPFailures         = 0
OutSuccesses          = 0
OutSPViolations       = 0
OutNotEnoughMemories  = 0
generalError          = 0
```

```
Switch:1>show ipsec statistics vlan 1
```

```

=====
                        Isec  Vlan  Stats
=====
Ifindex                = 2049
InSuccesses            = 0
InSPViolations         = 0
InNotEnoughMemories   = 0
InAHESPReplays        = 0
InAHFailures          = 0
InESPFailures         = 0
OutSuccesses          = 0
OutSPViolations       = 0
OutNotEnoughMemories  = 0
generalError          = 0
=====

```

### Display IPsec statistics for a loopback interface:

```
Switch:1>show ipsec statistics loopback 1
```

```

=====
                        Isec  LoopBack  Stats
=====
Ifindex                = 1344
InSuccesses            = 0
InSPViolations         = 0
InNotEnoughMemories   = 0
InAHESPReplays        = 0
InESPReplays          = 0
InAHFailures          = 0
InESPFailures         = 0
OutSuccesses          = 0
OutSPViolations       = 0
OutNotEnoughMemories  = 0
generalError          = 0
=====

```

```
Switch:1>show ipsec statistics mgmtethernet mgmt
```

```

=====
                        Isec  Port  Stats
=====
Ifindex                = 64
InSuccesses            = 0
InSPViolations         = 0
InNotEnoughMemories   = 0
InAHESPReplays        = 0
InESPReplays          = 0
InAHFailures          = 0
InESPFailures         = 0
OutSuccesses          = 0
OutSPViolations       = 0
OutNotEnoughMemories  = 0
generalError          = 0
=====

```

```
Switch:1>show ipsec statistics mgmtethernet mgmt2
```

```

=====
                        Isec  Port  Stats
=====
Ifindex                = 128
InSuccesses            = 0
InSPViolations         = 0
InNotEnoughMemories   = 0
InAHESPReplays        = 0
InESPReplays          = 0
=====

```

## Statistics

```
InAHFailures           = 0
InESPFailures          = 0
OutSuccesses           = 0
OutSPViolations        = 0
OutNotEnoughMemories   = 0
generalError           = 0
```

## Variable Definitions

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

Variable	Value
{slot/port[/sub-port][/-slot/port[/sub-port]][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
loopback <1-256>	Identifies the loopback interface.
mgmtethernet mgmt	Identifies the interface as the management interface.
system	Shows statistics for the entire system.
vlan <1-4059>	Specifies the VLAN.

## Job Aid

The following table describes the fields in the output for the `show ipsec statistics system` command.

Parameter	Description
InSuccesses	Specifies the number of ingress packets IPsec successfully carries.
InSPViolations	Specifies the number of ingress packets IPsec discards since boot time because of a security policy violation.
InNotEnoughMemories	Specifies the number of ingress packets IPsec discards since boot time because not enough memory is available.
InAHESPReplays	Specifies the number of ingress packets IPsec discards since boot time because the encapsulating security payload (ESP) replay check fails.
InAHFailures	Specifies the number of ingress packets IPsec discards since boot time because the AH authentication check fails.
InESPFailures	Specifies the number of ingress packets IPsec discards since boot time because the ESP authentication check fails.

*Table continues...*

Parameter	Description
OutSuccesses	Specifies the number of egress packets IPsec successfully carries since boot time.
OutSPViolations	Specifies the number of egress packets IPsec discards since boot time because a security policy violation occurs.
OutNotEnoughMemories	Specifies the number of egress packets IPsec discards since boot time because not enough memory is available since boot time.
generalError	Specifies a general error.
InAHSuccesses	Specifies the number of ingress packets IPsec carries because the AH authentication succeeds.
InESPSuccesses	Specifies the number of ingress packets IPsec carries since boot time because the ESP authentication succeeds.
OutAHSuccesses	Specifies the number of egress packets IPsec successfully carries since boot time.
OutESPSuccesses	Specifies the number of egress packets IPsec successfully carries since boot time.
OutKBytes	Specifies the total number of kilobytes on egress.
OutBytes	Specifies the total number of bytes on egress.
InKBytes	Specifies the total number of bytes on ingress.
InBytes	Specifies the total number of bytes on ingress.
TotalPacketsProcessed	Specifies the total number of packets processed.
TotalPacketsByPassed	Specifies the total number of packets bypassed.
OutAHFailures	Specifies the number of egress packets IPsec discards since boot time because the AH authentication check fails.
OutESPFailures	Specifies the number of egress packets IPsec discards since boot time because the ESP authentication check fails.
InMD5Hmacs	Specifies the number of inbound HMAC MD5 occurrences since boot time.
InSHA1Hmacs	Specifies the number of inbound HMAC SHA1 occurrences since boot time.
InAESXCBCs	Specifies the number of inbound AES XCBC MAC occurrences since boot time.
InAnyNullAuth	Specifies the number of inbound null authentication occurrences since boot time.
In3DESCBCs	Specifies the number of inbound 3DES CBC occurrences since boot time.

*Table continues...*

Parameter	Description
InAESCBCs	Specifies the number of inbound AES CBC occurrences since boot time.
InAESCTRs	Specifies the number of inbound AES CTR occurrences since boot time.
InAnyNullEncrypt	Specifies the number of inbound null occurrences since boot time. Used for debugging purposes.
OutMD5Hmacs	Specifies the number of outbound HMAC MD5 occurrences since boot time.
OutSHA1Hmacs	Specifies the number of outbound HMAC SHA1 occurrences since boot time.
OutAESXCBCs	Specifies the number of outbound AES XCBC MAC occurrences since boot time.
OutInAnyNullAuth	Specifies the number of outbound null authentication occurrences since boot time.
Out3DESCBCs	Specifies the number of outbound 3DES CBC occurrences since boot time.
OutAESCBCs	Specifies the number of outbound AES CBC occurrences since boot time.
OutAESCTRs	Specifies the number of outbound AES CTR occurrences since boot time.
OutInAnyNullEncrypt	Specifies the number of outbound null occurrences since boot time. Used for debugging purposes.

The following table describes the fields in the output for the `show ipsec statistics gigabitethernet {slot/port[-slot/port]} [, ...]` and `show statistics loopback <1-256>` commands.

Parameter	Description
Ifindex	Specifies the interface.
InSuccesses	Specifies the number of ingress packets IPsec successfully carries.
InSPViolations	Specifies the number of ingress packets IPsec discards since boot time because of a security policy violation.
InNotEnoughMemories	Specifies the number of ingress packets IPsec discards since boot time because not enough memory is available.
InAHESPReplays	Specifies the number of ingress packets IPsec discards since boot time because the encapsulating security payload (ESP) replay check fails.

*Table continues...*



Parameter	Description
InAHFailures	Specifies the number of ingress packets IPsec discards since boot time because the AH authentication check fails.
InESPFailures	Specifies the number of ingress packets IPsec discards since boot time because the ESP authentication check fails.
OutSuccesses	Specifies the number of egress packets IPsec successfully carries since boot time.
OutSPViolations	Specifies the number of egress packets IPsec discards since boot time because a security policy violation occurs.
OutNotEnoughMemories	Specifies the number of egress packets IPsec discards since boot time because not enough memory is available since boot time.
generalError	Specifies a general error.

The following table describes the fields in the output for the `show ipsec statistics vlan <1-4059>` command.

Parameter	Description
Ifindex	Specifies the interface.
InSuccesses	Specifies the number of ingress packets IPsec successfully carries.
InSPViolations	Specifies the number of ingress packets IPsec discards since boot time because of a security policy violation.
InNotEnoughMemories	Specifies the number of ingress packets IPsec discards since boot time because not enough memory is available.
InAHESPReplays	Specifies the number of ingress packets IPsec discards since boot time because the encapsulating security payload (ESP) replay check fails.
InAHFailures	Specifies the number of ingress packets IPsec discards since boot time because the AH authentication check fails.
InESPFailures	Specifies the number of ingress packets IPsec discards since boot time because the ESP authentication check fails.
OutSuccesses	Specifies the number of egress packets IPsec successfully carries since boot time.
OutSPViolations	Specifies the number of egress packets IPsec discards since boot time because a security policy violation occurs.

*Table continues...*

Parameter	Description
OutNotEnoughMemories	Specifies the number of egress packets IPsec discards since boot time because not enough memory is available since boot time.
generalError	Specifies a general error.

The following table describes the fields in the output for the `show ipsec statistics mgmtethernet` command.

Parameter	Description
lindex	Specifies the interface.
InSuccesses	Specifies the number of ingress packets IPsec successfully carries.
InSPViolations	Specifies the number of ingress packets IPsec discards since boot time because of a security policy violation.
InNotEnoughMemories	Specifies the number of ingress packets IPsec discards since boot time because not enough memory is available.
InAHESPReplays	Specifies the number of ingress packets IPsec discards since boot time because the AH replay check fails.
InESPReplays	Specifies the total number of ingress packets IPsec discards since boot time because the encapsulating security payload (ESP) replay check fails.
InAHFailures	Specifies the number of ingress packets IPsec discards since boot time because the AH authentication check fails.
InESPFailures	Specifies the number of ingress packets IPsec discards since boot time because the ESP authentication check fails.
OutSuccesses	Specifies the number of egress packets IPsec successfully carries since boot time.
OutSPViolations	Specifies the number of egress packets IPsec discards since boot time because a security policy violation occurs.
OutNotEnoughMemories	Specifies the number of egress packets IPsec discards since boot time because not enough memory is available since boot time.
generalError	Specifies a general error.

## Viewing IPv6 VRRP Statistics

View IPv6 VRRP statistics to monitor network performance

## Procedure

1. Log on to the switch to enter User EXEC mode.
2. View statistics for the device and for all interfaces:

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

## Example

View IPv6 VRRP statistics for VRID 1.

```
Switch:1(config)#show ipv6 vrrp statistics vrid 1

=====
                          VRRP Interface Stats - GlobalRouter
=====

VRID  P/V    BECOME_MASTER  ADVERTISE_RCV
-----
1      84      2              17372
1      85      2              17372
1      86      1              0
1      87      1              0
1      1001   2              17372

VRID  P/V    ADVERTISE_INT_ERR  TTL_ERR    PRIO_0_RCV
-----
1      84      0                  0          0
1      85      0                  0          0
1      86      0                  0          0
1      87      0                  0          0
1      1001   0                  0          0

VRID  P/V    PRIO_0_SENT    INVALID_TYPE_ERR  ADDRESS_LIST_ERR  UNKNOWN_AUTHTYPE
-----
--More-- (q = quit)
```

## 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.
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 37: show ipv6 vrrp statistics command output**

Heading	Description
CHK_SUM_ERR	Shows the number of VRRP packets received with an invalid VRRP checksum value.

*Table continues...*

Heading	Description
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. 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

*Table continues...*

Heading	Description
	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.

## Showing the EAPoL Status of the Device

Display the current device configuration.

### \* Note:

Use the `clear-stats` command to clear EAP or NEAP statistics.

### Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display the current device configuration by using the following command:

```
show eapol system
```

### Example

```
Switch:1#show eapol system
```

```

=====
                        Eapol System
=====
                        eap : enabled
                        non-eap-pwd-fmt : ip-addr.mac-addr.port-number
                        non-eap-pwd-fmt key :
                        non-eap-pwd-fmt padding : disabled
=====

```

## Showing EAPoL Authenticator Statistics

Display the authenticator statistics to manage network performance.

**\* Note:**

Use the `clear-stats` command to clear EAP or NEAP statistics.

**Procedure**

1. Log on to the switch to enter User EXEC mode.
2. Display the authenticator statistics:

```
show eapol auth-stats interface [gigabitEthernet [{slot/port[/sub-
port]}[-slot/port[/sub-port]][,...]]]
```

**Example**

**\* Note:**

Slot and port information can differ depending on hardware platform.

```
Switch:1#show eapol auth-stats interface
```

```
=====
                        Eap Authenticator Statistics
=====
```

PORT	EAP RCVD	AUTH-EAP TX	START RCVD	LOGOFF RCVD	INVALID FRAMES	LENGTH ERROR	LAST-RX VER	LAST-RX SRC
1/1	716	1074	0	0	0	0	1	18:a9:05:b1:04:ce
1/2	0	0	0	0	0	0	0	00:00:00:00:00:00
1/3	0	0	0	0	0	0	0	00:00:00:00:00:00
1/4	0	5	0	0	0	0	0	00:00:00:00:00:00
1/5	0	0	0	0	0	0	0	00:00:00:00:00:00
1/6	0	0	0	0	0	0	0	00:00:00:00:00:00
1/7	0	0	0	0	0	0	0	00:00:00:00:00:00
1/8	0	0	0	0	0	0	0	00:00:00:00:00:00
1/9	0	0	0	0	0	0	0	00:00:00:00:00:00
1/10	0	0	0	0	0	0	0	00:00:00:00:00:00

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

**Variable Definitions**

Use the data in the following table to use the `show eapol auth-stats interface` command.

Variable	Value
{slot/port[/sub-port]}[-slot/port[/sub-port]][,...]	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

**Job Aid**

The following table describes the output for the `show eapol auth-stats interface` command.

**Table 38: show eapol auth-stats interface field descriptions**

Parameter	Description
PORT	Displays the port number in use.
EAP RCVD	Displays the number of EAPoL-EAP frames received by this Authenticator.
AUTH-EAP TX	Displays the number of EAPoL-EAP frames transmitted by the Authenticator.
START RCVD	Displays the number of EAPoL start frames received by this Authenticator.
LOGOFF RCVD	Displays the number of EAPoL logoff frames received by this Authenticator.
INVALID FRAMES	Displays the number of EAPoL frames received by this Authenticator in which the frame type is not recognized.
LENGTH ERROR	Displays the number of EAPoL frames received by this Authenticator in which the Packet Body Length field is invalid.
LAST-RX VER	Displays the last received version of the EAPoL frame by this Authenticator.
LAST-RX SRC	Displays the source MAC address of the last received EAPoL frame by this Authenticator.

## Viewing EAPoL Session Statistics

View EAPoL session statistics to manage network performance.

### \* Note:

Use the `clear-stats` command to clear EAP/NEAP statistics.

### Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display the session statistics:

```
show eapol session-stats interface [gigabitEthernet [{slot/port[/sub-port]}[-slot/port[/sub-port]][,...]]
```

### Example

```
Switch:1#show eapol session-stats interface
```

```
=====
                        Eap Authenticator Session Statistics
=====
PORT  MAC      SESSION  AUTHENTIC  SESSION  TERMINATE  USER
NUM   ID       ID       METHOD      TIME     CAUSE      NAME
-----
1/1   18:a9:05:b1:04:ce  cb000000  remote-server  0 day(s), 05:58:16  not-terminated
sachin
1/4   00:00:00:00:00:01  cb000002  remote-server  0 day(s), 05:48:01  not-terminated
```

