ExtremeSwitching

Configuring System Monitoring on Ethernet Routing Switch 4900 and 5900 Series

Release 7.7
9035787 Rev.AA
June 2019

Legal Notice

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

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

Trademarks

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

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

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

Open Source Declarations

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

Contents

Chapter 1: About this document

Purpose

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

- Ethernet Routing Switch 4900 Series
- Ethernet Routing Switch 5900 Series

Operations include the following:

- Port mirroring
- Remote Network Monitoring (RMON)
- Remote Switch Port ANalyzer (RSPAN)

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

Conventions

This section discusses the conventions used in this guide.

Text Conventions

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

Table 1: Notice Icons

Icon	Alerts you to...
Important:	A situation that can cause serious inconvenience.
Note:	Important features or instructions.
Tip:	Helpful tips and notices for using the product.

Table continues…

Icon	Alerts you to...
Danger:	Situations that will result in severe bodily injury; up to and including death.
Warning:	Risk of severe personal injury or critical loss of data.
Caution:	Risk of personal injury, system damage, or loss of data.

Table 2: Text Conventions

Convention	Description
Angle brackets (< >)	Angle brackets (< >) indicate that you choose the text to enter based on the description inside the brackets. Do not type the brackets when you enter the command. If the command syntax is `cfm maintenance-domain maintenance-level <0-7>` , you can enter `cfm maintenance-domain maintenance-level 4`.
Bold text	Bold text indicates the GUI object name you must act upon. Examples: • Click **OK**. • On the **Tools** menu, choose **Options**.
Braces ({ })	Braces ({ }) indicate required elements in syntax descriptions. Do not type the braces when you enter the command. For example, if the command syntax is `ip address {A.B.C.D}`, you must enter the IP address in dotted, decimal notation.
Brackets ([ ])	Brackets ([ ]) indicate optional elements in syntax descriptions. Do not type the brackets when you enter the command. For example, if the command syntax is `show clock [detail]`, you can enter either `show clock` or `show clock detail`.
Ellipses (…)	An ellipsis (…) indicates that you repeat the last element of the command as needed. For example, if the command syntax is `ethernet/2/1 [ <parameter> <value> ]...`, you enter `ethernet/2/1` and as many parameter-value pairs as you need.

Table continues…

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

Documentation and Training

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

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

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/.

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

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.
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.

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 an analytics solution that combines the Deep Packet Inspection capabilities of Extreme Analytics Engine with sFlow data. This feature provides granular visibility into your network and monitors application performance, users, locations, and devices without the need for expensive sensors or collectors.

Application Telemetry uses policy rules to filter packets for analysis. This methodology enables Application Telemetry to monitor *all* application-level traffic flows at wire speed on *all* specified interfaces simultaneously.

Note:

Application Telemetry support with Extreme Management Center will be provided in a future release.

For more information, see:

- Application Telemetry Fundamentals on page 174
- Application Telemetry Configuration using CLI on page 178
- Application Telemetry Configuration using EDM on page 184

sFlow

sFlow content is now located to this document. See sFlow on page 164.

Chapter 3: System monitoring fundamentals

System monitoring is an important aspect of switch operation. The switch provides a wide range of system monitoring options that the administrator can use to closely follow the operation of a switch or stack.

This chapter describes two general system monitoring considerations, system logging and port mirroring, for the switch. Subsequent chapters provide information about specific system monitoring tools and their use.

CPU and memory utilization

The CPU utilization feature provides data for CPU and memory utilization. You can view CPU utilization information for the past 10 seconds, 1 minute, 10 minutes, 1 hour, 24 hours, or since system bootup. The switch displays CPU utilization as a percentage. You can use CPU utilization information to see how the CPU is used during a specific time interval.

The memory utilization provides you information on what percentage of the dynamic memory is currently used by the system. The switch displays memory utilization in terms of megabytes available since system bootup.

This feature does not require a configuration. It is a display-only feature.

Light Emitting Diode (LED)

The switch displays diagnostic and operation information through the LEDs on the unit.

For detailed information regarding the interpretation of the LEDs, see:

- Installing Ethernet Routing Switch 4900 Series.
- Installing Ethernet Routing Switch 5900 Series.

Remote logging

The remote logging feature provides an enhanced level of logging by replicating system messages on a syslog server. System log messages from several switches can be collected at a central location, alleviating the network manager from querying each switch individually to interrogate the log files.

You must configure the remote syslog server to log informational messages to this remote server. The User Datagram Protocol (UDP) packet is sent to port 514 of the configured remote syslog server.

After the IP address is in the system, syslog messages can be sent to the remote syslog server. If a syslog message is generated prior to capturing the IP address of the server, the system stores up to 10 messages that are sent after the IP address of the remote server is on the system.

You can configure this feature by enabling remote logging, specifying the IP address of the remote syslog server, and specifying the severity level of the messages to be sent to the remote server.

Dual syslog server support

You can enable dual syslog server support by configuring and enabling a secondary remote syslog server to run in tandem with the first. The system then sends syslog messages simultaneously to both servers to ensure that syslog messages are logged, even if one of the servers becomes unavailable. See Configuring remote system logging using EDM on page 51

SNMP traps

SNMP traps are configured as notification controls. For more information about notification controls, see Configuring Security on Ethernet Routing Switch 4900 and 5900 Series.

MIB Web page

With Web-based management, you can see the response of an SNMP Get and Get-Next request for an Object Identifier (OID) or object name.

With the SNMP walk, you can retrieve a subtree of the Management Information Base (MIB) that has the object as root by using Get-Next requests.

The MIB Web page does not support the following features:

- displaying SNMP SET requests
- displaying SNMP tables
- translating MIB enumerations (that is, displaying the name [interpretation] of number values of objects defined as enumerations in the MIB)

IGMP and the system event log

Internet Group Management Protocol (IGMP) uses the components provided by the syslog tool. Functions such as storing messages in the Non-volatile Random Access Memory (NVRAM) or remote host, and displaying these log messages through the CLI or Telnet is then carried out by the syslog tool on its own.

The IGMP log events can be classified into the following three categories based on their severity:

- critical
- serious
- informational

IGMP logs in the messages whenever any of the following types of events take place in the system:

- IGMP initialization
- configuration changes
- Stack join events
- IGMP messages: report, leave, and query messages received by the switch

Important:

Events such as reception of IGMP messages happen frequently in the switch, whenever a new host joins or leaves a group. Logging such messages consumes a lot of log memory. Therefore, such messages should not be logged all the time. By default, logging of such messages is disabled. You must enable this feature through the CLI.

In the table Table 3: IGMP syslog messages on page 16:

- %d represents a decimal value for the parameter preceding it, for example, 5 for Virtual Local Area Network (VLAN) 5
- %x represents a hexadecimal value for the parameter preceding it, for example, 0xe0000a01 for Group 224.0.10.1

The following table describes the IGMP syslog messages and their severity.

Table 3: IGMP syslog messages

Severity	Log Messages
Informational	IGMP initialization success
Critical	IGMP initialization failed: Error code %d
Informational	IGMP: policy initialization success
Informational	IGMP:policy initialization failed
Informational	IGMP configuration loaded successfully
Informational	IGMP configuration failed: Loaded to factory default

Table continues…

Severity	Log Messages
Informational	IGMP: Version %d Snooping enabled on VLAN %d
Informational	IGMP: Version %d Snooping disabled on VLAN %d
Informational	IGMP: Proxy enabled on VLAN %d
Informational	IGMP: IGMP version %d enabled on VLAN %d
Informational	IGMP: IGMP version %d disabled on VLAN %d
Informational	IGMP: Proxy disabled on VLAN %d
Informational	IGMP configuration changed: Query time set to %d on VLAN %d
Informational	IGMP configuration changed: Robust value set to %d on VLAN %d
Informational	IGMP configuration changed: Version %d router port mask 0x%x set on VLAN %d
Informational	IGMP configuration changed: Unknown multicast filter enabled
Informational	IGMP configuration changed: Unknown multicast filter disabled
Informational	IGMP: Added reserved multicast address
Informational	IGMP: Removed reserved multicast address
Informational	IGMP: Unable to add reserved multicast address
Informational	IGMP: Exceeded reserved multicast address range: #Addr %d * #VLANs %d > %d
Informational	IGMP configuration changed: Trunk %d created for IGMP
Informational	IGMP: Trunk %d created. IGMP groups added on all trunk ports
Informational	IGMP configuration changed: Trunk %d removed for IGMP ports
Informational	IGMP: Trunk %d removed. IGMP groups removed on all trunk ports
Informational	IGMP configuration changed: Mirror ports set
Informational	IGMP configuration changed: Port %d added to VLAN %d
Informational	IGMP configuration changed: Port %d removed from VLAN %d
Informational	IGMP new Querier IP %x learned on port %d
Informational	IGMP: Dynamic router port %d added
Informational	IGMP: Dynamic router port %d removed
Informational	IGMP: Config. database sent by unit %d
Informational	IGMP: Config. database received on unit %d from %d
Informational	IGMP: Config database exchanged between all units of the stack
Informational	IGMP: Error sending database from unit %d
Informational	IGMP stack join completed. Database synchronized
Serious	IGMP not able to join stack: Error code %d
Informational	IGMP: Group database sent by unit %d
Informational	IGMP Group database received on unit %d from %d
Informational	IGMP: Group database received from all non-base units
Informational	IGMP: Error sending group database from unit %d

Table continues…

Severity	Log Messages
Informational	IGMP: REPORT received for Group %s on VLAN %d and port %d
Informational	IGMP: LEAVE received for Group %s on VLAN %d and port %d
Informational	IGMP: QUERY received on port %d

Stack Monitor

You use the Stack Monitor feature to analyze the health of a stack by monitoring the number of active units in the stack.

With stacked switches, multilink trunking (MLT) links are often connected to separate units in a distributed MLT (DMLT). If the connections between switches in the stack fail, a situation can arise where the DMLT links are no longer connected to a stack, but to a combination of units that are no longer connected to each other. From the other end of the DMLT, the trunk links appear to be functioning properly. However, the traffic is no longer flowing across the cascade connections to all units, so the connectivity problems can occur.

With the Stack Monitor feature, when a stack is broken, the stack and any disconnected units from the stack, send Simple Network Management Protocol (SNMP) traps. If the stack or the disconnected units are still connected to the network, they generate log events and send trap messages to the management station to notify the administrator of the event. After the problem is detected, the stack and disconnected units continue to generate log events and send traps at a user-configurable interval until the situation is remedied (or the feature is disabled).

Local ports shutdown while stacking

When a switch is joining the stack, DMLT and dynamic Link Aggregation Groups (LAG) formed with Link Aggregation Protocol (LACP) can still be created because Link Layer Discovery Protocol Data Units (LACPDU) continue to be transmitted. This results in a temporary traffic delay (for a few seconds) until the switch fully joins the stack.

To solve this issue, the switch momentarily shuts down the local ports before joining the stack. After a reset or power up, if the switch detects power on its stacking cables and is connected to another unit, the switch shuts down all of its local ports. When the ports are disabled, the port LEDs blink, similar to ports that are shut down. The ports are reenabled when the unit finishes entering the stack formation or after a 60-second timeout, whichever comes first.

If the unit does not detect power on the stacking ports 20 seconds after it comes up, the local ports forward the traffic.

Stack loopback test

The stack loopback test feature allows the customer to quickly test the switch stack ports and the stack cables on the switches. This feature helps you while experiencing stack problems to determine whether the root cause is a bad stack cable or a damaged stack port and prevents potentially good switches being returned for service. You can achieve this by using two types of loopback tests:

- Internal loopback test
- External loopback test

Caution:

For accurate results, run the internal loopback test before the external loopback test.

Internal loopback test

Use the internal loopback test by putting each of stack links in loopback mode one by one, sending 1000 packets, and verifying that the packets are received back with the same content.

The purpose of the internal loopback test is to verify that all the stack ports are functional.

External loopback test

Use the external loopback test by connecting the stack uplink port, with the stack downlink port, sending 1000 packets from the uplink port and verifying that the packets are received back on the downlink port. The same tests are done by sending the packets from the downlink port and verifying that they are received back on the uplink port. The purpose of the external loopback test is to verify that the stack cable is functional.

Run the internal test before the external test and before the stack ports are verified to be functional.

On known good units and stack cables, no errors are returned by the internal and the external loopback test. The external loopback test returns an error if the stack cable is not present.

The main limitation of this feature is that it interferes with the normal functioning of the stack manager. Therefore, you must run both the tests on units that are taken off the stack.

Important:

Hardware Limitation: This feature is only useful for stackable switches.

Software Limitation: You can execute only one test at a time. If a test is started and not finished, a second test cannot be started until the first stops.

Stack Health Check

You can use Stack Health Check to:

- provide information on the stacking state of each switch rear port
- run a high-level test to monitor the rear port status for each unit
- confirm the number of switching units in stack
- detect whether the stack runs with a temporary base unit
- monitor the stack continuity

By default, the health check is enabled on all the stack units. You can use Stack Health Check in both user interfaces: CLI and EDM.

Port mirroring

You can designate a switch port to monitor traffic in the following ways:

- on any two specified switch ports, port-based
- to or from any two specified addresses that the switch learns, address-based

You must connect an Ethernet monitoring device to the designated monitor port to connect the mirrored traffic.

When you enable Port-Mirroring with one of the following modes, higher available precedence will be used for all ports:

- Asrc
- Adst
- AsrcBdst
- AsrcBdstOrBsrcAdst
- AsrcOrAdst
- XrxYtxOrYrxXtx
- XrxYtx
- Xtx
- Xrx
- XrxYtx
- XrxOrYtx
- XrxOrXtx

Note:

You cannot free resources used by Port Mirroring with the **`qos agent reset-default`** command

Important:

If a unit leaving the stack causes invalid port-mirroring instances or RSPAN destination instances, these instances will not be displayed in the ASCII running config file. The `show port-mirroring [rspan]` command output indicates invalid RSPAN or port-mirroring instances by marking them with an asterisk (*) character after the instance number.

The output may vary from unit to unit, for the same instance. For example, consider a port-mirroring instance with all configured ports residing on unit 2. When unit 2 leaves the stack, this instance becomes invalid on stack but remains valid on unit 2.

Note:

Each of the XrxorXtx, XrxOrYtx, ManyToOneRxTx modes needs twice the hardware resources of a usual port mirroring instance. This means whenever you use one or more of these modes, instead of configuring up to four port-mirroring instances, you can only configure up to:

- two instances, if both instances are of type XrxOrYtx or ManyToOneRxTx, in any combination
- three instances if one, and only one, of these instances is of type XrxorXtx or XrxOrYtx or ManyToOneRxTx

Four instances can be configured if two of these instances are of type XrxorXtx.

Many-to-Many Port Mirroring

You can use the many-to-many port mirroring feature to configure multiple sessions of mirroring configurations, each with a monitor port and mirrored ports.

You can provide a way to monitor more than one traffic pattern by using many-to-many port mirroring. You can use this feature to monitor multiple traffic patterns, which is important in networks which support a variety of complex user scenarios. As an example, you can set up port mirroring to allow duplication of VoIP traffic for call recording, another instance for intrusion detections, and still another instance for other activities or network troubleshooting.

You can configure this feature by using CLI or EDM. To configure each instance, you follow the same configuration process as the port mirroring configuration.

Port-based modes

The following port-based modes are supported:

- ManytoOneRx: Many-to-One port mirroring on ingress packets.
- ManytoOneTx: Many-to-One port mirroring on egress packets.

- ManytoOneRxTx: Many-to-One port mirroring on ingress and egress traffic.
- Xrx: Mirror packets received on port X.
- Xtx: Mirror packets transmitted on port X.
- XrxOrXtx: Mirror packets received or transmitted on port X.
- XrxYtx: Mirror packets received on port X and transmitted on port Y.
- XrxYtxOrYrxXtx: Mirror packets received on port X and transmitted on port Y or packets received on port Y and transmitted on port X.
- XrxOrYtx: Mirror packets received on port X or transmitted on port Y.

Address-based modes

The following address-based modes are supported:

- Asrc: Mirror packets with source MAC address A.
- Adst: Mirror packets with destination MAC address A
- AsrcOrAdst: Mirror packets with source or destination MAC address A.
- AsrcBdst: Mirror packets with source MAC address A and destination MAC address B.
- AsrcBdstOrBsrcAdst: Mirror packets with source MAC address A and destination MAC address B or packets with source MAC address B and destination MAC address A.

Many-to-Many Port Mirroring limitations and restrictions

You can use many-to-many port mirroring on both pure stacks and standalone boxes.

You cannot configure a monitor port (MTP) that is a mirrored port for another MTP. Frames mirrored to one MTP are not taken into account in MAC address-based mirroring on another MTP.

If you configure a port to be egress-mirrored in one instance, then that port cannot be egress-mirrored in another instance (to another MTP). Similarly, if you configure a port to be ingress-mirrored, then the system prohibits that port to be ingress-mirrored in another instance. The system allows a port to be ingress-mirrored in one instance and egress-mirrored in another.

The ports you configure as monitor ports may be allowed to participate in normal frame switching operation or be used as management ports, provided that you enable port mirroring with the allow traffic option.

You can configure up to four monitor ports.

You can configure multiple instances by using the existing interface in CLI or EDM. The system attaches one monitor port (MTP) to each instance. In some cases a monitor port can be used in more than one instance.

You cannot configure a port as a monitor port if it exists as part of an MLT group.

For MAC base modes: Asrc, Adst, AsrcBdst, AsrcBdstOrBsrcAdst, AsrcOrAdst and port based modes: XrxYtx, XrxYtxOrYrxXtx port-mirroring, you need to install filters to enable port mirroring. The application may not function in these modes if platform resource limits are reached.

Port-based mirroring configuration

The following image is an example of a port-based mirroring configuration in which port 44 is designated as the monitor port for ports 45 and 46 of Switch S1. Although this example shows ports 45 and 46 monitored by the monitor port (port 44), you can monitor any of the trunk members of T1 and T2.