```
000000000001
```

## Variable Definitions

Use the data in the following table to use the `show eapol session-stats interface` command.

Variable	Value
{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

## Job Aid

The following table describes the output for the `show eapol session-stats interface` command.

**Table 39: show eapol session-stats interface field descriptions**

Parameter	Description
PORT NUM	Displays the port number in use.
MAC	Displays the MAC address of the client.
USER NAME	Displays the user name of the Supplicant Authenticator Port Access Entity (PAE).
SESSION ID	Displays a unique identifier for the session.
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.

## Viewing Non-EAPoL MAC Information

Use this procedure to view non-EAPoL client MAC information on a port.

**\* Note:**

Use the `clear-stats` command to clear EAP/NEAP statistics.

### Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display the non-EAPoL MAC information:

```
show eapol multihost non-eap-mac status [vlan <1-4059>] [{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...]}
```



## Example

```
Switch:1#show eapol multihost non-eap-mac status
```

```
=====
                        Non-Eap Oper Status
=====
PORT  MAC                               STATE          VLAN
NUM  NUM                               ID
-----
1/3  00:00:00:11:22:33  RADIUS-Authenticated  250
=====
```

## Variable Definitions

Use the data in the following table to use the `show eapol multihost non-eap-mac status` command.

Variable	Value
{slot/port[/sub-port] [-slot/port[/sub-port]][,....]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.

## Job Aid

The following table describes the output for the `show eapol multihost non-eap-mac status` command.

**Table 40: show eapol multihost non-eap-mac status field descriptions**

Parameter	Description
PORT NUM	Displays the port number in use.
MAC	Displays the MAC address of the client.
STATE	Indicates the authentication status of the non EAP host that is authenticated using radius server.
VLAN ID	Indicates the VLAN assigned to the client.

## Viewing Port EAPoL Operation Statistics

Use this procedure to view port EAPoL operation statistics.

**\* Note:**

Use the `clear-stats` command to clear EAP/NEAP statistics.

**Procedure**

1. Log on to the switch to enter User EXEC mode.
2. Display the port EAPoL operation statistics information:

```
show eapol status interface [gigabitEthernet [{slot/port[/sub-port]
[-slot/port[/sub-port]][, ...]] [vlan <1-4059>]
```

**Example**

```
Switch:1#show eapol status interface
=====
                        Eap Oper Stats
=====
PORT  MAC                PAE          VLAN          PRI
NUM                                     STATUS        ID
-----
1/1   18:a9:05:b1:04:ce    authenticated  10            2
-----
Total Number of EAP sessions : 1
```

**Variable Definitions**

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

Variable	Value
{slot/port[/sub-port] [-slot/port[/sub-port]][, ...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
<1-4059>	Specifies the VLAN ID for which to show the statistics.

**Job Aid**

The following table describes the output for the `show eapol status interface` command.

**Table 41: show eapol status interface field descriptions**

Parameter	Description
PORT NUM	Displays the port number in use.
MAC	Displays the MAC address of the client.
PAE STATUS	Indicates the current state of the authenticator PAE state machine.
VLAN ID	Indicates the VLAN assigned to the client.
PRI	Indicates port priority.

## Viewing IP Multicast Threshold Exceeded Statistics

### Procedure

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

```
show sys stats ipmc-threshold-exceeded-cnt
```

### \* Note:

The command `show sys stats ipmc-threshold-exceeded-cnt` is not supported on all hardware platforms. For more information, see [Release Notes for VOSS](#).

### Example

```
Switch:1#show sys stats ipmc-threshold-exceeded-cnt
SourceGroupThresholdExceeded : 7372
EgressStreamThresholdExceeded : 7331
```

## Viewing NTP Statistics

### About this task

Use the `show ntp statistics` command to view the output for each NTP stratum regardless of whether authentication is enabled.

### Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View NTP statistics:

```
show ntp statistics
```

### Example

The output for the `show ntp statistics` command includes different information for NTPv3 and NTPv4.

For NTPv3:

```
Switch:1#show ntp statistics
NTP Server : 192.0.2.187
-----
Stratum : unknown
Version : unknown
Sync Status : not synchronized
Reachability : unreachable
Root Delay : unknown
Precision : unknown
Access Attempts : 2
Server Synch : 0
```

```
Server Fail : 2
Fail Reason : Server unreachable
```

**For NTPv4:**

```
Switch:1#show ntp statistics
```

```

      NTP Server : 192.0.2.187
-----
      Stratum : 16
      Version : NTPv4
      Broadcast : No
Auth Enabled : Disabled
Auth Status : Not-Auth
Sync Status : Rejected
Reachability : Unreachable
Root Delay : 0.000
Root Disp : 0.000
  Delay : 0.000
Dispersion : 15937.500
  Offset : 0.000
Precision : -23
  Jitter : 0.000
Last Event : Mobilize

      NTP Server : 192.0.2.201
-----
      Stratum : 4
      Version : NTPv4
      Broadcast : No
Auth Enabled : Enabled
Auth Status : Ok
Sync Status : Candidate
Reachability : Reachable
Root Delay : 18.448
Root Disp : 128.677
  Delay : 18.448
Dispersion : 0.366
  Offset : 0.202
Precision : -24
  Jitter : 1.041
Last Event : Popcorn
```

## Viewing Segmented Management Instance Statistics

View operational statistics for the Management Instance.

**Procedure**

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

```
show mgmt statistics [<clip | vlan>]
```
3. Enter Privileged EXEC mode:
 

```
enable
```

#### 4. (Optional) Clear the statistics for the Management Instance:

```
clear mgmt statistics [<clip | vlan>]
```

#### Example

```
Switch:1>show mgmt statistics
```

```
=====
                                Mgmt Interface Stats Information
=====
```

INST	DESCR	RX-PKTS	RX-ERROR	RX-DROP	TX-PKTS	TX-ERROR	TX-DROP
3	Mgmt-clip	124	0	22	100	0	12
4	Mgmt-vlan	765	0	1	434	0	0