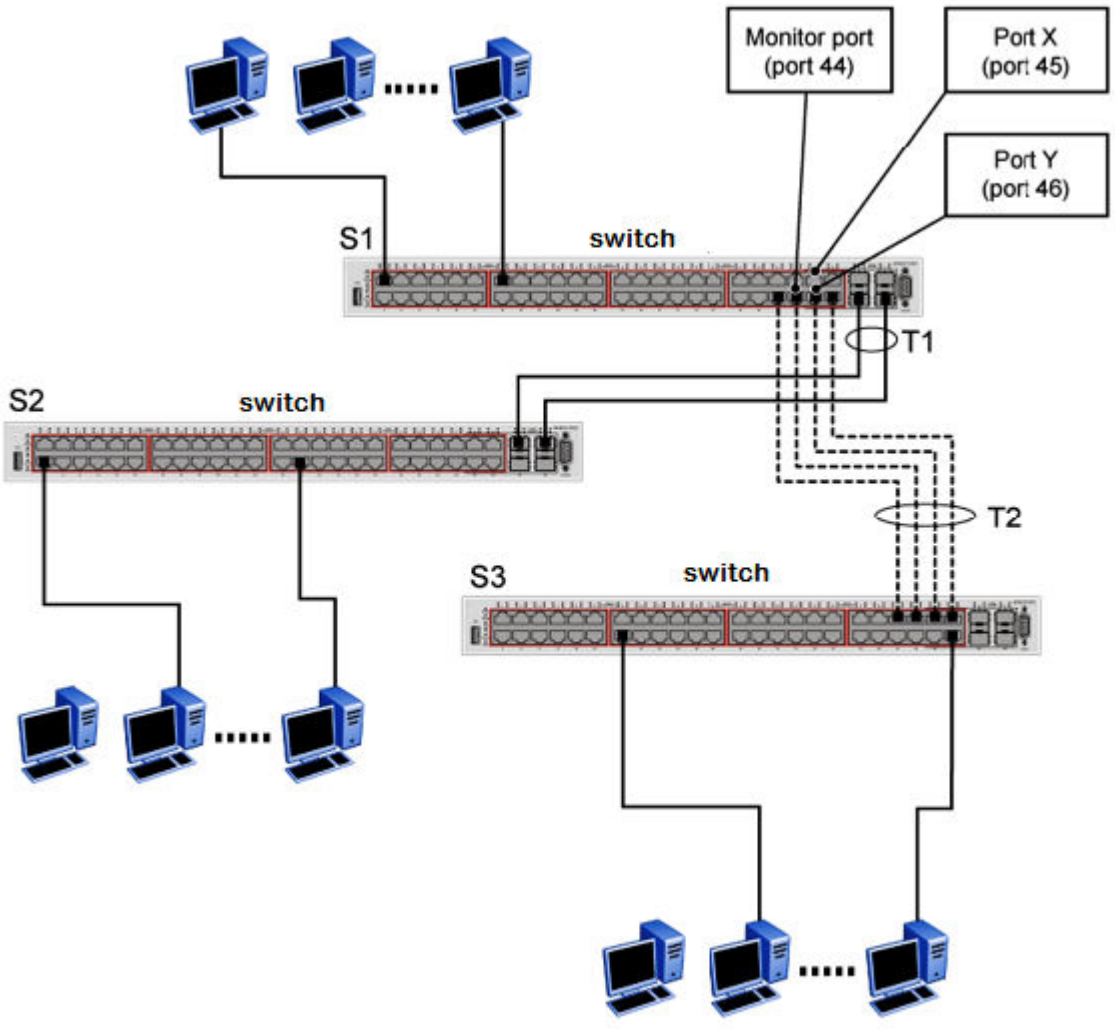

Figure 1: Port-based mirroring example

This example shows port X and port Y as members of Trunk T1 and Trunk T2. Port X and port Y are not required to always be members of Trunk T1 and Trunk T2.

Important:

You cannot configure trunk members as monitor port.

In the configuration example shown in the preceding figure, you can set the designated monitor port (port 44) to monitor traffic in any of the following modes:

- Monitor all traffic received by port X.
- Monitor all traffic transmitted by port X.
- Monitor all traffic received and transmitted by port X.
- Monitor all traffic received by port X or transmitted by port Y.
- Monitor all traffic received by port X (destined to port Y) and then transmitted by port Y.
- Monitor all traffic received/transmitted by port X and transmitted/received by port Y (conversations between port X and port Y).
- Monitor all traffic received on many ports (ManytoOneRX).
- Monitor all traffic transmitted on many ports (ManytoOneTX).
- Monitor all traffic received or transmitted on many ports (ManytoOneRxTX).

Address-based mirroring configuration

The following figure shows an example of an address-based mirroring configuration in which port 44, the designated monitor port for Switch S1, monitors traffic occurring between address A and address B.

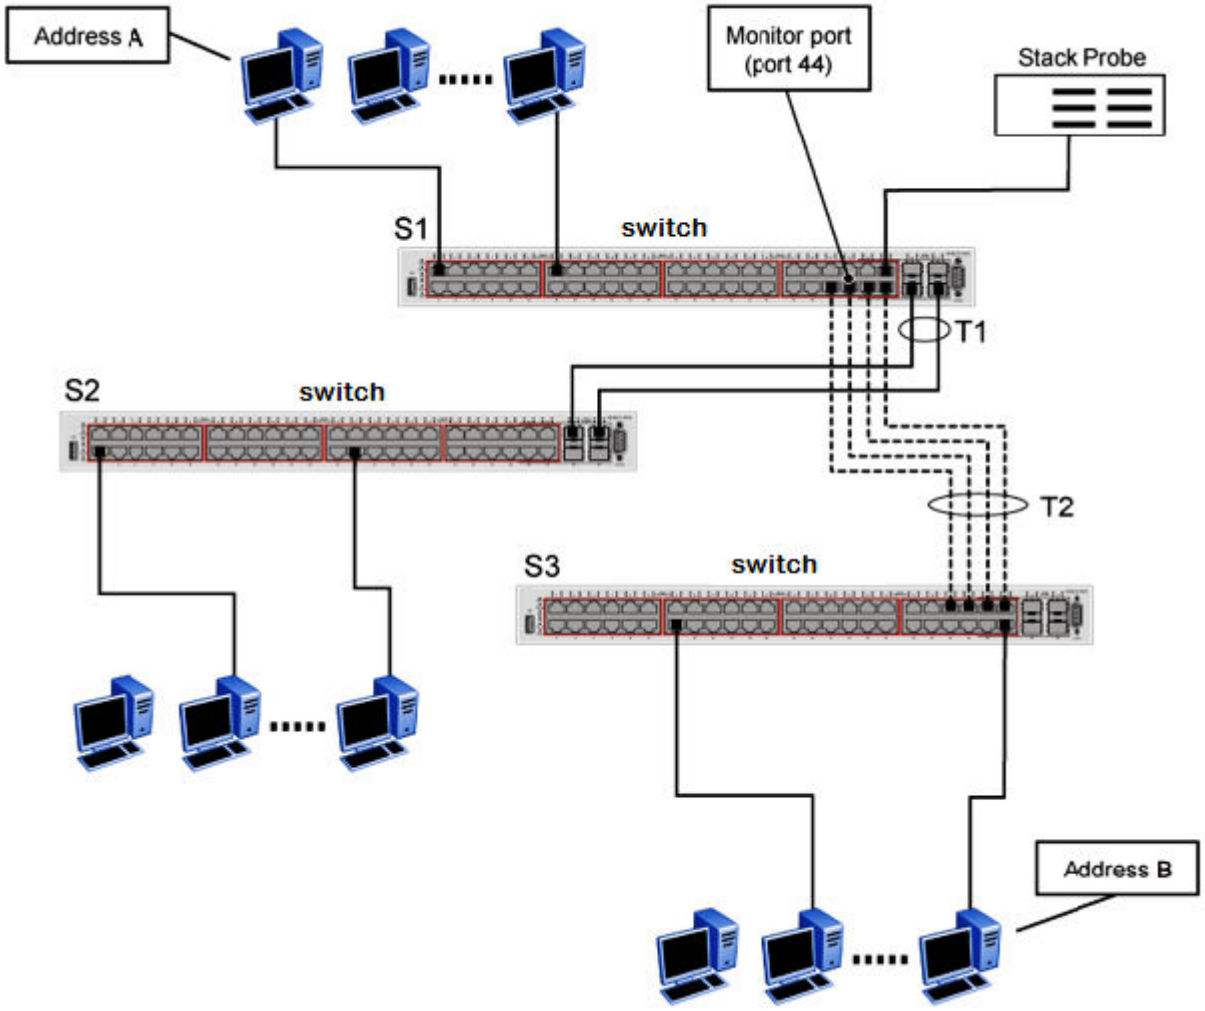

Figure 2: Address-based mirroring example

In this configuration, you can set the designated monitor port (port 44) to monitor traffic in any of the following modes:

- Monitor all traffic transmitted from address A to any address.
- Monitor all traffic received by address A from any address.
- Monitor all traffic received by or transmitted by address A.
- Monitor all traffic transmitted by address A to address B.
- Monitor all traffic between address A and address B (conversation between the two stations).

Port-mirroring on 802.1x (EAP) port

With this enhancement, you can enable or disable port mirroring on EAP ports.

A warning about the potential security risk is displayed in each of the following situations:

- when enabling the feature
- when mirroring an EAP port or trying to enable EAP on a mirror port
- when disabling the feature if there are mirroring ports with EAP enabled

Port VLAN based mirroring

Port VLAN based mirroring enhances port mirroring by enabling traffic to be mirrored from a specific VLAN by using filters. All port mirroring modes are supported.

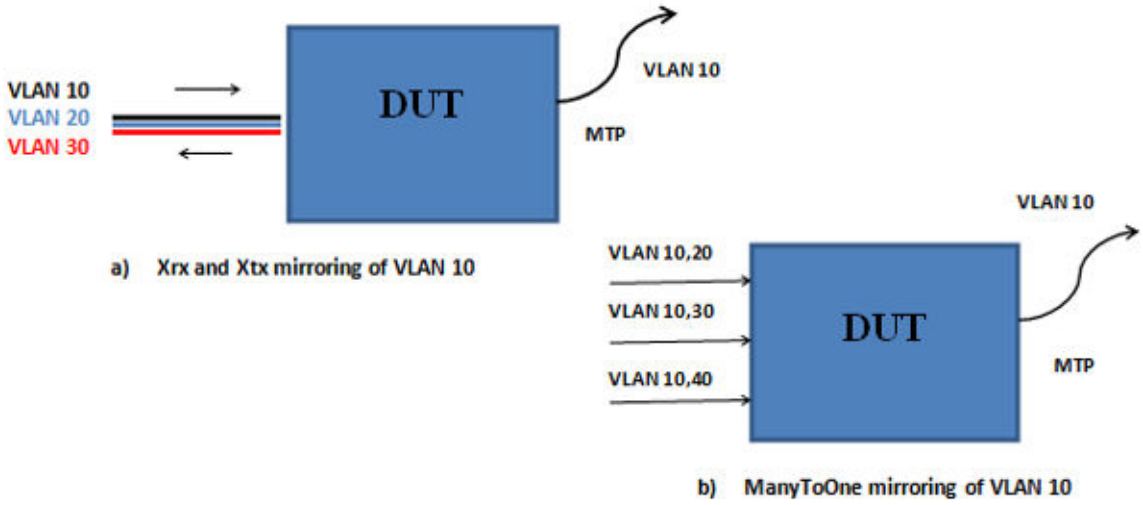

In the illustration, in example a, only the traffic from VLAN 10 (traffic received, sent, or both) is mirrored. In example b, there are multiple mirror ports, each of which is a member of several VLANs. This group of mirror ports is treated as a whole and only the traffic from VLAN 10 is copied on the sniffer port (MTP).

The following VLAN based mirroring modes are supported:

- Xrx: Mirror packets received on port X, from VLAN Y.
- Xtx: Mirror packets transmitted on port X, into VLAN Y.
- XrxOrXtx: Mirror packets received on port X from VLAN Y or packets transmitted on port X into VLAN Y.
- XrxOrYtx: Mirror packets received on port X from VLAN Z or packets transmitted on port Y into VLAN Z.

- XrxYtx: Mirror packets received on port X, VLAN Z and transmitted on port Y.
- XrxYtxOrYrxXtx: Mirror packets received on port X, VLAN Z and transmitted on port Y, or packets received on port Y, VLAN Z and transmitted on port X.
- ManytoOneRx: Many to one port mirroring ingress traffic from VLAN X.
- ManytoOneRxTx: Many to one port mirroring ingress traffic from VLAN X & egress traffic sent in VLAN X.
- ManytoOneTx: Many to one port mirroring egress traffic sent in VLAN X.
- Adst: Mirror packets from VLAN X with destination MAC address A
- Asrc: Mirror packets from VLAN X with source MAC address A
- AsrcBdst: Mirror packets from VLAN X with source MAC address A and destination MAC address B.
- AsrcBdstOrBsrcAdst: Mirror packets from VLAN X with source MAC address A and destination MAC address B, or packets from VLAN X with source MAC address B and destination MAC address A.
- AsrcOrAdst: Mirror packets from VLAN X with source or destination MAC address A.

RSPAN

Remote Switch Port Analyzer (RSPAN), also known as Remote Port Mirroring, enhances port mirroring by enabling mirrored traffic to be sent to one or more switches or stacks on the network. All participating switches must support the RSPAN feature.

For each RSPAN session, the mirrored traffic is carried over a user-specified RSPAN VLAN that is dedicated for that RSPAN session in all participating switches.

RSPAN consists of at least one RSPAN source session, an RSPAN VLAN, and at least one RSPAN destination session. The RSPAN traffic from the source ports is copied into the RSPAN VLAN and forwarded to a destination session monitoring the RSPAN VLAN. The final destination must always be a physical port on the destination switch. You can also include intermediate switches separating the RSPAN source and destination sessions. You separately configure RSPAN on the source switch, the intermediate switch(es), and on the destination switch.

You must create an RSPAN VLAN on each device involved in an RSPAN session.

RSPAN VLAN is a port based VLAN, carrying traffic between RSPAN source and destination sessions. You can have multiple RSPAN VLANs in a network at the same time, with each RSPAN VLAN defining a network-wide RSPAN session.

You can configure up to four RSPAN VLANs on a switch and up to three RSPAN instances.

For a minimal RSPAN configuration, you need:

- one RSPAN port on a source RSPAN session

- two ports on a destination RSPAN session (one port as a network port and one as an RSPAN destination port).

Note:

On an intermediate switch, it is recommended that you configure up to 12 ports.

Note:

Due to hardware limitations, RSPAN is not compatible with ERS 8800.

The following figure shows how the RSPAN is working for three connected devices:

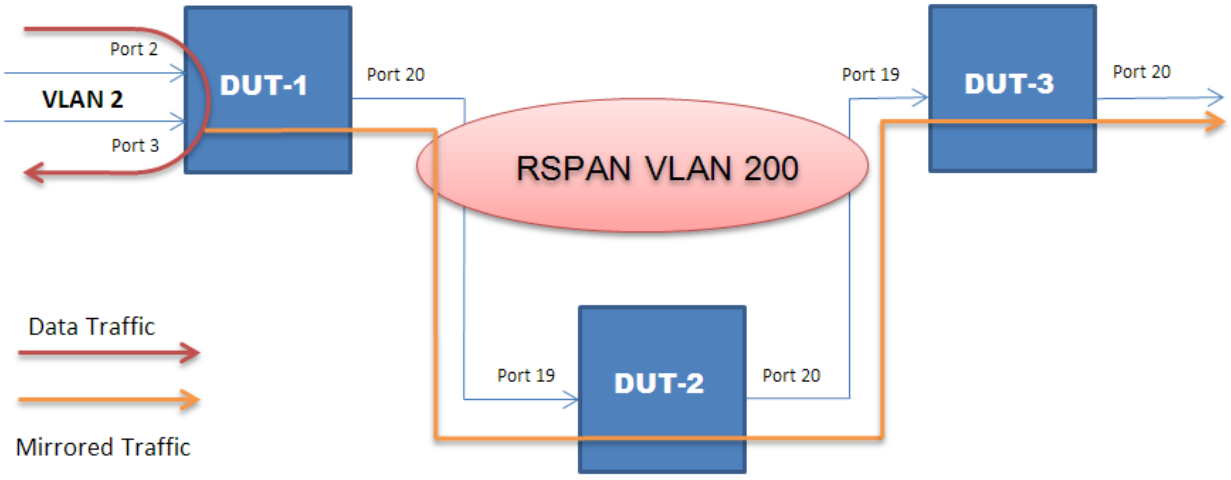

RSPAN source sessions

To configure an RSPAN source session on a source switch, you associate a port mirroring instance with an RSPAN VLAN. The output of this session is a stream of packets sent to the RSPAN VLAN. An RSPAN source session is very similar to a local port mirroring session, except that the packet stream is directed to the RSPAN VLAN. In an RSPAN instance, the mirrored packets are supplementary tagged with the RSPAN VLAN ID and directed to the destination switch. When exiting the source switch, the RSPAN traffic has both vlan labels (double tagging).

You can have more than one source session active in the same RSPAN VLAN, each source session on a separate switch. Multiple RSPAN source sessions anywhere in the network can contribute packets to the RSPAN session.

RSPAN destination sessions

An RSPAN destination session presents a copy of all RSPAN VLAN packets (except Layer 2 control packets) to the user for analysis.

To configure an RSPAN destination session on a destination switch, you associate the destination port with the RSPAN VLAN. The destination session collects all RSPAN VLAN traffic and sends it out the designated RSPAN destination port. An RSPAN destination session takes all packets

received on the RSPAN VLAN, strips off the VLAN tagging, and presents them on the destination port.

You can have more than one destination session active in the same Cisco compatible RSPAN VLAN. You can monitor the same RSPAN VLAN with multiple RSPAN destination sessions throughout the network. In this situation, you can regard the RSPAN VLAN ID as a network wide ID for a particular monitoring session.

When configuring an RSPAN destination session, if the destination port is not part of the RSPAN VLAN, the port is automatically moved in the RSPAN VLAN and set to untagged. If a previous VLAN configuration prevents port moving, an error message is displayed.

When an RSPAN destination interface is erased, the RSPAN port is removed from the RSPAN vlan and set to untagged state.

You can configure up to four RSPAN destination sessions on a destination switch. Each RSPAN instance holds a single destination port, meaning that you can configure up to four destination ports on a switch.

Note:

The RSPAN destination session does not occupy one of the four standard port-mirroring sessions. You can still configure up to four port-mirroring sessions on the destination switch.

RSPAN restrictions and interactions with other features

RSPAN interacts with the following features:

VLAN interactions

- You can configure up to four RSPAN VLANs on a switch.
- No MAC address learning occurs on the RSPAN VLAN, because all RSPAN VLAN traffic is always flooded.
- Mapping of an RSPAN VLAN over an SPB ISID and transport over an SPB cloud is not supported.
- You cannot:
  - remove an RSPAN destination port from the RSPAN VLAN while this port is involved in the RSPAN instance.
  - remove an RSPAN VLAN if it is used in an RSPAN instance. You must disable the RSPAN instance first.
  - change the membership of an RSPAN destination port without disabling first the instance.
  - set an SPBM B-VLAN or an spbm-switchedUni VLAN as an RSPAN VLAN.
  - set an RSPAN VLAN as a management VLAN.
  - use the same vlan or the same interface in another RSPAN instance.

Port-mirroring interactions

- You can configure up to 4 RSPAN destination sessions and up to 3 RSPAN source sessions on a switch or stack.

- Port Mirroring general limitations regarding VLAN tagging also apply to RSPAN.
- You can specify any ports within the stack as ports for RSPAN port-mirroring sessions, with the following exceptions:

  You cannot:

  - configure a port which has 802.1X enabled as an RSPAN destination port.
  - configure a port which is a member of MLT/DMLT/LAG as an RSPAN destination port.
  - configure a port as an RSPAN destination or Mirror To Port (MTP) if this port is an RSPAN source / mirrored port for another instance.
  - configure the allow-traffic option for port-mirroring along with RSPAN
- For Remote Port Mirroring with MAC base modes Asrc, Adst, AsrcBdst, AsrcBdstOrBsrcAdst, AsrcOrAdst, and port based modes XrxYtx, XrxYtxOrYrxXtx, you must install filters to enable an RSPAN source session. If platform resource limits are reached, the application may not function in these modes.
- For port based modes XrxYtx and XrxYtxOrYrxXtx, RSPAN can function only for unicast traffic.
- The RSPAN destination port is set as an untagged member of the RSPAN VLAN, to ensure that the RSPAN tag is stripped off.
- Mac-security cannot be enabled on RSPAN destination-ports, because a destination port is also a monitor port.

STP interactions

- The RSPAN destination port does not participate in STP.
- The RSPAN destination port follows the same rules as a local MTP in regard to STP and topology packets.
- Control packets are not mirrored by an RSPAN instance. The mirrored BPDUs may get mixed up with the actual BPDUs, resulting in STP loops and topology issues. Control packets are treated separately and may be discarded before reaching destination port.

MLT/LACP interactions

With RSPAN over MLT/LACP trunks you can add a monitor port that is part of a trunk to the RSPAN instance. As such, the entire trunk forwards mirrored traffic.

When an RSPAN over MLT/LACP instance is created, port-mirroring verifies if the chosen monitor port is part of a trunk and depending on the result the correct hardware configurations are made.

Port-mirroring is informed of any modification made to the trunk (ports are removed or are shutdown, or links are removed) and makes the appropriate changes. Even if the chosen monitor port goes down the remaining members of the trunk continue to forward the mirrored traffic. At reboot current configuration is expected to be reestablished.

RSPAN over MLT/LACP

RSPAN over MLT/LACP is configured similarly to an RSPAN instance except that in RSPAN over MLT/LACP the monitor port is part of a trunk.

The following figure shows how the RSPAN over MLT/LACP works with three connected devices.