```
-----
2 out of 2 Total Num of mgmt interfaces displayed
```

## Viewing Statistics Using EDM

Use statistics to help monitor the performance of the switch.

### About this task

To reset all statistics counters, click **Clear 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 Device Physical View tab, select the Device.
2. In the navigation pane, expand the **Configuration > Edit** folders.
3. Click **Chassis**.
4. Click the **System** tab.
5. 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 pane, expand the **Configuration > Graph** folders.

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
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 a graph of the 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. In the navigation pane, expand the **Configuration > Graph** folders, and then click **Port**.  
OR, use the following shortcut:  
Right-click the selected port or ports from Step 1, and choose **Graph**.
3. On the **Graph Port** tab for the selected port or ports, select the item you want to graph.
4. Click an icon to select the type of graph you require. The following list provides the graph types available:
  - Line Chart
  - Area Chart
  - Bar Chart
  - Pie Chart

## 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 pane, expand the **Configuration > Graph** folders.
3. Click **Chassis**.
4. Click the **System** tab.

## System Field Descriptions

The following table describes the fields on the System tab.

Name	Description
<b>MemUsed</b>	The percentage of memory 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>MemFree</b>	The amount in kilobytes of free memory.
<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 pane, expand the **Configuration > Graph** folders.
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.

*Table continues...*

Name	Description
<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>OutGetRequests</b>	The number of SNMP Get-Request PDUs that are generated by the SNMP protocol entity.
<b>InGetNexts</b>	The number of SNMP Get-Next PDUs the SNMP protocol accepts and processes.
<b>OutGetNexts</b>	The number of SNMP Get-Next PDUs that are generated by the SNMP protocol entity.
<b>InSetRequests</b>	The number of SNMP Set-Request PDUs the SNMP protocol accepts and processes.
<b>OutSetRequests</b>	The number of SNMP Set-Request PDUs that are generated by the SNMP protocol entity.
<b>InGetResponses</b>	The number of SNMP Get-Response PDUs the SNMP protocol accepts and processes.
<b>OutGetResponses</b>	The number of SNMP Get-Response PDUs that are generated by the SNMP protocol entity.
<b>InTraps</b>	The number of SNMP Trap PDUs the SNMP protocol accepts.
<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.
<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.

*Table continues...*



Name	Description
<b>InASNParseErrs</b>	The number of ASN.1 or BER errors the SNMP protocol encountered when decoding received SNMP messages.
<b>InTooBigs</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>InReadOnlys</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 pane, expand the **Configuration > Graph** folders.
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

*Table continues...*

Name	Description
	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). 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 pane, expand the **Configuration > Graph** folders.
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 pane, expand the **Configuration > Graph** folders.
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 Chassis TCP Statistics

View TCP statistics to monitor network performance.

### Procedure

1. In the Device Physical View, select the chassis.
2. In the navigation pane, expand the **Configuration > Graph** folders.
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.

*Table continues...*

Name	Description
<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 pane, expand the **Configuration > Graph** folders.
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>	<p>The number of received UDP datagrams with no application at the destination port.</p> <p>Discontinuities in the value of this counter can occur at reinitialization of the management system, and at other times as indicated by discontinuities in the value of sysUpTime.</p>
<b>InErrors</b>	<p>The number of received UDP datagrams that were not delivered for reasons other than the lack of an application at the destination port.</p> <p>Discontinuities in the value of this counter can occur at reinitialization of the management system and at other times as indicated by discontinuities in the value of sysUpTime.</p>
<b>InDatagrams</b>	<p>The number of UDP datagrams delivered to UDP users, for devices that can receive more than 1 000 000 UDP datagrams for each second.</p> <p>Discontinuities in the value of this counter can occur at reinitialization of the management system, and at other times as indicated by discontinuities in the value of sysUpTime.</p>
<b>OutDatagrams</b>	<p>The number of UDP datagrams sent from this entity.</p> <p>Discontinuities in the value of this counter can occur at reinitialization of the management system, and at other times as indicated by discontinuities in the value of sysUpTime.</p>
<b>HCInDatagrams</b>	<p>The number of TCP connections for which the current state is either ESTABLISHED or CLOSE-WAIT.</p>
<b>HCOutDatagrams</b>	<p>The number of UDP datagrams sent from this entity, for devices that can transmit more than 1 million UDP datagrams for each second.</p> <p>Discontinuities in the value of this counter can occur at reinitialization of the management system, and at other times as indicated by discontinuities in the value of sysUpTime.</p>

---

## 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 pane, expand the **Configuration > Graph** folders.
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>	Specifies the number of octets received on the interface, including framing characters.
<b>OutOctets</b>	Specifies the number of octets transmitted from the interface, including framing characters.
<b>InUcastPkts</b>	Specifies the number of packets delivered by this sublayer to a higher sublayer that were not addressed to a multicast or broadcast address at this sublayer.
<b>OutUcastPkts</b>	Specifies the number of packets that higher-level protocols requested be transmitted that were not addressed to a multicast address at this sublayer. The total number includes those packets discarded or not sent.
<b>InMulticastPkts</b>	Specifies the number of packets delivered by this sublayer to a higher sublayer that were addressed to a multicast address at this sublayer. For a MAC layer protocol, this number includes both group and functional addresses.
<b>OutMulticastPkts</b>	Specifies the number of packets that higher-level protocols requested be transmitted, and that are addressed to a multicast address at this sublayer, including those that were discarded or not sent. For a MAC layer protocol, this number includes both group and functional addresses.
<b>InBroadcastPkts</b>	Specifies 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>	Specifies the number of packets that higher-level protocols requested be transmitted, and that were addressed to a broadcast address at this sublayer, including those that were discarded or not sent.
<b>InDiscards</b>	Specifies the number of inbound packets that are discarded because of frames with errors or invalid frames or, in some cases, to fill up buffer space.
<b>InErrors</b>	For packet-oriented interfaces, specifies the number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol. For character-oriented or fixed-length interfaces, the number of inbound transmission units that contained errors preventing them from being deliverable to a higher-layer protocol.
<b>InUnknownProtos</b>	For packet-oriented interfaces, specifies the number of packets received through the interface that are discarded because of an unknown or unsupported protocol. For character-oriented or fixed-length interfaces that support protocol multiplexing, the number of transmission units received through the interface that were discarded because of an unknown or unsupported protocol. For any interface that does not support protocol multiplexing, this counter is always 0.
<b>HCInPfcPkts</b>	Specifies the total number of Priority Flow Control (PFC) packets received by this interface. This number does not increment for port-level flow control.

*Table continues...*

Name	Description
<b>HCOutPfcPkts</b>	Specifies the total number of PFC packets transmitted by this interface. This number does not increment for port-level flow control.
<b>InFlowCtrlPkts</b>	Specifies the number of port-level flow control packets received by this interface.
<b>OutFlowCtrlPkts</b>	Specifies the number of port-level flow control packets transmitted by this interface.
<b>InPfcPkts</b>	Specifies the total number of port-level flow control packets received by this interface.
<b>OutPfcPkts</b>	Specifies the total number of port-level flow control packets transmitted by this interface.
<b>NumStateTransition</b>	Specifies the number of times the port went in and out of service; the number of state transitions from up to down.

## Viewing Port Ethernet Errors Statistics

View port Ethernet errors statistics to manage network performance.

### Procedure

1. In the Device Physical View, select a port.
2. In the navigation pane, expand the **Configuration > Graph** folders.
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>	Specifies account of frames received on a particular interface that are not an integral number of octets in length and do not pass the FCS check. The count represented by an instance of this object increments when the alignmentError status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions obtain are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.
<b>FCSErrors</b>	Specifies a count of frames received on a particular interface that are an integral number of octets in length but do not pass the FCS check. The count represented by an instance of this object increments when the frameCheckError status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions obtained are, according to the conventions of IEEE 802.3

*Table continues...*



Name	Description
	Layer Management, counted exclusively according to the error status presented to the LLC.
<b>InternalMacTransmitErrors</b>	Specifies a count of frames for which transmission on a particular interface fails due to an internal MAC sublayer transmit error. A frame is only counted by an instance of this object if it is not counted by the corresponding instance of either the LateCollisions object, the ExcessiveCollisions object, or the CarrierSenseErrors object. The precise meaning of the count represented by an instance of this object is implementation-specific. In particular, an instance of this object can represent a count of transmission errors on a particular interface that are not otherwise counted.
<b>InternalMacReceiveErrors</b>	Specifies a count of frames for which reception on a particular interface fails due to an internal MAC sublayer receive error. A frame is only counted by an instance of this object if it is not counted by the corresponding instance of either the FrameTooLongs object, the AlignmentErrors object, or the FCSErrors object. The precise meaning of the count represented by an instance of this object is implementation-specific. In particular, an instance of this object can represent a count of receive errors on a particular interface that are not otherwise counted.
<b>CarrierSenseErrors</b>	Specifies the number of times that the carrier sense condition is lost or not asserted when the switch attempts to transmit a frame on a particular interface. The count represented by an instance of this object increments at most once for each transmission attempt, even if the carrier sense condition fluctuates during a transmission attempt.
<b>FrameTooLongs</b>	Specifies a count of frames received on a particular interface that exceed the maximum permitted frame size. The count represented by an instance of this object increments when the frameTooLong status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions obtained are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.
<b>SQETestErrors</b>	Specifies a count of times that the SQE TEST ERROR message is generated by the PLS sublayer for a particular interface. The SQE TEST ERROR message is defined in section 7.2.2.2.4 of ANSI/IEEE 802.3-1985 and its generation described in section 7.2.4.6 of the same document.
<b>DeferredTransmissions</b>	Specifies a count of frames for which the first transmission attempt on a particular interface is delayed because the medium is busy. The count represented by an instance of this object does not include frames involved in collisions.
<b>SingleCollisionFrames</b>	Specifies a count of successfully transmitted frames on a particular interface for which transmission is inhibited by exactly one collision. A frame that is counted by an instance of this object is also counted by

*Table continues...*

Name	Description
	the corresponding instance of either the UcastPkts, MulticastPkts, or BroadcastPkts objects and is not counted by the corresponding instance of the MultipleCollisionFrames object.
<b>MultipleCollisionFrames</b>	Specifies a count of successfully transmitted frames on a particular interface for which transmission is inhibited by more than one collision. A frame that is counted by an instance of this object is also counted by the corresponding instance of either the UcastPkts, MulticastPkts, or BroadcastPkts objects and is not counted by the corresponding instance of the SingleCollisionFrames object.
<b>LateCollisions</b>	Specifies the number of times that a collision is detected on a particular interface later than 512 bit-times into the transmission of a packet; 512 corresponds to 51.2 microseconds on a 10 Mb/s system. A (late) collision included in a count represented by an instance of this object is also considered as a (generic) collision for purposes of other collision-related statistics.
<b>ExcessiveCollisions</b>	Specifies a count of frames for which transmission on a particular interface fails due to excessive collisions.
<b>FrameTooShorts</b>	Specifies the number of frames, encountered on this interface, that are too short.
<b>LinkFailures</b>	Specifies the number of link failures encountered on this interface.
<b>PacketErrors</b>	Specifies the number of packet errors encountered on this interface.
<b>CarrierErrors</b>	Specifies the number of carrier errors encountered on this interface.
<b>LinkInactiveErrors</b>	Specifies the number of link inactive errors encountered on this interface.

## Viewing Port Bridging Statistics

View port bridging errors statistics to manage network performance.

 **Note:**

This tab is not available on all hardware platforms.

### Procedure

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

## Bridging Field Descriptions

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

Name	Description
<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 pane, expand the **Configuration > Graph** folders.
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.
<b>OutTcnBpdus</b>	The number of Topology Change Notifications BPDUs transmitted.

## Viewing Port Routing Statistics

View port routing statistics to manage network performance.

### \* Note:

This tab is not available on all hardware platforms.

### Procedure

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

4. Click the **Routing** tab.

## Routing Field Descriptions

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

Name	Description
<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 DHCP Statistics for an Interface

View DHCP statistics to manage network performance.

### Procedure

1. In the navigation pane, expand the **Configuration > IP** folders.
2. Click **DHCP Relay**.
3. Click the **Interfaces Stats** tab.

## Interfaces Stats Field Descriptions

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

Name	Description
<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.

---

## 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 pane, expand the **Configuration > Graph** folders.
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 pane, expand the **Configuration > Edit > Port** folders.
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 pane, expand the **Configuration > VLAN** folders.
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 pane, expand the **Configuration > IP** folders.
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 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. These reasons could be that the packet was received from an untrusted source or spoofing was detected. To

*Table continues...*

Name	Description
	determine the cause of the drop, you must enable trace on level 170.
<b>CircuitId</b>	Shows the value inserted in the packets as the circuit ID. The value is the index of the interface.
<b>AddedCircuitId</b>	Shows how many packets (requests from client to server) the circuit ID was inserted for that interface.  If you expect this value to increase but it does not, and the interface does not drop a packet, it is possible the packet does not have enough space to insert the option. You must enable trace on level 170 to determine the cause.
<b>RemovedCircuitId</b>	Shows how many packets (replies from server to client) the circuit id was removed for that interface.
<b>Remoteld</b>	Shows the value inserted in the packets as the remote ID. The value is the MAC address of the interface.
<b>AddedRemoteld</b>	Shows how many packets (requests from client to server) the remote ID was inserted for that interface.  If you expect this value to increase but it does not, and the interface does not drop a packet, it is possible the packet does not have enough space to insert the option. You must enable trace on level 170 to determine the cause.
<b>RemovedRemoteld</b>	Shows how many packets (replies from server to client) the remote ID was removed for that interface.

---

## Viewing Port OSPF Statistics

View port OSPF statistics to manage network performance.

### Procedure

1. On the Device Physical View, select a port.
2. In the navigation pane, expand the **Configuration > Graph** folders.
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>	Specifies the number of version mismatches received by this interface.