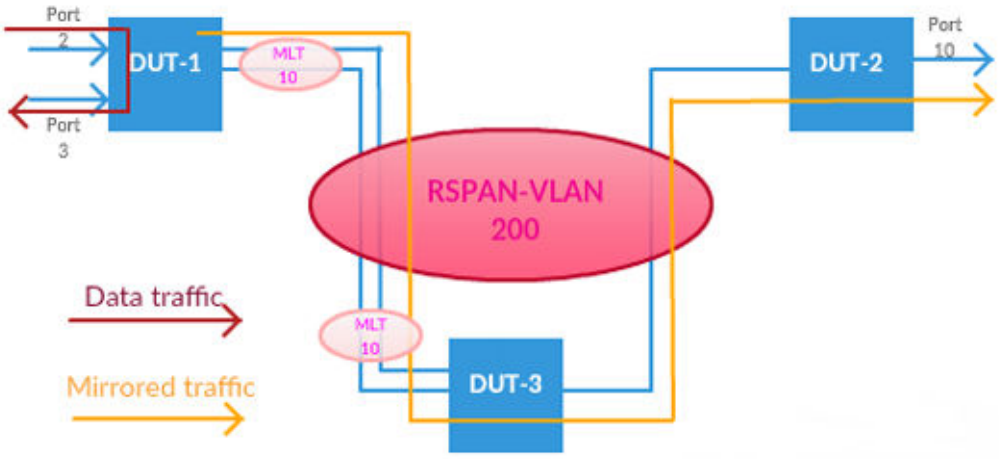

RSPAN over MLT/LACP source session

To configure an RSPAN over MLT/LACP source session on a device, you associate a port mirroring instance with an RSPAN VLAN. The trunk, which contains the chosen monitor port can be created prior to creating the mirroring instance or after. The output of this session is the stream of packets that are sent to the RSPAN VLAN. An RSPAN source session is very similar to a local port mirroring session, except for where the packet stream is directed. In an RSPAN instance, the mirrored packets are relabeled with the RSPAN VLAN ID and directed over normal trunk ports to the destination switch. When exiting the switch the traffic has both vlan labels (double tagging). Only RSPAN supports the monitor port to be part of MLT/LACP. If the created instance is not an RSPAN instance an error is returned when you try to add the monitor port to a trunk or when you create the mirroring instance if the trunk already exists.

Show Environmental

This feature provides an enhancement to display environmental information about the operation of the switch or units within a stack. The Show environmental command does not require any specific configuration, and it reports the following parameters for each switch:

- power supply status
- fan status
- switch system temperature

The Show Environmental command depends on the hardware of each unit. The command is available from any CLI mode, and you do not need to enable or activate this feature. The command displays information for a stand-alone switch and for each unit in a stack, regardless of how many units are in that stack.

You can configure the Show Environmental command in CLI, SNMP, and EDM.

The following table defines the various states of the environment of a switch.

Table 4: Environmental parameters

Measurement	State	Description
PSU1	Primary	If the power source is present and is the primary power source
PSU2	Redundant	If the power source is present and is the redundant power source
	N/A	If the power source is missing or not providing power
Fan	OK	If the fan is working properly
	FAIL	If any fan malfunction exists
	N/A	If the fan dose not exist
Temperature	OK	If temperature is lower than 65C
	HIGH	If temperature is higher than 65C but lower than 80C. **Note:** If temperature is higher than 80C, the switch shuts down.

SFP DDI information

SFP Digital Diagnostic Interface (DDI), also known as digital optical monitoring (DOM) or digital diagnostics monitoring (DDM) gives the end user the ability to monitor real-time parameters of the SFP(+), such as optical output power, optical input power, temperature, laser bias current, and transceiver supply voltage and also to receive warnings and alarms regarding current operating status of the transceiver.

Calibration and alarm/warning threshold data is written during device manufacture. There is also an operating status field that checks the alarm/warning bits and returns the adequate operating status. Measurements are calibrated over vendor specified operating temperature and voltage. Alarm and warning threshold values should be interpreted in the same manner as real time 16 bit data.

The Digital Diagnostic Interface (DDI) feature collects data and monitors alarms and warnings on all the supported SFP, SFP+, and GBIC transceivers. The DDI collects the following information:

- SFP vendor information (including type, wavelength, vendor name, vendor revision/serial, hardware options, CLEI code, and Product Code)
- DDI support information
- DDI alarm and warning threshold values
- temperature

- supply voltage
- transmit bias current
- TX/RX optical power
- transmit power
- receive power measurement
- transceiver calibration

This functionality is supported from the moment the switch finishes its initialization.

Note:

By default, logging is disabled for all ports.

MACsec statistics

MACsec is an 802.1AE IEEE 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.

Note:

If encryption is enabled, the following MACsec statistics are not incremented:

- *Octets Validated* for secure-channel inbound statistics
- *Octets Protected* for secure-channel outbound statistics

Table 5: 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 not operating in strict mode.
RxNoTagPkts	Specifies the number of received packets without the MAC security tag (SecTAG), with MACsec operating in strict mode.

Table continues…

Statistics	Description
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 not 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 6: Secure-channel inbound MACsec statistics

Statistics	Description
PortId	Specifies the port.
Late Packets	Specifies the number of packets received that have been discarded for this Secure Channel (SC) with Replay Protect enabled.
InvalidPkts	Specifies the summation of all packets received that were not valid for this SC, with MACsec operating in check mode.
Delayed Packets	Specifies the summation of packets for this SC, with the Packet Number (PN) of the packets lower than the lower bound replay protection PN.
Unchecked Packets	The total number of packets for this SC that: • were encrypted and had failed the integrity check • were not encrypted and had failed the integrity check • were received when MACsec validation was not enabled
Octets Validated	Specifies the number of octets of plaintext recovered from received packets that were integrity protected but not encrypted.
Octets Decrypted	Specifies the number of octets of plaintext recovered from received packets that were integrity protected and encrypted.

Table 7: Secure-channel outbound MACsec statistics

Statistics	Description
PortId	Specifies the port.

Table continues…

Statistics	Description
Octets Protected	Specifies the number of plain text octets that are integrity protected but not encrypted on the transmitting SC.
Octets Encrypted	Specifies the number of plain text octets that are integrity protected and encrypted on the transmitting SC.

Network monitoring configuration using CLI

This section describes the CLI commands that you use to configure network monitoring.

Viewing CPU utilization

About this task

View the CPU utilization of the switch or stack.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. View CPU utilization:

   ```
   show cpu-utilization
   ```

Example

```
Switch>enable
Switch#show cpu-utilization
-------------------------------------------------------------------
                      CPU Utilization
-------------------------------------------------------------------

Unit  Last 10 Sec, 1 Min, 10 Min, 60 Min, 24 Hrs, System Boot-Up
-------------------------------------------------------------------
1          23%     16%    15%     15%     15%     15%
```

Viewing memory utilization

About this task

View the memory utilization of the switch or stack.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View memory utilization:

```
show memory-utilization
```

Example

```
Switch>enable
Switch#show memory-utilization
----------------------------------------
          Memory Utilization
----------------------------------------
Unit  Total      Used          Free
----------------------------------------
1  256Mbytes     117Mbytes     139Mbytes
```

Viewing system logging information

About this task

Display system logging configuration information.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View system logging information:

```
show logging [config] [critical] [informational] [serious] [sort-
reverse] [unit <1-8>]
```

Example

```
Switch>enable
Switch#show logging config
Event Logging: Enabled
Volatile Logging Option: Overwrite
Event Types To Log: Critical, Serious, Informational
Event Types To Log To NV Storage: Critical, Serious
Remote Logging: Disabled
Remote Logging Address: 0.0.0.0
Secondary Remote Logging Address: 0.0.0.0
Event Types To Log Remotely: None
Facility: Daemon
```

Variable definitions

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

Variable	Value
config	Display local and remote system logging configuration status.
critical	Display critical log messages.
serious	Display serious log messages.

Table continues…

Variable	Value
informational	Display informational log messages.
sort-reverse	Display informational log messages in reverse chronological order (beginning with most recent).
unit <1-8>	Display log messages for a specific switch in a stack. **Important:** You cannot use this command variable for a standalone switch.

Configuring syslog capabilities

About this task

Display and clear the last software exception.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Display the last software exception:

   ```
   show system last-exception [unit{<1-8>|all}]
   ```

3. Clear the last software exception:

   ```
   clear last-exception [unit{<1-8>|all }]
   ```

Example

```
Switch>enable
Switch#show system last-exception
Last Saved Exception - Unit # 2
-------------------------------------
bld version:
```

Variable definitions

Use the data in the following table to use the **`show system last-exception`** command.

Variable	Value
unit <1-8> | all	The unit specified for the command. If you do not specify a unit, the last unit the command was run on is used.

Configuring system logging

About this task

Configure and manage the logging of system messages.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure system logging:

   ```
   logging [enable | disable] [level critical | serious | informational
   | none] [nv-level critical | serious | none] remote [address |
   enable | level] volatile [latch | overwrite]
   ```

Example

```
Switch>enable
Switch#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch(config)#logging enable level critical
```

Variable definitions

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

Variable	Value
enable | disable	Enables or disables the event log (enabled is the default setting).
level critical | serious | informational | none	Specifies the level of logging stored in Dynamic Random Access Memory (DRAM).
nv-level critical | serious | none	Specifies the level of logging stored in NVRAM.
remote	Configures remote logging parameters. Address: configure remote syslog address. Enable: enable remote logging. Level: configure remote logging level.
volatile	Configures options for logging to DRAM. Latch: latch DRAM log when it is full. Overwrite: overwrite DRAM log when it is full.

Disabling logging

About this task

Disable the system event log.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Disable logging:

   ```
   no logging
   ```

Default logging

About this task

Configure the system settings as the factory default settings for the system event log.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure default system settings for the system event log:

   ```
   default logging
   ```

Clearing log messages

About this task

Clear all log messages in DRAM.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Clear log messages:

   ```
   clear logging [non-volatile] [nv] [volatile]
   ```

Variable Definitions

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

Variable	Value
non-volatile	Clears log messages from NVRAM.
nv	Clears log messages from NVRAM and DRAM.
volatile	Clears log messages from DRAM.