<b>AreaMismatches</b>	Specifies the number of area mismatches received by this interface.
<b>AuthTypeMismatches</b>	Specifies the number of authentication type mismatches received by this interface.
<b>AuthFailures</b>	Specifies the number of authentication failures.
<b>NetmaskMismatches</b>	Specifies the number of net mask mismatches received by this interface.
<b>HelloIntervalMismatches</b>	Specifies the number of hello interval mismatches received by this interface.
<b>DeadIntervalMismatches</b>	Specifies the number of dead interval mismatches received by this interface.
<b>OptionMismatches</b>	Specifies the number of option mismatches in the hello interval or dead interval fields received by this interface.
<b>RxHellos</b>	Specifies the number of hello packets received by this interface.
<b>RxDBDescrs</b>	Specifies the number of database descriptor packets received by this interface.
<b>RxLSUpdates</b>	Specifies the number of link state update packets received by this interface.
<b>RxLSReqs</b>	Specifies the number of link state request packets received by this interface.
<b>RxLSAcks</b>	Specifies the number of link state acknowledge packets received by this interface.
<b>TxHellos</b>	Specifies the number of hello packets transmitted by this interface.
<b>TxDBDescrs</b>	Specifies the number of database descriptor packets transmitted by this interface.
<b>TxLSUpdates</b>	Specifies the number of link state update packets transmitted by this interface.
<b>TxLSReqs</b>	Specifies the number of link state request packets transmitted by this interface.
<b>TxLSAcks</b>	Specifies the number of link state acknowledge packets transmitted by this interface.

---

## Viewing LACP Port Statistics

View LACP port statistics to monitor the performance of the port.

### Procedure

1. In the Device Physical View, select a port.



2. In the navigation pane, expand the **Configuration > Graph** folders.
3. Click **Port**.
4. Click the **LACP** tab.
5. To change the poll interval, in the toolbar click the **Poll Interval** box, and then select a new interval.

## LACP Field Descriptions

Use the data in the following table to view the LACP statistics.

Name	Description
<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 LACPDUs 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.

This tab does not appear for all hardware models.

### Procedure

1. In the navigation pane, expand the **Configuration > Graph** folders.
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 onboard flash memory, as well as the number of files.

### Procedure

1. In the navigation pane, expand the **Configuration > Edit** folders.
2. Click **File System**.
3. Click the **Storage Usage** tab.

## Storage Usage Field Descriptions

Use the data in the following table to use the Storage Usage tab.

Name	Description
<b>IntflashBytesUsed</b>	Specifies the number of bytes used in internal flash memory.
<b>IntflashBytesFree</b>	Specifies the number of bytes available for use in internal flash memory.
<b>IntflashNumFiles</b>	Specifies the number of files in internal flash memory.
<b>UsbBytesUsed</b>	Specifies the number of bytes used in USB device.
<b>UsbBytesFree</b>	Specifies the number of bytes available for use in USB device.
<b>UsbNumFiles</b>	Specifies the number of files in USB device.

## Viewing ACE Port Statistics

### About this task

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

### Procedure

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

2. Click **Advanced Filters (ACE/ACLs)**.
3. Click the **ACL** tab.
4. 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 pane, expand the **Configuration > Security > Data Path** folders.
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.
<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.

*Table continues...*

Name	Description
<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 pane, expand the **Configuration > Security > Data Path** folders.
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 pane, expand the **Configuration > VLAN > Spanning Tree** folders.
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>	Specifies the number of times this port has transitioned to the forwarding state.
<b>RxMstBpduCount</b>	Specifies the number of MSTP BPDUs received on this port.
<b>RxRstBpduCount</b>	Specifies the number of RSTP BPDUs received on this port.
<b>RxConfigBpduCount</b>	Specifies the number of configuration BPDUs received on this port.
<b>RxTcnBpduCount</b>	Specifies the number of TCN BPDUs received on this port.
<b>TxMstBpduCount</b>	Specifies the number of MSTP BPDUs transmitted from this port.
<b>TxRstBpduCount</b>	Specifies the number of RSTP BPDUs transmitted from this port.
<b>TxConfigBpduCount</b>	Specifies the number of configuration BPDUs transmitted from this port.
<b>TxTcnBpduCount</b>	Specifies the number of TCN BPDUs transmitted from this port.
<b>InvalidMstBpduRxCount</b>	Specifies the number of Invalid MSTP BPDUs received on this port.
<b>InvalidRstBpduRxCount</b>	Specifies the number of Invalid RSTP BPDUs received on this port.
<b>InvalidConfigBpduRxCount</b>	Specifies the number of invalid configuration BPDUs received on this port.
<b>InvalidTcnBpduRxCount</b>	Specifies the number of invalid TCN BPDUs received on this port. The number of times this port has migrated from one STP protocol version to another. The relevant protocols are STP-compatible and RSTP/MSTP. A trap is generated on the occurrence of this event.
<b>ProtocolMigrationCount</b>	Specifies the number of times this port has migrated from one STP protocol version to another. The relevant protocols are STP-compatible and RSTP. A trap is generated on the occurrence of this event.

---

## Viewing VLAN and Spanning Tree MSTI Statistics

### About this task

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

**Procedure**

1. In the navigation pane, expand the **Configuration > VLAN > Spanning Tree** folders.
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>	Specifies the number of times this port has transitioned to the forwarding state for this specific instance.
<b>ReceivedBPDUs</b>	Specifies the number of BPDUs received by this port for this spanning tree instance.
<b>TransmittedBPDUs</b>	Specifies the number of BPDUs transmitted on this port for this spanning tree instance.
<b>InvalidBPDUsRcvd</b>	Specifies the number of invalid BPDUs received on this port for this spanning tree instance.

---

**Viewing VRRP Interface Statistics**

**About this task**

View VRRP statistics to manage network performance.

**Procedure**

1. In the navigation pane, expand the **Configuration > IP** folders.
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>AdvertiseRcvd</b>	Specifies the number of VRRP advertisements received by this virtual router.

*Table continues...*

Name	Description
<b>AdvertiseIntervalErrors</b>	Specifies the number of received VRRP advertisement packets with a different interval is than configured for the local virtual router.
<b>IPtTtlErrors</b>	Specifies the number of VRRP packets received by the virtual router with IP TTL (Time-To-Live) not equal to 255.
<b>PriorityZeroPktsRcvd</b>	Specifies the number of VRRP packets received by the virtual router with a priority of 0.
<b>PriorityZeroPktsSent</b>	Specifies the number of VRRP packets sent by the virtual router with a priority of 0'.
<b>InvalidTypePktsRcvd</b>	Specifies the number of VRRP packets received by the virtual router with an invalid value in the 'type' field.
<b>AddressListErrors</b>	Specifies the packets received address list the address list does not match the locally configured list for the virtual router.
<b>AuthTypeMismatch</b>	Specifies the count of authentication type mismatch messages.
<b>PacketLengthErrors</b>	Specifies the count of packet length errors.
<b>AuthFailures</b>	Specifies the count of authentication failure messages.

## Viewing VRRP Statistics

### About this task

View VRRP statistics to monitor network performance.

### Procedure

1. In the navigation pane, expand the **Configuration > IP** folders.
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>	Specifies the number of VRRP packets received with an invalid VRRP checksum value.
<b>VersionErrors</b>	Specifies the number of VRRP packets received with an unknown or unsupported version number.

*Table continues...*

Name	Description
VrIDErrors	Specifies the number of VRRP packets received with an invalid VrID for this virtual router.

## Viewing SMLT Statistics

View SMLT statistics to manage network performance.

### Procedure

1. In the navigation pane, expand the **Configuration > VLAN** folders.
2. Click **MLT/LACP**.
3. Select the **Ist/SMLT Stats** tab.

## IST/SMLT Stats Field Descriptions

The following table describes parameters on the IST/SMLT Stats tab.

Name	Description
SmltIstDownCnt	The number of times the session between the two peering switches has gone down since last boot.
SmltHelloTxMsgCnt	The count of transmitted hello messages.
SmltHelloRxMsgCnt	The count of received hello messages.
SmltLearnMacAddrTxMsgCnt	The count of transmitted learned MAC address messages.
SmltLearnMacAddrRxMsgCnt	The count of received learned MAC address messages.
SmltMacAddrAgeOutTxMsgCnt	The count of transmitted aging out MAC address messages.
SmltMacAddrAgeOutRxMsgCnt	The count of received aging out MAC address messages.
SmltMacAddrAgeExpTxMsgCnt	The count of transmitted MAC address age expired messages.
SmltMacAddrAgeExpRxMsgCnt	The count of received MAC address age expired messages.
SmltStgInfoTxMsgCnt	The count of transmitted STG information messages.
SmltStgInfoRxMsgCnt	The count of received STG information messages.
SmltDelMacAddrTxMsgCnt	The count of transmitted MAC address deleted messages.
SmltDelMacAddrRxMsgCnt	The count of received MAC address received messages.
SmltSmltDownTxMsgCnt	The count of transmitted SMLT down messages.

*Table continues...*



Name	Description
SmltSmltDownRxMsgCnt	The count of received SMLT down messages
SmltUpTxMsgCnt	The count of transmitted SMLT up messages.
SmltUpRxMsgCnt	The count of received SMLT up messages.
SmltSendMacTbITxMsgCnt	The count of sent send MAC table messages.
SmltSendMacTbIRxMsgCnt	The count of received send MAC table messages.
SmltIcmpTxMsgCnt	The count of sent IGMP messages.
SmltIcmpRxMsgCnt	The count of received IGMP messages.
SmltPortDownTxMsgCnt	The count of sent port down messages.
SmltPortDownRxMsgCnt	The count of received port down messages.
SmltReqMacTbITxMsgCnt	The count or sent MAC table request messages.
SmltReqMacTbIRxMsgCnt	The count of received MAC table request messages.
SmltRxUnknownMsgTypeCnt	The count of received unknown message type messages.
SmltPortTbISyncReqTxMsgCnt	The count of sent sync request messages.
SmltPortTbISyncReqRxMsgCnt	The count of received sync request messages.
SmltPortTbISyncTxMsgCnt	The count of sent sync messages.
SmltPortTbISyncRxMsgCnt	The count of received sync messages.
SmltPortUpdateTxMsgCnt	The count of sent update messages.
SmltPortUpdateRxMsgCnt	The count of received update messages.
SmltEntryUpdateTxMsgCnt	The count of sent entry update messages.
SmltEntryUpdateRxMsgCnt	The count of received entry update messages.
SmltDialectNegotiateTxMsgCnt	The count of sent protocol ID messages.
SmltDialectNegotiateRxMsgCnt	The count of received protocol ID messages.
SmltUpdateRespTxMsgCnt	The count of sent update response messages.
SmltUpdateRespRxMsgCnt	The count of received update response messages.
SmltTransQHighWaterMarkMsgCnt	The count of transaction queue high watermark messages.
SmltPollCountHighWaterMarkCnt	The count of poll count high watermark.

## Viewing RSTP Status Statistics

### About this task

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

### Procedure

1. In the navigation pane, expand the **Configuration > VLAN > Spanning Tree** folders.
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>	Specifies the number of RSTP BPDUs this port received.
<b>RxConfigBpduCount</b>	Specifies the number of configuration BPDUs this port received.
<b>RxTcnBpduCount</b>	Specifies the number of TCN BPDUs this port received.
<b>TxRstBpduCount</b>	Specifies the number of RSTP BPDUs this port transmitted.
<b>TxConfigBpduCount</b>	Specifies the number of Config BPDUs this port transmitted.
<b>TxTcnBpduCount</b>	Specifies the number of TCN BPDUs this port transmitted.
<b>InvalidRstBpduRxCount</b>	Specifies the number of invalid RSTP BPDUs this port received. A trap is generated on the occurrence of this event.
<b>InvalidConfigBpduRx Count</b>	Specifies the number of invalid configuration BPDUs this port received. A trap is generated on the occurrence of this event.
<b>InvalidTcnBpduRxCount</b>	Specifies the number of invalid TCN BPDUs this port received. A trap is generated on the occurrence of this event.
<b>ProtocolMigrationCount</b>	Specifies 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 pane, expand the **Configuration > VLAN** folders.
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>	Specifies the total number of octets received on the MLT interface, including framing characters.
<b>OutOctets</b>	Specifies the total number of octets transmitted out of the MLT interface, including framing characters.
<b>InUcastPkts</b>	Specifies the number of packets delivered by this MLT to higher level protocols that were not addressed to a multicast or broadcast address at this sublayer.
<b>OutUcastPkts</b>	Specifies the number of packets that higher-level protocols requested be transmitted that were not addressed to a multicast address at this MLT. This total number includes discarded or unsent packets.
<b>InMulticastPkt</b>	Specifies the number of packets delivered to this MLT that were addressed to a multicast address at this sublayer. For a MAC layer protocol, this number includes both Group and Functional addresses.
<b>OutMulticast</b>	Specifies the total number of packets that higher-level protocols requested be transmitted, and that were addressed to a multicast address at this MLT, including those that were discarded or unsent. For a MAC layer protocol, this number includes both Group and Functional addresses.
<b>InBroadcastPkt</b>	Specifies the number of packets delivered to this MLT that were addressed to a broadcast address at this sublayer.
<b>OutBroadcast</b>	Specifies the total number of packets that higher-level protocols requested be transmitted, and that were addressed to a broadcast address at this MLT, including those that were discarded or not sent.
<b>InLsmPkts</b>	Specifies the total number of Link State Messaging (LSM) packets delivered on this MLT.
<b>OutLsmPkts</b>	Specifies the total number of Link State Messaging (LSM) packets transmitted on this MLT.

---

## Viewing MLT Ethernet Error Statistics

### About this task

Use MLT Ethernet error statistics to view the error statistics.

### Procedure

1. In the navigation pane, expand the **Configuration > VLAN** folders.
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>	Specifies the frame count frames received on a particular MLT that is not an integral number of octets in length and does not pass the FCS check. The count represented by an instance of this object increments when the alignmentError status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions occur are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.
<b>FCSErrors</b>	Specifies the frame count received on an MLT that is an integral number of octets in length, but does not pass the Frame Check Sequence (FCS) check. The count represented by an instance of this object increments when the FrameCheckError status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions occur are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.
<b>IMacTransmitError</b>	Specifies the frame count for which transmission on a particular MLT fails due to an internal MAC sublayer transmit error. A frame is only counted by an instance of this object if it is not counted by the corresponding instance of either the LateCollisions object, the ExcessiveCollisions object, or the CarrierSenseErrors object.
<b>IMacReceiveError</b>	<p>Specifies the frame count for which reception on a particular MLT fails due to an internal MAC sublayer receive error. A frame is only counted by an instance of this object if it is not counted by the corresponding instance of either the FrameTooLongs object, the AlignmentErrors object, or the FCSErrors object.</p> <p>The precise meaning of the count represented by an instance of this object is implementation specific. In particular, an instance of this object can represent receive errors on a particular interface that are not otherwise counted.</p>
<b>CarrierSenseError</b>	Specifies the number of times that the carrier sense condition was lost or never asserted when attempting to transmit a frame on a particular MLT. The count represented by an instance of this object increments at most once for each transmission attempt, even if the carrier sense condition fluctuates during a transmission attempt.
<b>FrameTooLong</b>	Specifies the frame count received on a particular MLT that exceeds the maximum permitted frame size. The count represented by an instance of this object increments when the frameTooLong status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions occur are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.
<b>SQETestError</b>	Specifies the number of times that the SQE test error message is generated by the PLS sublayer for a particular MLT. The SQE test error message is defined in section 7.2.2.2.4 of ANSI/ IEEE 802.3-1985.

*Table continues...*

Name	Description
<b>DeferredTransmiss</b>	Specifies the frame count for which the first transmission attempt on a particular MLT is delayed because the medium is busy. The count represented by an instance of this object does not include frames involved in collisions.
<b>SingleCollFrames</b>	Specifies a count of successfully transmitted frames on a particular MLT for which transmission is inhibited by exactly one collision. A frame that is counted by an instance of this object is also counted by the corresponding instance of either the ifOutUcastPkts object, the ifOutMulticastPkts object, or the ifOutBroadcastPkts object, and is not counted by the corresponding instance of the MultipleCollisionFrames object.
<b>MultipleCollFrames</b>	Specifies the successfully transmitted frame count on a particular MLT for which transmission is inhibited by more than one collision. A frame that is counted by an instance of this object is also counted by the corresponding instance of either the ifOutUcastPkts object, the ifOutMulticastPkts object, or the ifOutBroadcastPkts object, and is not counted by the corresponding instance of the SingleCollisionFrames object.
<b>LateCollisions</b>	Specifies the number of times that a collision is detected on a particular MLT later than 512 bit-times into the transmission of a packet; 512 corresponds to 51.2 microseconds on a 10 Mb/s system. A late collision included in a count represented by an instance of this object is also considered as a generic collision for purposes of other collision-related statistics.
<b>ExcessiveCollis</b>	Specifies the frame count for which transmission on a particular MLT fails due to excessive collisions.

---

## Viewing RIP Statistics

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

### Procedure

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

*Table continues...*

Name	Description
<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

Use statistics to help you monitor Open Shortest Path First (OSPF) performance. You can also graph statistics for all OSPF packets transmitted by the switch.

### Procedure

1. In the navigation pane, expand the **Configuration > IP** folders.
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>	Specifies the number of entries in the link state database table.
<b>TxPackets</b>	Specifies the number of packets transmitted by OSPF.
<b>RxPackets</b>	Specifies the number of packets received by OSPF.
<b>TxDropPackets</b>	Specifies the number of packets dropped before being transmitted by OSPF.
<b>RxDropPackets</b>	Specifies the number of packets dropped before they are received by OSPF.
<b>RxBadPackets</b>	Specifies the number of packets received by OSPF that are bad.
<b>SpfRuns</b>	Specifies the number of SPF calculations performed by OSPF.
<b>BuffersAllocated</b>	Specifies the number of buffers allocated for OSPF.
<b>BuffersFreed</b>	Specifies the number of buffers freed by OSPF.
<b>BufferAllocFailures</b>	Specifies the number of times that OSPF has failed to allocate buffers.
<b>BufferFreeFailures</b>	Specifies the number of times that OSPF has failed to free buffers.
<b>Routes</b>	Specifies the count of OSPF routes.
<b>Adjacencies</b>	Specifies the count of OSPF adjacencies.
<b>Areas</b>	Specifies the count of OSPF areas.

## Graphing OSPF Statistics for a VLAN

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

### Procedure

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

*Table continues...*

Name	Description
<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.

## Graphing OSPF Statistics for a Port

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

### Procedure

1. On the Device Physical View, select a port.
2. In the navigation pane, expand the **Configuration > Edit > Port** folders.
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.

*Table continues...*



Name	Description
<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.
<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 BGP Global Stats

View BGP global stats.

### Procedure

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

2. Click **BGP**.
3. Click the **Global Stats** tab.

## Global Stats Field Descriptions

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

Name	Description
<b>AbsoluteValue</b>	Displays the counter value.
<b>Cumulative</b>	Displays the total value since you opened the Stats tab.
<b>Average/sec</b>	Displays the average value for each second.
<b>Minimum/sec</b>	Displays the minimum value for each second.
<b>Maximum/sec</b>	Displays the maximum value for each second.
<b>LastVal/sec</b>	Displays the last value for each second.
<b>Starts</b>	Displays the number of times the BGP connection started.
<b>Stops</b>	Displays the number of times the BGP connection stopped.
<b>Opens</b>	Displays the number of times BGP opens TCP.
<b>Closes</b>	Displays the number of times BGP closes TCP.
<b>Fails</b>	Displays the number of times TCP attempts failed.
<b>Fatals</b>	Displays the number of times TCP crashes due to fatal error.
<b>ConnExps</b>	Displays the number of times the TCP retry timer expired.
<b>HoldExps</b>	Displays the number of times the hold timer expired.
<b>KeepExps</b>	Displays the number of times the keepalive timer expired.
<b>RxOpens</b>	Displays the number of open instances BGP receives.
<b>RxKeeps</b>	Displays the number of keepalive instances BGP receives.
<b>RxUpdates</b>	Displays the number of update instances BGP receives.
<b>RxNotifys</b>	Displays the number of notification instances BGP receives.
<b>TxOpens</b>	Displays the number of open instances BGP transmitted.
<b>TxKeeps</b>	Displays the number of keepalive instances BGP transmitted.

*Table continues...*

Name	Description
<b>TxUpdates</b>	Displays the number of updates instances BGP transmits.
<b>TxNotifys</b>	Displays the number of notification instances BGP transmits.
<b>BadEvents</b>	Displays the number of invalid events FSM received.
<b>SyncFails</b>	Displays the number of times FDB sync failed.
<b>TrEvent</b>	Displays the trace event.
<b>RxECodeHeader</b>	Displays the total header errors received.
<b>RxECodeOpen</b>	Displays the total open errors received.
<b>RxECodeUpdate</b>	Displays the total update errors received.
<b>RxECodeHoldtimer</b>	Displays the total hold timer errors received.
<b>RxECodeFSM</b>	Displays the total FSM errors received.
<b>RxECodeCease</b>	Displays the total cease errors received.
<b>RxHdrCodeNoSync</b>	Displays the header not synchronized errors received.
<b>RxHdrCodeInvalidMsgLen</b>	Displays the header invalid message length errors received.
<b>RxHdrCodeInvalidMsgType</b>	Displays the header invalid message type errors received.
<b>RxOpCodeBadVer</b>	Displays the open errors received for Bad Version.
<b>RxOpCodeBadAs</b>	Displays the open errors received for le Bad AS Number.
<b>RxOpCodeBadRtID</b>	Displays the open errors received for Bad BGP Rtr ID.
<b>RxOpCodeUnsuppOption</b>	Displays the open errors received for Unsupported Option.
<b>RxOpCodeAuthFail</b>	Displays the open errors received for Auth Failures.
<b>RxOpCodeBadHold</b>	Displays the open errors received for Bad Hold Value.
<b>RxUpdCodeMalformedAttrList</b>	Displays the update errors received for Malformed Attr List.
<b>RxUpdCodeWelKnownAttrUnrecog</b>	Displays the update errors received for Welknown Attr Unrecog.
<b>RxUpdCodeWelknownAttrMiss</b>	Displays the update errors received for Welknown Attr Missing.
<b>RxUpdCodeAttrFlagError</b>	Displays the update errors received for Attr Flag Error.
<b>RxUpdCodeAttrLenError</b>	Displays the update errors received for Attr Len Error.

*Table continues...*

Name	Description
<b>RxUpdCodeBadORIGINAttr</b>	Displays the update errors received for Bad ORIGIN Attr.
<b>RxUpdCodeASRoutingLoop</b>	Displays the update errors received for AS Routing Loop.
<b>RxUpdCodeBadNHAttr</b>	Displays the update errors received for Bad NEXT-HOP Attr.
<b>RxUpdCodeOptionalAttrError</b>	Displays the update errors received for Optional Attr Error.
<b>RxUpdCodeBadNetworkField</b>	Displays the update errors received for Bad Network Field.
<b>RxUpdCodeMalformedASPath</b>	Displays the update errors received for Malformed AS Path.
<b>TxECodeHeader</b>	Displays the total Header errors transmitted.
<b>TxECodeOpen</b>	Displays the total Open errors transmitted.
<b>TxECodeUpdate</b>	Displays the total Update errors transmitted.
<b>TxECodeHoldtimer</b>	Displays the total Hold timer errors transmitted.
<b>TxECodeFSM</b>	Displays the total FSM errors transmitted.
<b>TxECodeCease</b>	Displays the total Cease errors transmitted.
<b>TxHdrCodeNoSync</b>	Displays the header Not Synchronized errors transmitted.
<b>TxHdrCodeInvalidMsgLen</b>	Displays the header Invalid msg len errors transmitted.
<b>TxHdrCodeInvalidMsgType</b>	Displays the header Invalid msg type errors transmitted.
<b>TxOpCodeBadVer</b>	Displays the open errors transmitted for Bad Version.
<b>TxOpCodeBadAs</b>	Displays the open errors transmitted for Bad AS Number.
<b>TxOpCodeBadRtrID</b>	Displays the open errors transmitted for Bad BGP Rtr ID.
<b>TxOpCodeUnsuppOption</b>	Displays the open errors transmitted for Unsupported Option.
<b>TxOpCodeAuthFail</b>	Displays the open errors transmitted for Auth Failures.
<b>TxOpCodeBadHold</b>	Displays the open errors transmitted for Bad Hold Value.
<b>TxUpdCodeMalformedAttrList</b>	Displays the update errors transmitted for Malformed Attr List.
<b>TxUpdCodeWelknownAttrUnrecog</b>	Displays the update errors transmitted for Welknown Attr Unrecog.

*Table continues...*

Name	Description
<b>TxUpdCodeWelknownAttrMiss</b>	Displays the update errors transmitted for Welknown Attr Missing.
<b>TxUpdCodeAttrFlagError</b>	Displays the update errors transmitted for Attr Flag Error.
<b>TxUpdCodeAttrLenError</b>	Displays the update errors transmitted for Attr Len Error.
<b>TxUpdCodeBadORIGINAttr</b>	Displays the update errors transmitted for Bad ORIGIN Attr.
<b>TxUpdCodeASRoutingLoop</b>	Displays the update errors transmitted for AS Routing Loop
<b>TxUpdCodeBadNHAttr</b>	Displays the update errors transmitted for Bad NEXT-HOP Attr
<b>TxUpdCodeOptionalAttrError</b>	Displays the update errors transmitted for Optional Attr Error.
<b>TxUpdCodeBadNetworkField</b>	Displays the update errors transmitted for Bad Network Field.
<b>TxUpdCodeMalformedASPath</b>	Displays the update errors transmitted for 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 pane, expand the **Configuration > IP** folders.
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>StatRouteEntries</b>	Specifies the number of routes for this VRF.
<b>StatFIBEntries</b>	Specifies the number of Forwarding Information Base (FIB) entries for this VRF.

## 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 pane, expand the **Configuration > Security > Control Path** folders.
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 the type of IP address. RADIUS supports IPv4 addresses only.
<b>Address</b>	Shows the IP address of the RADIUS server.
<b>Used by</b>	Identifies the client.
<b>AccessRequests</b>	Shows the number of access-response packets sent to the server; does not include retransmissions.
<b>AccessAccepts</b>	Shows the number of access-accept packets, valid or invalid, received from the server.
<b>AccessRejects</b>	Shows the number of access-reject packets, valid or invalid, received from the server.
<b>BadResponses</b>	Shows the number of invalid access-response packets received from the server.
<b>PendingRequests</b>	Shows the access-request packets sent to the server that have not yet received a response or that have timed out.
<b>ClientRetries</b>	Shows the number of authentication retransmissions to the server.
<b>AcctOnRequests</b>	Shows the number of accounting on requests sent to the server.
<b>AcctOffRequests</b>	Shows the number of accounting off requests sent to the server.
<b>AcctStartRequests</b>	Shows the number of accounting start requests sent to the server.
<b>AcctStopRequests</b>	Shows the number of accounting stop requests sent to the server.
<b>AcctInterimRequests</b>	<p>Number of Accounting Interim requests sent to the server.</p> <p> <b>Important:</b></p> <p>The AcctInterimRequests counter increments only if you select AcctIncludeCli from the RADIUS Global tab.</p>
<b>AcctBadResponses</b>	Shows the number of Invalid responses discarded from the server.

*Table continues...*

Name	Description
<b>AcctPendingRequests</b>	Shows the number of requests waiting to be sent to the server.
<b>AcctClientRetries</b>	Shows the number of retries made to this server.
<b>RoundTripTime</b>	Shows the time difference between the instance when a RADIUS request is sent and the corresponding response is received.
<b>AccessChallenges</b>	Shows the number of RADIUS access-challenges packets sent to this server. This does not include retransmission.
<b>NasIpAddress</b>	Shows the RADIUS client NAS Identifier for 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 pane, expand the **Configuration > Security > Control Path** folders.
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>OutTooBig</b>	Shows 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>	Shows 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>	Shows 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>	Shows 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>	Shows the number of SNMP messages that were delivered to the SNMP protocol entity and were for an unsupported SNMP version.
<b>InBadCommunityNames</b>	Shows the number of SNMP messages delivered to the SNMP protocol entity that used an SNMP community name not known to the entity.

*Table continues...*

Name	Description
<b>InBadCommunityUsers</b>	Shows the number of SNMP messages delivered to the SNMP protocol entity that represented an SNMP operation not allowed by the SNMP community named in the message.
<b>InASNParseErrs</b>	Shows the number of ASN.1 or BER errors encountered by the SNMP protocol entity when decoding received SNMP messages.
<b>InTooBigs</b>	Shows the number of SNMP PDUs that were delivered to the SNMP protocol entity and for which the value of the error-status field is tooBig.
<b>InNoSuchNames</b>	Shows the number of SNMP PDUs that were delivered to the SNMP protocol entity and for which the value of the error-status field is noSuchName.
<b>InBadValues</b>	Shows the number of SNMP PDUs that were delivered to the SNMP protocol entity and for which the value of the error-status field is badValue.
<b>InReadOnlys</b>	Shows the number of valid SNMP PDUs delivered to the SNMP protocol entity and for which the value of the error-status field is "read-only". It is a protocol error to generate an SNMP PDU that contains the value "read-only" in the error-status field; this object is provided as a means of detecting incorrect implementations of the SNMP.
<b>InGenErrors</b>	Shows the number of SNMP PDUs delivered to the SNMP protocol entity and for which the value of the error-status field is "genErr."

---

## Enabling RMON Statistics

### About this task

Enable Ethernet statistics collection for RMON.

### Procedure

1. In the navigation pane, expand the **Configuration > Serviceability > RMON** folders.
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 pane, expand the **Configuration > Graph** folders.
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>	Specifies the number of octets of data (including those in bad packets) received on the network (excluding framing bits but including FCS octets).  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

*Table continues...*

Name	Description
	<p>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>	Specifies the number of packets (including bad packets, broadcast packets, and multicast packets) received.
<b>BroadcastPkts</b>	Specifies the number of good packets received that were directed to the broadcast address. This number does not include multicast packets.
<b>MulticastPkts</b>	Specifies 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>CRCAlignErrors</b>	Specifies 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>	Specifies 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>	Specifies 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>	<p>Specifies 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).</p> <p>It is entirely normal for Fragments to increment because it counts both runs (which are normal occurrences due to collisions) and noise hits.</p>
<b>Collisions</b>	<p>Specifies 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</p>

*Table continues...*

Name	Description
	<p>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>

## Displaying IS-IS System Statistics

Use the following procedure to display Intermediate-System-to-Intermediate-System (IS-IS) system statistics.

### Procedure

1. In the navigation pane, expand the **Configuration > IS-IS** folders.
2. Click **Stats**.
3. Click the **System Stats** tab.

## System Stats Field Descriptions

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

Name	Description
<b>CorrLSPs</b>	Indicates the number of corrupted in-memory link-state packets (LSPs) detected. LSPs received from the wire with a bad checksum are silently dropped and not counted.
<b>AuthFails</b>	Indicates the number of authentication key failures recognized by this Intermediate System.
<b>LSPDbaseOloads</b>	Indicates the number of times the LSP database has become overloaded.
<b>ManAddrDropFromAreas</b>	Indicates the number of times a manual address has been dropped from the area.
<b>AttmptToExMaxSeqNums</b>	Indicates the number of times the IS has attempted to exceed the maximum sequence number.
<b>SeqNumSkips</b>	Indicates the number of times a sequence number skip has occurred.
<b>OwnLSPPurges</b>	Indicates the number of times a zero-aged copy of the system's own LSP is received from some other node.
<b>IDFieldLenMismatches</b>	Indicates the number of times a PDU is received with a different value for ID field length to that of the receiving system.
<b>PartChanges</b>	Indicates partition changes.

*Table continues...*

Name	Description
<b>AbsoluteValue</b>	Displays the counter value.
<b>Cumulative</b>	Displays the total value since you opened the Stats tab.
<b>Average/sec</b>	Displays the average value for each second.
<b>Minimum/sec</b>	Displays the minimum value for each second.
<b>Maximum/sec</b>	Displays the maximum value for each second.
<b>LastVal/sec</b>	Displays the last value for each second.

## Displaying IS-IS Interface Counters

Use the following procedure to display IS-IS interface counters.

### Procedure

1. In the navigation pane, expand the **Configuration > IS-IS** folders.
2. Click **Stats**.
3. Click the **Interface Counters** tab.

## Interface Counters Field Descriptions

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

Name	Description
<b>Index</b>	Shows a unique value identifying the IS-IS interface.
<b>AdjChanges</b>	Shows the number of times an adjacency state change has occurred on this circuit.
<b>InitFails</b>	Shows the number of times initialization of this circuit has failed. This counts events such as PPP NCP failures. Failures to form an adjacency are counted by isisCircRejAdjs.
<b>RejAdjs</b>	Shows the number of times an adjacency has been rejected on this circuit.
<b>IDFieldLenMismatches</b>	Shows the number of times an IS-IS control PDU with an ID field length different to that for this system has been received.
<b>MaxAreaAddrMismatches</b>	Shows the number of times an IS-IS control PDU with a max area address field different to that for this system has been received.