Configuring remote system logging

About this task

Configure and manage the logging of system messages on a remote server.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure the remote system log:

   ```
   logging remote [address <A.B.C.D|WORD>] [secondary-address <A.B.C.D|
   WORD>] [enable] [level <critical|informational|none|serious>]
   [facility <daemon| local0 | local1 | local2 | local3 | local4 |
   local5 | local6 | local7>]
   ```

Variable definitions

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

Variable	Value
`address <A.B.C.D|WORD>`	Specifies the primary remote system log server IP address. • A.B.C.D—the IPv4 address of the remote server • WORD—the remote host IPv6 address. Value is a character string with a maximum of 45 characters.
`secondary-address <A.B.C.D|WORD>`	Specifies the secondary remote system log server IP address. • A.B.C.D—the IPv4 address of the remote server • WORD—the remote host IPv6 address. Value is a character string with a maximum of 45 characters.

Table continues…

Variable	Value
`enable`	Enables system message logging on the remote server. You must configure either the primary or secondary remote server IP address before you can enable remote logging.
`facility <daemon|local0|local1|local2|local3|local4|local5|local6|local7>`	Specifies remote logging facility.
`level <critical|informational|none|serious>`	Specifies the level of system messages to send to the remote system log server. • critical—only messages classified as critical are sent to the remote system log server • serious—only messages classified as serious are sent to the remote system log server • informational—only messages classified as informational are sent to the remote system log server • none—no system log messages are sent to the remote system log server

Disabling remote system logging

About this task

Disable the logging of system messages on a remote server.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Disable the remote system log:

   ```
   no logging remote [address] [secondary-address] [enable] [level]
   ```

Variable definitions

Use the data in the following table to use the **`no logging remote`** command.

Variable	Value
`address`	Clears the primary remote system log server IP address.

Table continues…

Variable	Value
`secondary-address`	Clears the secondary remote system log server IP address.
`enable`	Disables system message logging on the remote server.
`level`	Clears the remote server logging level.

Restoring remote system logging to default

About this task

Restore the logging of system messages on a remote server to factory defaults.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Disable the remote system log:

   ```
   default logging remote [address] [secondary-address] [level]
   ```

Variable definitions

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

Variable	Value
`address`	Restores the primary remote system log server IP address to the factory default (0.0.0.0).
`secondary-address`	Restores the secondary remote system log server IP address to factory the default (0.0.0.0).
`level`	Restores the remote server logging level to the factory default (none).

Displaying environmental status

Use this procedure to view the environmental status of the switch or stack.

Procedure

1. Log on to CLI to enter User EXEC mode.
2. At the command prompt, enter the following command:

   ```
   show environmental
   ```

Note:

You can use the command from Global Configuration mode or User EXEC mode.

Example

The following figure provides a sample of **show environmental** command.

```
Switch>show environmental
Unit# FAN1 FAN2 FAN3 FAN4 Temperature
----- ---- ---- ---- ---- -----------
1     OK   N/A  N/A  N/A   OK   39.5C
```

Network monitoring configuration using Enterprise Device Manager

This section provides the procedures to configure network monitoring using Enterprise Device Manager (EDM).

Prerequisites

- Open one of the supported browsers.
- Enter the IP address of the switch to open an EDM session.

Viewing CPU and memory utilization using EDM

Use the following procedure to view both CPU and memory utilization.

Prerequisites

- Open one of the supported browsers.
- Enter the IP address of the switch to open an EDM session.

Procedure steps

1. From the navigation tree, double-click **Edit**.
2. In the Edit tree, double-click **Chassis**.
3. In the Chassis tree, double-click **Chassis**.
4. In the work area, click the **CPU/Mem Utilization** tab.
5. On the tool bar, click **Refresh** to update the data.

Variable definitions

The following table describes the fields on the CPU/Mem Utilization tab.

Variable	Value
Unit	Indicates the numerical representation of the unit.
Last10Seconds	Indicates the CPU usage, in percentage, for the last 10 seconds.
Last1Minute	Indicates the CPU usage, in percentage, for the last minute.
Last10Minutes	Indicates the CPU usage, in percentage, for the last 10 minutes.
Last1Hour	Indicates the CPU usage, in percentage, for the last hour.
Last24Hours	Indicates the CPU usage, in percentage, for the last 24 hours.
TotalCPUUsage	Indicates the CPU usage in percentage, since system start up.
MemoryTotalMB	Indicates the total memory present, in megabytes, on the unit.
MemoryAvailableMB	Indicates the memory remaining available on the unit.
MemoryUsedMB	Indicates the memory being used on the unit.

Switch stack information management using EDM

Use the information in the following sections to display and edit switch stack information.

Viewing stack information using EDM

Use this procedure to display information about the operating status of stack switches.

Procedure steps

1. From the navigation tree, double-click **Edit**.
2. In the Edit tree, double-click **Chassis**.
3. In the Chassis tree, double-click **Switch/Stack**.
4. On the work area, click the **Stack Info** tab.

Variable Definitions

Use the information in the following table to help you understand the stack information display.

Variable	Value
Unit	Indicates the unit number and the value range is from 1 to 8. For a stack environment, a **Unit** value of 1 specifies the base unit. For a standalone switch, the **Unit** value is 1.
Description	Describes the component or subcomponent. If not available, the value is a zero length string.
Location	Indicates the geographic location of a component in a system modeled as a chassis, but possibly physically implemented with geographically separate devices connected to exchange management information. Chassis modeled in this manner are

Table continues…

Variable	Value
	sometimes referred to as virtual chassis. An example value is: **4th flr wiring closet in blg A**. **Important:** This field applies only to components that are in either the Board or Unit groups. If the information is unavailable, for example, the chassis is not modeling a virtual chassis or component is not in a Board or Unit group, the value is a zero-length string. If this field is applicable and is not assigned a value through a SNMP SET PDU when the row is created, the value defaults to the value of the object s5ChasComSerNum.
LstChng	Indicates the value of sysUpTime when it was detected that the component or sub-component was added to the chassis. If this action has not occurred since the cold or warm start of the agent, the value is zero.
AdminState	Indicates the state of the component or subcomponent. • enable: enables operation • reset: resets component
OperState	Indicates the current operational state of the component. The possible values are • other: another state • notAvail: state not available • removed: component removed • disabled: operation disabled • normal: normal operation • resetInProg: reset in progress • testing: performing a self test • warning: operating at warning level • nonFatalErr: operating at error level • fatalErr: error stopped operation The component type determines the allowable (and meaningful) values.
IpAddress	Indicates the IP address of the component or subcomponent.
Ipv6Address	Indicates the IPv6 address of the component or subcomponent.
Ipv6NetMask	Indicates the IPv6 Netmask address of the component or subcomponent.
Ver	Indicates the version number of the component or subcomponent. If not available, the value is a zero-length string.

Table continues…

Variable	Value
SerNum	Indicates the serial number of the component or subcomponent. If not available, the value is a zero-length string.
BaseNumPorts	Indicates the number of base ports of the component or subcomponent.
TotalNumPorts	Indicates the number of ports of the component or subcomponent.
RunningSoftwareVer	Indicates the software version running on the switch.

Editing stack information using EDM

Use this procedure to change the information about the switch units in the stack.

Procedure steps

1. From the navigation tree, double-click **Edit**.
2. In the Edit tree, double-click **Chassis**.
3. In the Chassis tree, double-click **Switch/Stack**.
4. In the work area, click the **Stack info** tab.
5. To select a switch unit for which to edit information, click a switch row.
6. In the row, double-click the cell in the **Location** column.
7. Type a location.
8. In the row, double-click the cell in the **AdminState** column.
9. Select a value from the list.
10. On the toolbar, click **Apply**.

Variable Definitions

Use the information in the following table to help you understand the stack information display.

Variable	Value
Unit	Indicates the unit number and the value range is from 1 to 8. For a stack environment, a **Unit** value of 1 specifies the base unit. For a standalone switch, the **Unit** value is 1.
Description	Describes the component or subcomponent. If not available, the value is a zero length string.
Location	Indicates the geographic location of a component in a system modeled as a chassis, but possibly physically implemented with geographically separate devices connected to exchange management information. Chassis modeled in this manner are sometimes referred to as virtual chassis. An example value is: **4th flr wiring closet in blg A**.

Table continues…

Variable	Value
	Important: This field applies only to components that are in either the Board or Unit groups. If the information is unavailable, for example, the chassis is not modeling a virtual chassis or component is not in a Board or Unit group, the value is a zero-length string. If this field is applicable and is not assigned a value through a SNMP SET PDU when the row is created, the value defaults to the value of the object s5ChasComSerNum.
LstChng	Indicates the value of sysUpTime when it was detected that the component or sub-component was added to the chassis. If this action has not occurred since the cold or warm start of the agent, the value is zero.
AdminState	Indicates the state of the component or subcomponent. • enable: enables operation • reset: resets component
OperState	Indicates the current operational state of the component. The possible values are • other: another state • notAvail: state not available • removed: component removed • disabled: operation disabled • normal: normal operation • resetInProg: reset in progress • testing: performing a self test • warning: operating at warning level • nonFatalErr: operating at error level • fatalErr: error stopped operation The component type determines the allowable (and meaningful) values.
IpAddress	Indicates the IP address of the component or subcomponent.
Ipv6Address	Indicates the IPv6 address of the component or subcomponent.
Ipv6NetMask	Indicates the IPv6 Netmask address of the component or subcomponent.
Ver	Indicates the version number of the component or subcomponent. If not available, the value is a zero-length string.
SerNum	Indicates the serial number of the component or subcomponent. If not available, the value is a zero-length string.

Table continues…

Variable	Value
BaseNumPorts	Indicates the number of base ports of the component or subcomponent.
TotalNumPorts	Indicates the number of ports of the component or subcomponent.
RunningSoftwareVer	Indicates the software version running on the switch.

Viewing pluggable ports using EDM

Use this procedure to display pluggable port information.

Procedure steps

1. From the navigation tree, double-click **Edit**.
2. In the Edit tree, double-click **Chassis**.
3. In the Chassis tree, double-click **Switch/Stack**.
4. In the work area, click the **Pluggable Ports** tab to display the current stack information.
5. To configure pluggable ports for a switch, double-click table cells as required.
6. On the toolbar, click **Apply**.
7. On the toolbar, you can click **Refresh** to verify the port configuration.

Variable definitions

Use the data in the following table to help you understand the pluggable ports display.