<b>AuthFails</b>	Shows the number of times an IS-IS control PDU with the correct auth type has failed to pass authentication validation.
<b>LANDesiChanges</b>	Shows the number of times the Designated IS has changed on this circuit at this level. If the circuit is point to point, this count is zero.

## Displaying IS-IS Interface Control Packets

Use the following procedure to display IS-IS interface control packets.

### Procedure

1. In the navigation pane, expand the **Configuration > IS-IS** folders.
2. Click **Stats**.
3. Click the **Interface Control Packets** tab.

## Interface Control Packets Field Descriptions

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

Name	Description
<b>Index</b>	Shows a unique value identifying the Intermediate-System-to-Intermediate-System (IS-IS) interface.
<b>Direction</b>	Indicates whether the switch is sending or receiving the PDUs.
<b>Hello</b>	Indicates the number of IS-IS Hello frames seen in this direction at this level.
<b>LSP</b>	Indicates the number of IS-IS LSP frames seen in this direction at this level.
<b>CSNP</b>	Indicates the number of IS-IS Complete Sequence Number Packets (CSNP) frames seen in this direction at this level.
<b>PSNP</b>	Indicates the number of IS-IS Partial Sequence Number Packets (PSNP) frames seen in this direction at this level.

## Graphing IS-IS Interface Counters

Use the following procedure to graph IS-IS interface counters.

### Procedure

1. In the navigation pane, expand the **Configuration > IS-IS** folders.
2. Click **IS-IS**.
3. Click the **Interfaces** tab.
4. Select an existing interface.
5. Click the **Graph** button.

## Interface Counters Field Descriptions

The following table describes the fields in the **Interface Counters** tab.

Name	Description
<b>InitFails</b>	Indicates the number of times initialization of this circuit has failed. This counts events such as PPP NCP failures.
<b>RejAdjs</b>	Indicates the number of times an adjacency has been rejected on this circuit.
<b>IDFieldLenMismatches</b>	Indicates the number of times an Intermediate-System-to-Intermediate-System (IS-IS) control PDU with an ID field length different from that for this system has been received.
<b>MaxAreaAddrMismatches</b>	Indicates the number of times an IS-IS control PDU with a max area address field different from that for this system has been received.
<b>AuthFails</b>	Indicates the number of times an IS-IS control PDU with the correct auth type has failed to pass authentication validation.
<b>LANDesISChanges</b>	Indicates the number of times the Designated IS has changed on this circuit at this level. If the circuit is point to point, this count is zero.
<b>AbsoluteValue</b>	Displays the counter value.
<b>Cumulative</b>	Displays the total value since you opened the Stats tab.
<b>Average/Sec</b>	Displays the average value for each second.
<b>Minimum/Sec</b>	Displays the minimum value for each second.
<b>Maximum/Sec</b>	Displays the maximum value for each second.
<b>Last Val/Sec</b>	Displays the last value for each second.

## Graphing IS-IS Interface Sending Control Packet Statistics

Use the following procedure to graph IS-IS interface receiving control packet statistics.

### Procedure

1. In the navigation pane, expand the **Configuration > IS-IS** folders.
2. Click **IS-IS**.
3. Click the **Interfaces** tab.
4. Select an existing interface.
5. Click the **Graph** button.
6. Click the **Interface Sending Control Packets** tab.

## Interface Sending Control Packets Field Descriptions

The following table describes the fields in the **Interface Sending Control Packets** tab.

Name	Description
<b>Hello</b>	Indicates the number of IS-IS Hello (IIH) PDUs seen in this direction at this level. Point-to-Point IIH PDUs are counted at the lowest enabled level: at L1 on L1 or L1L2 circuits, and at L2 otherwise.
<b>LSP</b>	Indicates the number of IS-IS LSP frames seen in this direction at this level.
<b>CSNP</b>	Indicates the number of IS-IS Complete Sequence Number Packet (CSNP) frames seen in this direction at this level.
<b>PSNP</b>	Indicates the number of IS-IS Partial Sequence Number Packets (PSNPs) seen in this direction at this level.
<b>AbsoluteValue</b>	Displays the counter value.
<b>Cumulative</b>	Displays the total value since you opened the Stats tab.
<b>Average/Sec</b>	Displays the average value for each second.
<b>Minimum/Sec</b>	Displays the minimum value for each second.
<b>Maximum/Sec</b>	Displays the maximum value for each second.
<b>Last Val/Sec</b>	Displays the last value for each second.

## Graphing IS-IS Interface Receiving Control Packet Statistics

Use the following procedure to graph IS-IS interface sending control packet statistics.

### Procedure

1. In the navigation pane, expand the **Configuration > IS-IS** folders.
2. Click **IS-IS**.
3. Click the **Interfaces** tab.
4. Select an existing interface.
5. Click the **Graph** button.
6. Click the **Interface Receiving Control Packets** tab.

## Interface Receiving Control Packets Field Descriptions

The following table describes the fields in the **Interface Receiving Control Packets** tab.

Name	Description
<b>Hello</b>	Indicates the number of IS-IS Hello PDUs seen in this direction at this level. Point-to-Point IIH PDUs are counted at the lowest enabled level: at L1 on L1 or L1L2 circuits, and at L2 otherwise.
<b>LSP</b>	Indicates the number of IS-IS link-state packet (LSP) frames seen in this direction at this level.

*Table continues...*

Name	Description
<b>CSNP</b>	Indicates the number of IS-IS Complete Sequence Number Packet (CSNP) frames seen in this direction at this level.
<b>PSNP</b>	Indicates the number of IS-IS Partial Sequence Number Packets (PSNPs) seen in this direction at this level.
<b>AbsoluteValue</b>	Displays the counter value.
<b>Cumulative</b>	Displays the total value since you opened the Stats tab.
<b>Average/Sec</b>	Displays the average value for each second.
<b>Minimum/Sec</b>	Displays the minimum value for each second.
<b>Maximum/Sec</b>	Displays the maximum value for each second.
<b>Last Val/Sec</b>	Displays the last value for each second.

---

## Graphing Stat Rate Limit Statistics for a Port

View stat rate limit statistics to view the total dropped packets and bytes.

### Procedure

1. In the Device Physical View, select a port.
2. In the navigation pane, expand the **Configuration > Graph** folders.
3. Click **Port**.
4. Click the **Stat Rate Limit** tab.
5. Select one or more values.
6. Click the type of graph to create.

## Stat Rate Limit Field Descriptions

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

Name	Description
<b>DropPktRate</b>	Indicates the drop packet rate.
<b>DropByteRate</b>	Indicates the drop byte rate.
<b>DropTotalBytes</b>	Indicates the total bytes dropped.
<b>DropTotalPkts</b>	Indicates the total packets dropped.

---

## Viewing IPv6 Statistics for an Interface

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



## Procedure

1. In the navigation pane, expand the **Configuration > IPv6** folders.
2. Click **IPv6**.
3. Click the **Interfaces** tab.
4. Select an interface.
5. Click **IfStats**.
6. **(Optional)** Select one or more values, and then click on the type of graph to graph the data.

## Statistics Field Descriptions

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

Name	Description
<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 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.

*Table continues...*

Name	Description
<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 <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

*Table continues...*

Name	Description
	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 ICMP Statistics

View ICMP statistics for ICMP configuration information.

### Procedure


1. In the navigation pane, expand the **Configuration > IPv6** folders.
2. Click **IPv6**.
3. Click **Interfaces** tab.
4. Select the interface on which you want to view the ICMP statistics.
5. Click **ICMPstats** option from the menu.

## ICMP stats Field Descriptions

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

Name	Description
<b>InMsgs</b>	Specifies the total number of ICMP messages which the entity received.

*Table continues...*

Name	Description
	<p> <b>Note:</b> This counter includes all those counted by icmpInErrors.</p>
<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>InTimeExclds</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.
<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.

*Table continues...*

Name	Description
<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 destination 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.
<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 IPv6 OSPF Statistics

View OSPF statistics to analyze trends. You can also graph statistics for all OSPF packets transmitted by the switch.

### Procedure

1. In the navigation pane, expand the **Configuration > IPv6** folders.
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.
<b>Routes</b>	Shows the number of OSPF routes added to the routing table.
<b>Adjacencies</b>	Shows the number of existing adjacencies.
<b>Areas</b>	Shows the number of configured areas.
<b>Nbrs</b>	Shows the number of OSPF neighbors.

## Viewing IPv6 VRRP Statistics

View IPv6 VRRP statistics to monitor network performance.

### Procedure

1. In the navigation pane, expand the **Configuration > IPv6** folders.
2. Click **VRRP**.

- Click the **Stats** tab.

## Stats Field Descriptions

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

Name	Description
<b>InetAddrType</b>	Shows the type of IP address (IPv4 or IPv6).
<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>VrldErrors</b>	Shows 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

- In the navigation pane, expand the **Configuration > IPv6** folders.
- Click **VRRP**.
- Click the **Interface** tab.
- Select an interface.
- 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.

*Table continues...*

Name	Description
<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 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.



## Configuring IPv6 VRRP Statistics

View IPv6 VRRP statistics for a VLAN or port.

### Before you begin

Change the VRF instance as required to view IPv6 VRRP statistics on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

### Procedure

1. In the navigation pane, expand the **Configuration > Edit > Port** folders.
2. Click **IPv6**.
3. Click the **VRRP** tab.
4. Select an interface.
5. Click **Statistics**.

## Viewing IP VRRPv3 Statistics

### About this task

Use the following procedure to view IPv6 VRRPv3 statistics for monitoring the network performance.

### Procedure

1. In the navigation pane, expand the **Configuration --> IP** folders.
2. Click **VRRP**.
3. Click the **V3 Stats** tab.

## V3 Stats Field Descriptions

Use the data in the following table to interpret the **V3 Stats** tab.

Name	Description
<b>InetAddrType</b>	Shows that the address type of the statistical entry is IPv4.
<b>ChecksumErrors</b>	Specifies the total number of VRRP packets received with an invalid VRRP checksum value.
<b>VersionErrors</b>	Specifies the total number of VRRP packets received with an unknown or unsupported version number.
<b>VrIdErrors</b>	Specifies the total number of VRRP packets received with an invalid VRID for the virtual router.

## Graphing IPv6 VRRP Statistics

### About this task

Use the following procedure to graph IPv6 VRRPv3 statistics for monitoring the network performance.

### Procedure

1. In the navigation pane, expand the **Configuration** --> **IPv6** folders.
2. Click **VRRP**.
3. Click the **Stats** tab.
4. Select an interface, and click **Graph**.
5. Select one or more values.
6. Select a graph type, click one of the icons in the upper-left corner of the menu bar. Your choices are:
  - Line Chart
  - Area Chart
  - Bar Chart
  - Pie Chart

## Stats Field Descriptions

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

Name	Description
<b>InetAddrType</b>	Shows the type of IP address (IPv4 or IPv6).
<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.

## Graphing IP VRRPv3 Statistics

### About this task

Use the following procedure to view and graph IP VRRPv3 statistics for monitoring the network performance.

## Procedure


1. In the navigation pane, expand the **Configuration --> IP** folders.
2. Click **VRRP**.
3. Click the **V3 Interface** tab.
4. Select an interface, and click **Graph**.
5. Select one or more values.
6. Select a graph type, click one of the icons in the upper-left corner of the menu bar. Your choices are:
  - Line Chart
  - Area Chart
  - Bar Chart
  - Pie Chart

## V3 Interface Field Descriptions

Use the data in the following table to use the **V3 Interface** tab.

Name	Description
<b>IfIndex</b>	Shows the index value that uniquely identifies the interface to which this entry applies.
<b>VrId</b>	Specifies a number that uniquely identifies a virtual router on a VRRP router.
<b>PrimaryIpAddr</b>	Specifies the virtual address assigned to the VRRP.
<b>VirtualMacAddr</b>	Specifies the MAC address of the virtual router interface.
<b>State</b>	Specifies the state of the virtual router interface: <ul style="list-style-type: none"> <li>• Initialize—waiting for a startup event</li> <li>• Backup—monitoring availability and state of the master router</li> <li>• Master—functioning as the forwarding router for the virtual router IP addresses.</li> </ul>
<b>Control</b>	Specifies whether VRRP is enabled or disabled for the port (or VLAN). The default is enabled.
<b>Priority</b>	Specifies the priority value used by this VRRP router. Set a value from 1 to 255, where 255 is reserved for the router that owns the IP addresses associated with the virtual router. The default is 100.
<b>AdvInterval</b>	Specifies the time interval (in seconds) between sending advertisement messages. The range is 1 to

*Table continues...*

Name	Description
	255 seconds with a default of 1 second. Only the master router sends advertisements. The default is 1.
<b>UpTime</b>	Specifies the time interval (in hundredths of a second) since the virtual router was initialized.
<b>CriticalIpAddr</b>	<p>This command specifies an IP interface on the local router, which is configured so that a change in its state causes a role switch in the virtual router (for example, from master to backup) in case the interface stops responding.</p> <p> <b>Note:</b></p> <p>In this context, local implies an address from the same VRF as the IP interface where VRRP is being configured.</p>
<b>CriticalIpAddrEnabled</b>	Configures the IP interface on the local router to enable or disable the backup. The default is disabled.
<b>BackUpMaster</b>	Enables the backup VRRP system traffic forwarding. The default is disabled.
<b>BackUpMasterState</b>	Indicates whether the backup VRRP system traffic forwarding is enabled or disabled. The default is disabled.
<b>FasterAdvIntervalEnabled</b>	Enables or disables the Fast Advertisement Interval. When disabled, the regular advertisement interval is used. The default is disable.
<b>FasterAdvInterval</b>	Configures the Fast Advertisement Interval between sending VRRP advertisement messages. The interval is between 200 and 1000 milliseconds, and you must enter the same value on all participating routers. The default is 200. You must enter the values in multiples of 200 milliseconds.
<b>PreemptMode</b>	Issued to preempt the existing router. If a new router is added to the network with its priority higher than the existing routers, then the new router becomes the master.
<b>Action</b>	<p>Lists options to override the delay timer manually and force preemption:</p> <ul style="list-style-type: none"> <li>• <b>none</b> does not override the timer</li> <li>• <b>preemptHoldDownTimer</b> preempts the timer</li> </ul>
<b>HoldDownTimer</b>	Indicates the hold-down state of this VRRP interface. If the hold-down timer is operational, this variable is

*Table continues...*

Name	Description
	set to active; otherwise, this variable is set to dormant.
<b>HoldDownTimeRemaining</b>	Indicates the amount of time (in seconds) left before the HoldDownTimer expires.
<b>MasterAdvInterval</b>	On the VRRPv3 master, the master advertisement interval is same as the advertisement interval. On the VRRPv3 Backup, the master advertisement interval is set to the Advertisement configured on the Master (received in the packet).

## Viewing IPv6 DHCP Relay Statistics for a Port

Display individual IPv6 DHCP Relay statistics for specific ports to manage network performance. You can also create a graph of selected statistical values.

### Procedure

1. On the Device Physical view, select a port.
2. In the navigation pane, expand the **Configuration > IPv6** folders.
3. Click the **DHCP Relay** tab.
4. Click the **Interface** tab.
5. Select the interface on which you want to view the IPv6 DHCP Relay statistics.
6. Click **Statistics**.
7. Select one or more values.
8. Click the type of graph.

## Statistics Field Descriptions

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

Name	Description
<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.

## Displaying IPsec Interface Statistics

Use this procedure to view IPsec statistics and counter values for each IPsec-enabled interface.

**About this task**

If you select an interface on the **Stats** tab, you can click **Graph** to graph particular statistics for that interface.

**Procedure**

1. In the navigation pane, expand the following folders: **Configuration > Security > Control Path**.
2. Click **IPSec**.
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>	Shows the interface index for which the statistic is captured.
<b>InSuccesses</b>	Specifies the number of ingress packets IPsec successfully carries.
<b>InSPViolations</b>	Specifies the number of ingress packets IPsec discards since boot time because of a security policy violation.
<b>InNotEnoughMemories</b>	Specifies the number of ingress packets IPsec discards since boot time because not enough memory is available.
<b>InAHESPReplays</b>	Specifies the number of ingress packets IPsec discards since boot time because the AH replay check fails.
<b>InESPReplays</b>	Specifies the number of ingress packets IPsec discards since boot time because the ESP replay check fails.
<b>InAHFailures</b>	Specifies the number of ingress packets IPsec discards since boot time because the AH authentication check fails.
<b>InESPFailures</b>	Specifies the number of ingress packets IPsec discards since boot time because the ESP authentication check fails.
<b>OutSuccesses</b>	Specifies the number of egress packets IPsec successfully carries since boot time.
<b>OutSPViolations</b>	Specifies the number of egress packets IPsec discards since boot time because a security policy violation occurs.

*Table continues...*

Name	Description
<b>OutNotEnoughMemories</b>	Specifies the number of egress packets IPsec discards since boot time because not enough memory is available since boot time.
<b>generalError</b>	Specifies a general error.
<b>InAhSuccesses</b>	Specifies the number of ingress packets IPsec carries because the AH authentication succeeds.
<b>OutAHSuccesses</b>	Specifies the number of egress packets IPsec successfully carries since boot time.
<b>InESPSuccesses</b>	Specifies the number of ingress packets IPsec carries since boot time because the ESP authentication succeeds.
<b>OutESPSuccesses</b>	Specifies the number of egress packets IPsec successfully carries since boot time.
<b>OutKBytes</b>	Specifies the total number of kilobytes on egress.
<b>OutBytes</b>	Specifies the total number of bytes on egress.
<b>InKBytes</b>	Specifies the total number of bytes on ingress.
<b>InBytes</b>	Specifies the total number of bytes on ingress.
<b>TotalPacketsProcessed</b>	Specifies the total number of packets processed.
<b>TotalPacketsByPassed</b>	Specifies the total number of packets bypassed.
<b>OutAHFailures</b>	Specifies the number of egress packets IPsec discards since boot time because the AH authentication check fails.
<b>OutESPFailures</b>	Specifies the number of egress packets IPsec discards since boot time because the ESP authentication check fails.
<b>InMD5Hmac</b>	Specifies the number of inbound HMAC MD5 occurrences since boot time.
<b>InSHA1Hmac</b>	Specifies the number of inbound HMAC SHA1 occurrences since boot time.
<b>InAESXCBCs</b>	Specifies the number of inbound AES XCBC MAC occurrences since boot time.
<b>InAnyNullAuth</b>	Specifies the number of inbound null authentication occurrences since boot time.
<b>In3DESCBCs</b>	Specifies the number of inbound 3DES CBC occurrences since boot time.
<b>InAESCBCs</b>	Specifies the number of inbound AES CBC occurrences since boot time.
<b>InAESCTRs</b>	Specifies the number of inbound AES CTR occurrences since boot time.

*Table continues...*

Name	Description
<b>InAnyNulEncrypt</b>	Specifies the number of inbound null occurrences since boot time. Used for debugging purposes.
<b>OutMD5Hmacs</b>	Specifies the number of outbound HMAC MD5 occurrences since boot time.
<b>OutSHA1Hmacs</b>	Specifies the number of outbound HMAC SHA1 occurrences since boot time.
<b>OutAESXCBCs</b>	Specifies the number of outbound AES XCBC MAC occurrences since boot time.
<b>OutInAnyNullAuth</b>	Specifies the number of outbound null authentication occurrences since boot time.
<b>Out3DESCBCs</b>	Specifies the number of outbound 3DES CBC occurrences since boot time.
<b>OutAESCBCs</b>	Specifies the number of outbound AES CBC occurrences since boot time.
<b>OutAESCTRs</b>	Specifies the number of outbound AES CTR occurrences since boot time.
<b>OutInAnyNullEncrypt</b>	Specifies the number of outbound null occurrences since boot time. Used for debugging purposes.

## Graphing IPsec Interface Statistics

Use this procedure to graphically view IPsec statistics and counter values for each IPsec-enabled interface.

### About this task

If you select an interface on the **Stats** tab, you can click **Graph** to graph particular statistics for that interface.

### Procedure

1. In the navigation pane, expand the **Security > Control Path** folders.
2. Click **IPSec**.
3. Click the **Interface Stats** tab.
4. Select a row, and click **Graph**.
5. Select one of the parameters, and click the appropriate icon in the upper-left corner of the menu bar to draw a line chart, area chart, bar chart, or a pie chart.
6. To clear existing counters, and fix a reference point in time to restart the counters, click **Clear Contents**.
7. To export the statistical data to a file, click **Export**.
8. To configure a poll interval, select an appropriate value from the **Poll Interval** drop-down list.



## Displaying Switch Level Statistics for IPsec-Enabled Interfaces

Use this procedure to view IPsec statistics and counter values at the switch level for all IPsec-enabled interfaces.

### Procedure

1. In the navigation pane, expand the following folders: **Configuration > Security > Control Path**.
2. Click **IPSec**.
3. Click the **Global Stats** tab.

### Global Stats Field Descriptions

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

Name	Description
<b>InSuccesses</b>	Specifies the number of ingress packets IPsec successfully carries.
<b>InSPViolations</b>	Specifies the number of ingress packets IPsec discards since boot time because of a security policy violation.
<b>InNotEnoughMemories</b>	Specifies the number of ingress packets IPsec discards since boot time because not enough memory is available.
<b>InAHESPReplays</b>	Specifies the number of ingress packets IPsec discards since boot time because the AH replay check fails.
<b>InESPReplays</b>	Specifies the number of ingress packets IPsec discards since boot time because the ESP replay check fails.
<b>InAHFailures</b>	Specifies the number of ingress packets IPsec discards since boot time because the AH authentication check fails.
<b>InESPFailures</b>	Specifies the number of ingress packets IPsec discards since boot time because the ESP authentication check fails.
<b>OutSuccesses</b>	Specifies the number of egress packets IPsec successfully carries since boot time.
<b>OutSPViolations</b>	Specifies the number of egress packets IPsec discards since boot time because a security policy violation occurs.

*Table continues...*

Name	Description
<b>OutNotEnoughMemories</b>	Specifies the number of egress packets IPsec discards since boot time because not enough memory is available since boot time.
<b>generalError</b>	Specifies a general error.
<b>InAHSuccesses</b>	Specifies the number of ingress packets IPsec carries because the AH authentication succeeds.
<b>OutAHSuccesses</b>	Specifies the number of egress packets IPsec successfully carries since boot time.
<b>InESPSuccesses</b>	Specifies the number of ingress packets IPsec carries since boot time because the ESP authentication succeeds.
<b>OutESPSuccesses</b>	Specifies the number of egress packets IPsec successfully carries since boot time.
<b>OutKBytes</b>	Specifies the total number of kilobytes on egress.
<b>OutBytes</b>	Specifies the total number of bytes on egress.
<b>InKBytes</b>	Specifies the total number of bytes on ingress.
<b>InBytes</b>	Specifies the total number of bytes on ingress.
<b>TotalPacketsProcessed</b>	Specifies the total number of packets processed.
<b>TotalPacketsByPassed</b>	Specifies the total number of packets bypassed.
<b>OutAHFailures</b>	Specifies the number of egress packets IPsec discards since boot time because the AH authentication check fails.
<b>OutESPFailures</b>	Specifies the number of egress packets IPsec discards since boot time because the ESP authentication check fails.
<b>InMD5Hmacs</b>	Specifies the number of inbound HMAC MD5 occurrences since boot time.
<b>InSHA1Hmacs</b>	Specifies the number of inbound HMAC SHA1 occurrences since boot time.
<b>InAESXCBCs</b>	Specifies the number of inbound AES XCBC MAC occurrences since boot time.
<b>InAnyNullAuth</b>	Specifies the number of inbound null authentication occurrences since boot time.
<b>In3DESCBCs</b>	Specifies the number of inbound 3DES CBC occurrences since boot time.
<b>InAESCBCs</b>	Specifies the number of inbound AES CBC occurrences since boot time.
<b>InAESCTRs</b>	Specifies the number of inbound AES CTR occurrences since boot time.

*Table continues...*

Name	Description
<b>InAnyNulEncrypt</b>	Specifies the number of inbound null occurrences since boot time. Used for debugging purposes.
<b>OutMD5Hmacs</b>	Specifies the number of outbound HMAC MD5 occurrences since boot time.
<b>OutSHA1Hmacs</b>	Specifies the number of outbound HMAC SHA1 occurrences since boot time.
<b>OutAESXCBCs</b>	Specifies the number of outbound AES XCBC MAC occurrences since boot time.
<b>OutInAnyNullAuth</b>	Specifies the number of outbound null authentication occurrences since boot time.
<b>Out3DESCBCs</b>	Specifies the number of outbound 3DES CBC occurrences since boot time.
<b>OutAESCBCs</b>	Specifies the number of outbound AES CBC occurrences since boot time.
<b>OutAESCTRs</b>	Specifies the number of outbound AES CTR occurrences since boot time.
<b>OutInAnyNullEncrypt</b>	Specifies the number of outbound null occurrences since boot time. Used for debugging purposes.

## Viewing EAPoL Authenticator Statistics

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

### Procedure

1. On the Device Physical View, select the port you want to graph.  
A yellow outline appears around the selected ports  
If you want to select multiple ports, press Ctrl and hold down the key while you click the ports you want to configure. A yellow outline appears around the selected ports.
2. In the navigation pane, expand the **Configuration > Graph** folders.
3. Click **Port**.
4. Click **EAPOL Stats**.
5. 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>InvalidFramesRx</b>	Displays the number of EAPoL frames received by this Authenticator in which the frame type is not recognized.
<b>EapLengthErrorFramesRx</b>	Displays the number of EAPoL frames received by this Authenticator in which the Packet Body Length field is invalid.
<b>StartFramesRx</b>	Displays the number of EAPoL start frames received by this Authenticator.
<b>EapFramesRx</b>	Displays the number of EAPoL-EAP frames received by this Authenticator.
<b>LogoffFramesRx</b>	Displays the number of EAPoL Logoff frames received by this Authenticator.
<b>LastRxFrameVersion</b>	Displays the last received version of the EAPoL frame by this Authenticator.
<b>LastRxFrameSource</b>	Displays the source MAC address of the last received EAPoL frame by this Authenticator.
<b>AuthEapFramesTx</b>	Displays the number of EAPoL-EAP frames transmitted by the Authenticator.

## Viewing Multihost Status Information

Use the following procedure to display multiple host status for a port.

### Procedure

1. In the navigation pane, expand the **Configuration --> Security --> Data Path** folders.
2. Click **802.1x-EAPOL**.
3. Click the **MultiHost Status** tab.

## MultiHost Status Field Descriptions

The following table describes values on the **MultiHost Status** tab.

Name	Description
<b>PortNumber</b>	Indicates the port number associated with this port.
<b>ClientMACAddr</b>	Indicates the MAC address of the client.
<b>PaeState</b>	Indicates the current state of the authenticator PAE state machine.
<b>VlanId</b>	Indicates the VLAN assigned to the client.
<b>Priority</b>	Specifies the priority associated with this client MAC. This priority could be the Radius assigned priority or the port QOS level.

## Viewing EAP Session Statistics

Use the following procedure to display multiple host session information for a port.

### Procedure

1. In the navigation pane, expand the **Configuration --> Security --> Data Path** folders.
2. Click **802.1x-EAPOL**.
3. Click the **MultiHost Session** tab.

## MultiHost Session Field Descriptions

The following table describes values on the **MultiHost Session** tab.

Name	Description
<b>StatsPortNumber</b>	Indicates the port number associated with this port.
<b>StatsClientMACAddr</b>	Indicates the MAC address of the client.
<b>Id</b>	Indicates the unique identifier for the session.
<b>AuthenticMethod</b>	Indicates the authentication method used to establish the session.
<b>Time</b>	Indicates the elapsed time of the session.
<b>TerminateCause</b>	Indicates the cause of the session termination.
<b>UserName</b>	Indicates the user name that represents the identity of the supplicant PAE.

## Viewing NEAP MAC Information

Use this procedure to view NEAP client MAC information on a port.

### Procedure

1. In the navigation pane, expand the **Configuration --> Security --> Data Path** folders.
2. Click **802.1x-EAPOL**.
3. Click the **NEAP Radius** tab.

## NEAP Radius Field Descriptions

The following table describes values on the **NEAP Radius** tab.

Name	Description
<b>MacPort</b>	Indicates the port number associated with this port.
<b>MacAddr</b>	Indicates the MAC address of the client.

*Table continues...*

Name	Description
<b>MacStatus</b>	Indicates the authentication status of the NEAP host that is authenticated using the RADIUS server.
<b>VlanId</b>	Indicates the VLAN assigned to the client.
<b>MacClear</b>	Clears the non EAP MAC entry associated with a specific index.
<b>MacPriority</b>	Specifies the priority associated with this Non-EAP client MAC. This priority could be the Radius assigned priority or the port QOS level.

## Viewing Secure Channel (SC) Outbound Statistics

Use this procedure to view the secure channel (SC) outbound statistics using EDM.

This feature is not supported on all hardware platforms. If you do not see this command in EDM, the feature is not supported on your hardware. For more information about feature support, see [Release Notes for VOSS](#).

### Procedure

1. In the Device Physical View tab, select the port for which you need to view the SC outbound statistics.

The switch supports MACsec on specific ports. For more information, see your hardware documentation.

2. In the navigation pane, expand the **Edit > Port > General** folders.
3. Click the **SC Outbound Stats** tab.

**\* Note:**

Use the **Clear Stats** button to clear single-port secure channel outbound statistics. The **Clear Stats** button is not available to clear multiple-port secure channel outbound statistics.

## SC Outbound Stats Field Descriptions

The following table describes the fields in the **SC Outbound Stats** tab.

Field	Description
<b>ProtectedPkts</b>	Specifies the number of integrity protected but not encrypted packets for this transmitting SC.
<b>EncryptedPkts</b>	Specifies the number of integrity protected and encrypted packets for this transmitting SC.

*Table continues...*

Field	Description
<b>OctetsProtected</b>	Specifies the number of plain text octets that are integrity protected but not encrypted on the transmitting SC.
<b>OctetsEncrypted</b>	Specifies the number of plain text octets that are integrity protected and encrypted on the transmitting SC.

## Viewing Secure Channel (SC) Inbound Statistics

Use this procedure to view the secure channel (SC) inbound statistics using EDM.

This feature is not supported on all hardware platforms. If you do not see this command in EDM, the feature is not supported on your hardware. For more information about feature support, see [Release Notes for VOSS](#).

### Procedure

1. In the Device Physical View tab, select the port for which you need to view the SC inbound statistics.

The switch supports MACsec on specific ports. For more information, see your hardware documentation.

2. In the navigation pane, expand the **Edit > Port > General** folders.
3. Click the **SC Inbound Stats** tab.

#### **Note:**



Use the **Clear Stats** button to clear single-port secure channel inbound statistics. The **Clear Stats** button is not available to clear multiple-port secure channel inbound statistics.

## SC Inbound Stats Field Descriptions

The following table describes the fields in the **SC Inbound Stats** tab.

Field	Description
<b>UnusedSAPkts</b>	Specifies the summary of received unencrypted packets on all SAs of this secure channel, with MACsec <i>not</i> in strict mode.
<b>NoUsingSAPkts</b>	Specifies the summary of received packets that were discarded along with either encrypted packets or packets that were received with MACsec operating in strict mode.

*Table continues...*

Field	Description
<b>LatePkts</b>	<p>Specifies the number of packets received that have been discarded for this secure channel (SC) with Replay Protect enabled.</p> <p> <b>Note:</b> The switch does not support Replay Protect.</p>
<b>NotValidPkts</b>	<p>Specifies the summary of packets that were discarded in all SAs of the SC because they were not valid with one of the following conditions:</p> <ul style="list-style-type: none"> <li>• MACsec was operating in strict mode.</li> <li>• The packets received were encrypted but contained erroneous fields.</li> </ul>
<b>InvalidPkts</b>	<p>Specifies the summary of all packets received that were not valid for this SC, with MACsec operating in <i>check</i> mode.</p>
<b>DelayedPkts</b>	<p>Specifies the summary of packets for this SC, with the packet number (PN) of the packets lower than the lower bound replay protection PN.</p> <p> <b>Note:</b> The switch does not support Replay Protect.</p>
<b>UncheckedPkts</b>	<p>The total number of packets for this SC that:</p> <ul style="list-style-type: none"> <li>• Were encrypted and had failed the integrity check.</li> <li>• Were <i>not</i> encrypted and had failed the integrity check.</li> <li>• Were received when MACsec validation was not enabled.</li> </ul>
<b>AcceptedPkts</b>	<p>Specifies the total number of Integrity Check Validated (ICV) packets for all SAs of this Secure Channel. The number of octets of User Data recovered from received frames that were integrity protected but not encrypted.</p>
<b>OctetsValidated</b>	<p>Specifies the number of octets of plaintext recovered from received packets that were integrity protected but not encrypted.</p>
<b>OctetsDecrypted</b>	<p>Specifies the number of octets of plaintext recovered from received packets that were integrity protected and encrypted.</p>

## Viewing MACsec Interface Statistics

Use this procedure to view the MACsec interface statistics using EDM.



This feature is not supported on all hardware platforms. If you do not see this command in EDM, the feature is not supported on your hardware. For more information about feature support, see [Release Notes for VOSS](#).

## Procedure

1. In the Device Physical View tab, select the port for which you need to view the MACsec interface statistics.  
  
The switch supports MACsec on specific ports. For more information, see your hardware documentation.
2. In the navigation pane, expand the **Edit > Port > General** folders.
3. Click the **MacSec Interface Stats** tab.

### **Note:**

Use the **Clear Stats** button to clear MACsec interface statistics. The **Clear Stats** button is available to clear single-port as well as multiple-port MACsec interface statistics.