Variable	Value
IfIndex	Indicates the interface index.
GbicType	Indicates the type of SFP or SFP+ connector
Wavelength	Indicates the wavelength in nm of the SFP or SFP+.
VendorName	Indicates the name of the SFP or SFP+ manufacturer.
VendorOui	Indicates the vendor ID of the SFP or SFP+ manufacturer.
VendorPartNo	Indicates the model of the SFP or SFP+.
VendorRevision	Indicates the manufacturer revision level for the SFP or SFP+.
VendorSerial	Indicates the manufacturer serial number for the SFP or SFP+.
HwOptions	Indicates hardware options set for the SFP or SFP+.
DateCode	Indicates the manufacturer date code for the SFP or SFP+.

Table continues…

Variable	Value
CleiCode	Indicates the Telcordia register assignment CLEI code.
ProductCode	Indicates the PEC code of the device.

Viewing stack health using EDM

Use this procedure to display stack health information.

Procedure steps

1. From the navigation tree, double-click **Edit**.
2. In the Edit tree, double-click **Chassis**.
3. In the Chassis tree, double-click **Switch/Stack**.
4. In the work area, click the **Stack Health** tab to display the stack health.

Variable definitions

Use the data in the following table to help you understand the stack health.

Variable	Value
Switch Units Found	Indicates the number of switch units in the stack.
Stack Health Check	Indicates the stack health.
Stack Diagnosis	Indicates the stack mode.
Unit	Indicates the unit number.
Description	Describes each unit in the stack.
Cascade Up	Indicates the cascade up link status.
Cascade Down	Indicates the cascade down link status.
Stack Role	Indicates which unit is the base unit.

Configuring the system log using EDM

Use the following procedure to configure and manage the logging of system messages.

Prerequisites

- Open one of the supported browsers.
- Enter the IP address of the switch to open an EDM session.

Procedure Steps

1. From the navigation tree, double-click **Edit**.

2. In the Edit tree, double-click **Diagnostics**.
3. In the Diagnostics tree, double-click **System Log**.
4. In the work area, click the **System Log Settings** tab.
5. Choose the operation in the **Operation** field.
6. Choose the buffer space allocation in the **BufferFullAction** field.
7. Choose the type of system messages to save in volatile memory in the **SaveTargets**field.
8. Choose the type of system messages to save in non-Volatile memory in the **SaveTargets** field.
9. Choose the types of system log messages to delete from volatile and non-volatile memory in the **ClearMessageBuffers** field.
10. On the tool bar, Click **Apply**.

Variable definitions

Use the data in the following table to configure the system log.

Variable	Value
Operation	Enables (on) or disables (off) the system log.
BufferFullAction	Specifies the action for the system to take when the buffer space allocated for system log messages is exhausted. • overwrite—previously logged messages are overwritten • latch—halts the saving of system log messages until overwrite is selected, or buffer space is made available by other means (for example, clearing the buffer).
Volatile - CurSize	Indicates the number of messages currently stored in volatile memory.
Volatile - SaveTargets	Specifies the type of system messages to save in volatile memory. • critical—only messages classified as critical are saved in volatile memory • critical/serious—only messages classified as critical and serious are saved in volatile memory • critical/serious/inform—only messages classified as critical, serious, and informational are saved in volatile memory • none—no system log messages are saved in volatile memory
non-Volatile - CurSize	Indicates the number of messages currently stored in non-volatile memory.

Table continues…

Variable	Value
non-Volatile - SaveTargets	Specifies the type of system messages to save in non-volatile memory. • critical—only messages classified as critical are saved in volatile memory • critical/serious—only messages classified as critical and serious are saved in non-volatile memory • critical/serious/inform—only messages classified as critical, serious, and informational are saved in non-volatile memory • none—no system log messages are saved in volatile memory
ClearMessageBuffers	Specifies the types system log messages to delete from volatile and non-volatile memory. • volCritical—only messages classified as critical are deleted from volatile memory • volSerious—only messages classified as serious are deleted from volatile memory • volInformational—only messages classified as informational are deleted from volatile memory • nonVolCritical—only messages classified as critical are deleted from non-volatile memory • nonVolSerious—only messages classified as serious are deleted from non-volatile memory
CLI Audit Log	Enables or disables the CLI audit log. By default, the audit log is enabled.

Configuring remote system logging using EDM

Use this procedure to configure and manage the logging of system messages on a secondary, remote syslog server.

Prerequisites

- Open one of the supported browsers.
- Enter the IP address of the switch to open an EDM session.

Procedure Steps

1. From the navigation tree, double-click **Edit**.
2. In the Edit tree, double-click **Diagnostics**.
3. In the Diagnostics tree, double-click **System Log**.
4. In the work area, click the **Remote System Log** tab.

5. Choose the type of IP address of the remote system log server in the **RemoteSyslogAddressType** field.
6. In the **RemoteSyslogAddress** box, enter a IP address of the remote system log server to send system log messages.
7. Choose the type of IP address of the secondary remote system log server in the **SecondarySyslogAddressType** field.
8. In the **SecondarySyslogAddress** box, enter a IP address of the secondary remote system log server to send system log messages.
9. Choose the **Enabled** checkbox to enable remote system logging.

   OR

   Clear the **Enabled** checkbox to disable remote system logging.
10. In the**Save Targets** section, click the type of system messages.
11. In the **Facility** section, click the type of facility required.
12. On the toolbar, click **Apply**.

Variable definitions

Use the data in the following table to configure the remote system log.

Variable	Value
RemoteSyslogAddressType	Specifies the type of IP address of the remote system log server.
RemoteSyslogAddress	Specifies the IP address of the remote system log server to send system log messages to.
SecondarySyslogAddressType	Specifies the type of IP address of the secondary remote system log server.
SecondarySyslogAddress	Specifies the IP address of the secondary remote system log server to send system log messages to.
Enabled	Enables or disables the remote logging of system messages.
SaveTargets	Specifies the type of system messages to send to the remote system log server. • critical—only messages classified as critical are sent to the remote system log server • critical/serious—only messages classified as critical and serious are sent to the remote system log server • critical/serious/inform—only messages classified as critical, serious, and informational are sent to the remote system log server • none—no system log messages are sent to the remote system log server

Table continues…

Variable	Value
Facility	Specifies the remote logging facility. • Daemon • Local0 • Local1 • Local2 • Local3 • Local4 • Local5 • Local6 • Local7 DEFAULT: Daemon

Viewing system logs using EDM

Use the following procedure to display system log information.

Prerequisites

- Open one of the supported browsers.
- Enter the IP address of the switch to open an EDM session.

Procedure Steps

1. From the navigation tree, double-click **Edit**.
2. In the Edit tree, double-click **Diagnostics**.
3. In the Diagnostics tree, double-click **System Log**.
4. In the work area, click the **System Logs** tab.

Variable definitions

Use the data in the following table to help you understand the system log display.

Variable	Value
OrigUnitNumber	Indicates the slot or unit number of the originator of a log message.
MsgTime	Indicates the time (in one hundredths of a second) between system initialization and the appearance of a log message in the system log.

Table continues…

Variable	Value
MsgIndex	Indicates a sequential number the system assigns to a log message when it enters the system log.
MsgSrc	Indicates whether a log message was loaded from non-volatile memory at system initialization or was generated since system initialization.
MsgType	Indicates the type of message: Critical, Serious, or Information.
MsgString	Indicates the log message originator and the reason the log message was generated.

EDM MIB Web page

Use the information in this section to use the EDM MIB Web page to monitor network SNMP characteristics.

Using the EDM MIB Web page for SNMP Get and Get-Next

You can use the EDM Management Information Base (MIB) Web page to view the response of an SNMP Get and Get-Next request for any Object Identifier (OID).

Procedure steps

1. From the navigation tree, double-click **Administration**.
2. In the Administration tree, double-click **MIB Web Page**.
3. In the **MIB Name/ OID** box, enter the object name or OID.
4. Click **Get**.

   The result of the request appears in the Result area of the window. If the request is unsuccessful, a description of the received error appears.
5. Click **Get Next** to retrieve the information of the next object in the MIB.
6. Repeat step 3 as required.

Using the EDM MIB Web page for SNMP walk

You can use SNMP walk to retrieve a subtree of the MIB that has the SNMP object as root.

Perform this procedure to request the result of MIB Walk.

Procedure steps

1. From the navigation tree, double-click **Administration**.
2. In the Administration tree, double-click **MIB Web Page**.
3. In the **MIB Name/ OID** box, enter the object name or OID.
4. Click **Walk**.

   The result of the request appears in the Result area. If the request is unsuccessful, a description of the received error appears.

sFlow configuration using EDM

This section describes sFlow configuration.

Enabling sFlow globally

Use the following procedure to enable sFlow globally for a switch or stack.

Procedure

1. From the navigation pane, double-click **Serviceability**.
2. In the **Serviceability** tree, click **sFlow**.
3. In the sFlow work area, click the **Global** tab.
4. In the sFlow work area, click **Enable** check box to enable sFlow globally.
5. On the tool bar, click **Apply**.
6. On the tool bar, click **Refresh** to verify sFlow global configuration.

Configuring sFlow collectors

Use this procedure to configure sFlow collectors for a switch or stack.

Note:

UDP port cannot be changed from EDM for an existing collector. If UDP port for a collector is needed, it must be configured while creating the collector.

Procedure

1. From the navigation pane, double-click **Serviceability**.
2. In the **Serviceability** tree, click **sFlow**.
3. In the sFlow work area, click the **Collectors** tab.
4. Select and configure Collectors parameters as required.
5. On the toolbar, click **Apply**.
6. On the toolbar, you can click **Refresh** to verify the sFlow configuration.

Configuring sFlow interfaces

Use the following procedure to configure sFlow interfaces for a switch or stack.

Procedure

1. Follow one of the following paths:
   - From the **Device Physical View**, use Ctrl-click to select more than one port, right-click **Edit**, then click the **sFlow** tab.
   - From the **Device Physical View**, use Ctrl-click to select more than one port, then follow the navigation tree to **Edit > Chassis > Ports > sFlow** tab.

- From the navigation tree, select **Serviceability > sFlow > Interfaces** tab.