## MacSec Interface Stats Field Descriptions

The following table describes the fields in the MacSec Interface Stats tab.

Field	Description
<b>TxUntaggedPkts</b>	Specifies the number of transmitted packets without the MAC security tag (SecTAG), with MACsec disabled on the interface.
<b>TxTooLongPkts</b>	Specifies the number of transmitted packets discarded because the packet length is greater than the maximum transmission unit (MTU) of the common port interface.
<b>RxUntaggedPkts</b>	Specifies the number of received packets without the MAC security tag (SecTAG), with MACsec <i>not</i> operating in strict mode.
<b>RxNoTagPkts</b>	Specifies the number of received packets without the MAC security tag (SecTAG), with MACsec operating in strict mode.
<b>RxBadTagPkts</b>	Specifies the number of received packets discarded with an invalid SecTAG, or with a zero value packet number (PN), or invalid Integrity Check Value (ICV).
<b>RxUnknownSCIPkts</b>	Specifies the number of packets received with an unknown secure channel identifier (SCI), and with MACsec <i>not</i> operating in strict mode.
<b>RxNoSCIPkts</b>	Specifies the number of packets received with an unknown secure channel identifier (SCI), and with MACsec operating in strict mode.

*Table continues...*

Field	Description
RxOverrunPkts	Specifies the number of packets discarded because the number of received packets exceeded the cryptographic performance capabilities.

## Viewing Segmented Management Instance Statistics

View operational statistics for the Management Instance.

### Procedure

1. In the navigation pane, expand the **Configuration > Edit** folders.
2. Click **Mgmt Instance**.
3. Click the **Mgmt Instance** tab.
4. Select a Management Instance by placing the cursor in a cell within the applicable row.
5. Click **Graph**.

## Interface Counters Field Descriptions

Use the data in the following table to use the Interface Counters tab.

Name	Description
RxPkts	Counts the packets received on the Segmented Management Instance.
RxError	Counts the packets received with errors on the Segmented Management Instance.
RxDrop	Counts the packets received and dropped on the Segmented Management Instance.
TxPkts	Counts the packets transmitted on the Segmented Management Instance.
TxError	Counts the packets transmitted with errors on the Segmented Management Instance.
TxDrop	Counts the packets dropped before transmission on the Segmented Management Instance.

## Displaying RADIUS CoA Reauthenticate Statistics

### About this task

Use this procedure to display RADIUS CoA Reauthenticate statistics.

## Procedure

1. In the navigation tree, expand the following folders: **Configuration > Security > Control Path**.
2. Click **RADIUS CoA**.
3. Click the **Reauthenticate Stats** tab.

## Reauthenticate Stats Field Descriptions

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

Name	Description
<b>Requests</b>	Specifies the number of RADIUS Reauthentication-Requests received from this Dynamic Authorization Client. This also includes the Reauthentication requests that have a Service-Type attribute with a value of Authorize Only.
<b>AuthOnlyRequests</b>	Specifies the number of RADIUS Reauthentication-Requests that include a Service-Type attribute with value Authorize Only received from this Dynamic Authorization Client.
<b>DupRequests</b>	Specifies the number of duplicate RADIUS Reauthentication-Request packets received from this Dynamic Authorization Client.
<b>Acks</b>	Specifies the number of incoming Reauthentication packets from this Dynamic Authorization Client silently discarded by the server application for some reason other than malformed, bad authenticators, or unknown types.
<b>Nacks</b>	Specifies the number of RADIUS Reauthentication-NAK packets sent to this Dynamic Authorization Client. This includes the RADIUS Reauthentication-NAK packets sent with a Service-Type attribute value of Authorize Only, and the RADIUS Reauthentication-NAK packets sent because no session context was found.
<b>NacksAuthOnlyRequests</b>	Specifies the number of RADIUS Reauthentication-NAK packets that include a Service-Type attribute with a value of Authorize Only sent to this Dynamic Authorization Client.
<b>NacksNoSess</b>	Specifies the number of RADIUS Reauthentication-NAK packets sent to this Dynamic Authorization Client because no session context was found.
<b>SessReauthenticated</b>	Specifies the number of user sessions reauthenticated for the Reauthentication-Requests received from this Dynamic Authorization Client.

*Table continues...*

Name	Description
	Depending on site-specific policies, a single Reauthentication-Request can change the authorization of multiple user sessions. In cases where the Dynamic Authorization Server has no knowledge of the number of user sessions that are affected by a single request, each CoA-Request counts as a single affected user session.
<b>Malformed</b>	Specifies the number of malformed RADIUS Reauthentication-Request packets received from the Dynamic Authorization Client. Bad authenticators and unknown types are not included as malformed Reauthentication-Requests.
<b>Dropped</b>	Specifies the number of incoming Reauthentication packets from the Dynamic Authorization Client silently discarded by the server application for some reason other than malformed, bad authenticators, or unknown types.
<b>BadAuths</b>	Specifies the number of RADIUS Reauthentication-Request packets that contained an invalid Authenticator field received from the Dynamic Authorization Client.
<b>UnknownTypes</b>	Specifies the number of incoming requests that have an invalid ID type.

## Displaying RADIUS CoA Disconnect Statistics

### About this task

Use this procedure to display RADIUS CoA Disconnect statistics.

### Procedure

1. In the navigation tree, expand the following folders: **Configuration > Security > Control Path**.
2. Click **RADIUS CoA**.
3. Click the **Disconnect Stats** tab.

## Disconnect Stats Field Descriptions

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

Name	Description
<b>DisconRequests</b>	Specifies the number of RADIUS disconnect requests received from this Dynamic Authorization Client. This also includes the RADIUS disconnect

*Table continues...*

Name	Description
	requests that have a Service-Type attribute with a value of Authorize Only.
<b>DisconAuthOnlyRequests</b>	Specifies the number of RADIUS disconnect requests that include a Service-Type attribute with a value of Authorize Only received from this Dynamic Authorization Client.
<b>DupDisconRequests</b>	Specifies the number of duplicate RADIUS Disconnect-Request packets received from this Dynamic Authorization Client.
<b>DisconAcks</b>	Specifies the number of RADIUS Disconnect-ACK packets sent to this Dynamic Authorization Client.
<b>DisconNaks</b>	Specifies the number of RADIUS Disconnect-NAK packets sent to this Dynamic Authorization Client. This includes the RADIUS Disconnect-NAK packets sent with a Service-Type attribute with a value of Authorize Only and the RADIUS Disconnect-NAK packets sent because no session context was found.
<b>DisconNakAuthOnlyRequests</b>	Specifies the number of RADIUS Disconnect-NAK packets that include a Service-Type attribute with a value of Authorize Only sent to this Dynamic Authorization Client.
<b>DisconNakSessNoContext</b>	Specifies the number of RADIUS Disconnect-NAK packets sent to this Dynamic Authorization Client because no session context was found.
<b>DisconUserSessRemoved</b>	Specifies the number of user sessions removed for the disconnect requests received from this Dynamic Authorization Client. Depending on site-specific policies, a single disconnect request can remove multiple user sessions. In cases where this Dynamic Authorization Server has no knowledge of the number of user sessions that are affected by a single request, each such disconnect request counts as a single affected user session only.
<b>MalformedDisconRequests</b>	Specifies the number of malformed RADIUS Disconnect-Request packets received from this Dynamic Authorization Client. Bad authenticators and unknown types are not included as malformed disconnect requests
<b>DisconBadAuthenticators</b>	Specifies the number of RADIUS Disconnect-Request packets that contain an invalid Authenticator field received from this Dynamic Authorization Client.
<b>DisconPacketsDropped</b>	Specifies the number of incoming disconnect requests from this Dynamic Authorization Client

*Table continues...*

Name	Description
	silently discarded by the server application for some reason other than malformed packets, bad authenticators, or unknown types.
<b>UnknownTypes</b>	Specifies the number of incoming requests that have an invalid ID type.

## Displaying RADIUS CoA Statistics

### About this task

Use this procedure to display information about RADIUS CoA statistics.

### Procedure

1. In the navigation tree, expand the following folders: **Configuration > Security > Control Path**.
2. Click **RADIUS CoA**.
3. Click the **CoA Stats** tab.

## CoA Stats Field Descriptions

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

Name	Description
<b>AddressType</b>	Specifies the RADIUS Dynamic Authorization Client IP address type.
<b>Address</b>	Specifies the RADIUS Dynamic Authorization Client IP address.
<b>DisconRequests</b>	Specifies the number of RADIUS Disconnect-Requests received from this Dynamic Authorization Client. This also includes the Disconnect requests that have a Service-Type attribute with a value of Authorize Only.
<b>DisconAuthOnlyRequests</b>	Specifies the number of RADIUS Disconnect-Requests that include a Service-Type attribute with value Authorize Only received from this Dynamic Authorization Client.
<b>DupDisconRequests</b>	Specifies the number of duplicate RADIUS Disconnect-Request packets received from this Dynamic Authorization Client.
<b>DisconAcks</b>	Specifies the number of RADIUS Disconnect-ACK packets sent to this Dynamic Authorization Client.

*Table continues...*

Name	Description
<b>DisconNaks</b>	Specifies the number of RADIUS Disconnect-ACK packets sent to this Dynamic Authorization Client. This includes the RADIUS Disconnect-ACK packets sent with a Service-Type attribute value of Authorize Only, and the RADIUS Disconnect-ACK packets sent because no session context was found.
<b>DisconNakAuthOnlyRequests</b>	Specifies the number of RADIUS Disconnect-NAK packets that include a Service-Type attribute with a value of Authorize Only sent to this Dynamic Authorization Client.
<b>DisconNakSessNoContext</b>	Specifies the number of RADIUS Disconnect-NAK packets sent to this Dynamic Authorization Client because no session context was found.
<b>DisconUserSessRemoved</b>	Specifies the number of user sessions removed for the Disconnect-Requests received from this Dynamic Authorization Client. Depending on site-specific policies, a single Disconnect-Request can change the authorization of multiple user sessions. In cases where the Dynamic Authorization Server has no knowledge of the number of user sessions that are affected by a single request, each Disconnect-Request counts as a single affected user session.
<b>MalformedDisconRequests</b>	Specifies the number of malformed RADIUS Disconnect-Request packets received from the Dynamic Authorization Client. Bad authenticators and unknown types are not included as malformed Disconnect-Requests.
<b>UnknownTypes</b>	Specifies the number of incoming requests that have an invalid ID type.

# Glossary

<b>American Standard Code for Information Interchange (ASCII)</b>	A code to represent 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 Autonomous System, and using an EGP to route packets to other Autonomous Systems.
<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 specific 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>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 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>External Data Representation (XDR)</b>	An IETF standard, RFC 1832, for the description and encoding of data.
<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>Intermediate System to Intermediate System (IS-IS)</b>	<p>Intermediate System to Intermediate System( IS-IS) is a link-state, interior gateway protocol. ISO terminology refers to routers as Intermediate Systems (IS), hence the name Intermediate System to Intermediate System (IS-IS). IS-IS operation is similar to Open Shortest Path First (OSPF).</p> <p>In Shortest Path Bridging MAC (SPBM) networks, IS-IS discovers network topology and builds shortest path trees between network nodes that IS-IS uses for forwarding unicast traffic and determining the forwarding table for multicast traffic. SPBM employs IS-IS as the interior gateway protocol and implements additional Type-Length-Values (TLVs) to support additional functionality.</p>
<b>Internet Control Message Protocol (ICMP)</b>	A collection of error conditions and control messages exchanged by IP modules in both hosts and gateways.
<b>Internet Group Management Protocol (IGMP)</b>	IGMP is a host membership protocol used to arbitrate membership in multicast services. IP multicast routers use IGMP to learn the existence of host group members on their directly attached subnets.

<b>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>	Layer 2 is the Data Link Layer of the OSI model. Examples of Layer 2 protocols are Ethernet and Frame Relay.
<b>Layer 3</b>	Layer 3 is the Network Layer of the OSI model. An example of a Layer 3 protocol is Internet Protocol (IP).
<b>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>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>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>	The MIB 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 flow statistics, deriving further flow properties, detecting flow expiration, passing flow records to the exporting process, and deleting flow records.
<b>multiplexing</b>	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).

<b>nanometer (nm)</b>	One billionth of a meter ( $10^{-9}$ meter). A unit of measure commonly used to express the wavelengths of light.
<b>NonVolatile Random Access Memory (NVRAM)</b>	Random Access Memory that retains its contents after electrical power turns off.
<b>Observation domain</b>	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.
<b>Observation point</b>	An observation point is a network location where you can observe IP packets. Examples include a port or a VLAN.
<b>Options data record</b>	The data record that contains values and scope information of the flow measurement parameters that correspond to an options template record.
<b>Options template flowset</b>	One or more options template records in an export packet.
<b>Options template record</b>	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.
<b>policing</b>	Ensures that a traffic stream follows the domain service-provisioning policy or service-level agreement (SLA).
<b>Port Access Entity (PAE)</b>	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).
<b>Protocol Data Units (PDUs)</b>	A unit of data that is specified in a protocol of a specific layer and that consists of protocol-control information of the specific layer and possibly user data of that layer.
<b>QSFP+</b>	A hot pluggable, quad small form-factor pluggable plus (QSFP+) transceiver, which is used in 40 Gbps and 4x10 Gbps Ethernet applications. 4x10 Gbps requires channelization support.
<b>QSFP28</b>	A hot pluggable, quad small form-factor pluggable 28 (QSFP28) transceiver, which is used in 100 Gbps and 4x25 Gbps Ethernet applications. 4x25 Gbps requires channelization support. It is similar in physical appearance to QSFP+ transceivers.
<b>quality of service (QoS)</b>	QoS features reserve resources in a congested network, allowing you to configure a higher priority for certain devices. For example, you can configure a higher priority for IP deskphones, which need a fixed bit rate

and split the remaining bandwidth between data connections if calls in the network are more important than the file transfers.

**Random Access Memory (RAM)**

Memory into which you can write and read data. A solid state memory device used for transient memory stores. You can enter and retrieve information from storage position.

**remote login (rlogin)**

An application that provides a terminal interface between hosts (usually UNIX) that use the TCP/IP network protocol. Unlike Telnet, rlogin assumes the remote host is, or behaves like, a UNIX host.

**remote monitoring (RMON)**

A remote monitoring standard for Simple Network Management Protocol (SNMP)-based management information bases (MIB). The Internetwork Engineering Task Force (IETF) proposed the RMON standard to provide guidelines for remote monitoring of individual LAN segments.

**sFlow agent**

Provides the interface for the sFlow instance. The agent maintains the measurement session with, and sends sFlow datagrams to, the sFlow collector.

**sFlow collector**

Receives sFlow datagrams from one or more sFlow agents.

**sFlow datagram**

A User Datagram Protocol (UDP) packet that contains the measurement information. The sFlow datagram also includes information about the source and process.

**SFP**

A hot pluggable, small form-factor pluggable (SFP) transceiver, which is used in Ethernet applications up to 1 Gbps.

**SFP+**

A hot pluggable, small form-factor pluggable plus (SFP+) transceiver, which is used in Ethernet applications up to 10 Gbps. It is similar in physical appearance to SFP transceivers.

**Shortest Path Bridging MAC (SPBM)**

Shortest Path Bridging MAC (SPBM) uses the Intermediate-System-to-Intermediate-System (IS-IS) link-state routing protocol to provide a loop-free Ethernet topology that creates a shortest-path topology from every node to every other node in the network based on node MAC addresses. SPBM uses the 802.1ah MAC-in-MAC frame format and encapsulates the source bridge identifier into the B-MAC header. SPBM eliminates the need for multiple overlay protocols in the core of the network by reducing the core to a single Ethernet-based link-state protocol, which can provide virtualization services, both layer 2 and layer 3, using a pure Ethernet technology base.

**shortest path first (SPF)**

A class of routing protocols that use Dijkstra's algorithm to compute the shortest path through a network, according to specified metrics, for efficient transmission of packet data.

<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.