2. Configure the parameters as required in the port row.
3. Optionally, to configure parameters for multiple ports, you can use the Multiple Port Configuration section as below.
4. In the work area, in the Make Selection section of the Multiple Port Configuration pane, click the Switch/Stack/Ports ellipsis (...) to open the Port Editor dialog. If there is no Switch/Stack/Ports selection and you have already selected ports from the **Device Physical View**, proceed to the next step.

   a. In the Port Editor window, click the ports you want to configure. If you want to configure all ports, click **All**.

   b. Click **OK** to return to the Make Selection pane.

   The ports you selected appear in the Switch/Stack/Ports box.
5. To change the configuration of the selected ports, in the Multiple Port Configuration pane, double-click the cell beneath the column heading that represents the parameter you want to change and do one of the following:

   - If applicable, select a value from a drop-down list.
   - Otherwise, type a value in the cell.
6. In the Make Selection pane, click **Apply Selection**.

   The changes appear in the table.
7. **(Optional)** Click **Clear Selection** to clear Multiple Port Configurations or click **Hide Non-Editable** to display only those parameters that are editable in the Multiple Port Configuration pane for the selected ports.

Chapter 4: System diagnostics and statistics

This chapter provides procedures to configure system diagnotics and statistics using Command Line Interface (CLI) and Enterprise Device Manager (EDM).

System diagnostics and statistics using CLI

This section provides procedures to perform system diagnostics and gather statistics using CLI.

Configuring diagnostics quick mode

The diagnostics quick mode flag enables you to choose the diagnostic test behavior during boot. You can enable quick mode boot tests or all the diagnostic tests. The impact to boot time is 15 to 20 seconds when all diagnostic tests run during startup.

The diagnostic quick mode is disabled by default.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. To enable the diagnostics quick mode flag, enter the following command at the command prompt:

   ```
   diagnostics-quick-mode enable
   ```

3. To disable the diagnostics quick mode flag, enter either of the following commands at the command prompt:

   ```
   no diagnostics-quick-mode enable
   ```

   OR

   ```
   default diagnostics-quick-mode
   ```

4. To display the configuration, enter the following command at the command prompt:

```
show diagnostics-quick-mode
```

Example

Enable diagnostics quick mode:

```
Switch>enable
Switch#configure terminal
Switch(config)#diagnostics-quick-mode enable
Switch(config)#show diagnostics-quick-mode
2013-10-02 08:53:27 GMT+00:00

Diagnostics quick mode:  Enabled
```

Trace diagnosis of problems

The following sections describe how to use trace to diagnose problems.

Using trace to diagnose problems

About this task

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

Caution:

Risk of traffic loss

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

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. **(Optional)** Display the available module or submodule options:

   ```
   trace module hint
   ```

   OR

   ```
   trace module {ospf | igmp | pim | rip |eap | ntp | sflow | dvmrp}
   submodule hint
   ```

3. Set the trace level:

   ```
   trace module {ospf | igmp | pim | rip |eap | ntp | sflow | dvmrp}
   [submodule <submodule_ID>] level {critical | error | warning | info
   | debug | no-display} [unit <1-8>]
   ```

4. Enable the trace screen:

   ```
   trace screen enable
   ```

5. Disable the trace screen:

   ```
   trace screen disable
   ```

6. Disable the trace:

   ```
   trace shutdown
   ```

Variable definitions

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

Variable	Value
`module`	Sets the trace module: • OSPF • IGMP • PIM • RIP • EAP • NTP • SFLOW • DVMRP
level	Sets the trace level: • critical—displays only critical level errors • error—displays critical and error levels • warning—displays errors from critical to warning • info—displays informational errors • debug— displays all errors • no-display—disables error displaying
`screen <enable | disable>`	Enables or disables the trace screen. You can use this command to control the trace output to the console. The default is disable.
`shutdown`	Disables the trace. Shutdown sets all the modules level to *no-display*, and produces a "NO_DISPLAY" message.

Viewing the trace level

About this task

Use this procedure to view the trace level information for the modules.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. DIsplay the trace level:

```
show trace level
```

Example

```
Switch#show trace level
SERIAL:
Module      | Submodule   | Trace level
-------------------------------------------------
Default Trace Level: NO-DISPLAY
```

Variable definitions

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

Variable	Value
Module	Indicates the Trace module.
Submodule	Indicates the Trace submodule.
Trace level	Indicates the trace level.

Viewing the trace mode ID list

About this task

Use this procedure to view the supported module list for the trace feature.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. DIsplay the trace mode ID list:

```
show trace list
```

Example

```
Switch>enable
Switch#show trace list
Name
----------
OSPF
IGMP
PIM
RIP
EAP
NTP
SFLOW
DVMRP
```

Variable definitions

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

Variable	Value
Name	Indicates the name of the mode.

Viewing port-statistics

About this task

Use this procedure to view the statistics for the port on both received and transmitted traffic.

Procedure

1. Enter Global Configuration mode:

   ```
   enable

   configure terminal
   ```

2. Display port statistics:

   ```
   show port-statistics [port <portlist>]
   ```

Example

```
Switch>enable
Switch#show port-statistics
Port: 1
--------
Received
    Packets:                6578385
    Multicasts:             5625998
    Broadcasts:             623614
    Total Octets:           459044352
    MTU Exceeded:           0
    FCS Errors:             0
    Undersized Packets:     0
    Oversized Packets:      0
    Filtered Packets:       14
    Pause Frames:           0
Transmitted
    Packets:                438276
    Multicasts:             415540
    Broadcasts:             168
    Total Octets:           49355869
    Collisions:             0
    Single Collisions:      0
    Multiple Collisions:    0
    Excessive Collisions:   0
----More (q=Quit, space/return=Continue)----
```

Variable Definitions

Use the data in the following table to use the **`show port-statistics`** command.

Variable	Value
port <portlist>	The ports to display statistics for. When no port list is specified, all ports are shown.

Configuring Stack Monitor

The following CLI commands are used to configure the Stack Monitor.

Viewing the stack-monitor

About this task

Use this procedure to display the status of the Stack Monitor.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. View the stack monitor status:

   ```
   show stack-monitor
   ```

Example

```
Switch>enable
Switch#show stack-monitor
Status: disabled
Stack size: 2
Trap interval: 60
```

Configuring the stack-monitor

About this task

Use this procedure to configure the Stack Monitor.

Important:

If you do not specify a parameter for this command, all Stack Monitor parameters are set to the default values.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure Stack Monitor:

   ```
   stack-monitor [enable] [stack-size <2-8>] [trap-interval <30-300>
   ```

Variable Definitions

Use the data in the following table to use the **`stack-monitor`** command.

Variable	Value
enable	Enables stack monitoring.
stack-size <2-8>	Sets the size of the stack to monitor. Valid range is from 2–8. By default the stack size is 2.
trap-interval <30-300>	Sets the interval between traps, in seconds. Valid range is from 30 to 300 seconds. By default the trap-interval is 60 seconds.

Setting default stack-monitor values

About this task

Use this procedure to set the Stack Monitor parameters to the default values.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Set Stack Monitor parameters to default:

   ```
   default stack-monitor
   ```

Disabling the stack monitor

About this task

Use this procedure to disable the stack monitor.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Disable Stack Monitor:

   ```
   no stack monitor
   ```

Configure Stack Health Monitoring and Recovery

Use the following procedures to configure Stack Health Monitoring and Recovery.

Rebooting stack units on failure

About this task

Use this procedure to reboot stack units when the system detects failure of stacking.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure the system to reboot failed stacking units:

```
stack reboot-on-failure
```

Displaying the status of stack reboot on failure

About this task

Use this procedure to display the status of rebooting of stack units on failure.

Note:

By default, stack reboot-on-failure is enabled on the switch.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display the status of stack reboot-on-failure:

```
show stack reboot-on-failure
```

Example

```
Switch>enable
Switch#show stack reboot-on-failure
Stack Reboot on Failure:  Enabled
```

Disabling stack reboot on failure

About this task

Use this procedure to disable stack reboot on failure.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Disable stack reboot on failure:

```
no stack reboot-on-failure
```

Configuring stack retry count

About this task

Use this procedure to configure the number of times the system attempts to reach a unit before it indicates that the unit is down.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure stack retry count:

```
Stack retry-count [retry-count]
```

Variable definitions

Use the data in the following table to use the **`stack retry-count`** command.

Variable	Value
retry count	Sets the retry count for the stack. The retry count is a value in a range from 0 to 4,294,967,295. Default value: 0 **Note:** To use the command, you must enter a value.

Displaying stack retry count

About this task

Use this procedure to display the stack retry count value.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display stack retry count:

```
show stack retry-count
```

Displaying stack health

About this task

Use this procedure to display stack health information.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display stack health:

```
show stack health
```

Example

The following figure is an example of the show stack health command output when the stack is formed but the initialization process is not complete.

```
#show stack health

--------------------------------------------------------------------------
UNIT#       Switch Model        Cascade Up      Cascade Down
--------------------------------------------------------------------------
1 (Base)     <Switch#>             0K           0K
2            <Switch#>             0K           0K
3            <Switch#>             0K           0K
4            <Switch#>             0K           0K
5            <Switch#>             0K           0K
6            <Switch#>             0K           0K
7            <Switch#>             0K           0K
8            <Switch#>             0K           0K

--------------------------------------------------------------------------

Switch Units Found = 8
Stack Health Check = OK - RESILIENT
Stack Diagnosis = Stack in full resilient mode.
```

Where, <Switch#> is the switch model.

The following figure is an example of the show stack health command output when the stack is formed and initialized and there are damaged/missing rear links.

```
#show stack health

--------------------------------------------------------------------------
UNIT#       Switch Model        Cascade Up            Cascade Down
--------------------------------------------------------------------------
1 (Base)     <Switch#>             0K                    0K
2            <Switch#>             0K                    0K
3            <Switch#>             0K                    0K
4            <Switch#>             0K             LINK DOWN or MISSING
6            <Switch#>      LINK DOWN or MISSING          0K
7            <Switch#>              0K                    0K
8            <Switch#>              0K                    0K

--------------------------------------------------------------------------
Switch Units Found = 7
Stack Health Check = WARNING - NON-RESILIENT
Stack Diagnosis = Stack in non-resilient mode.
Recommend to add/replace the identified cable(s).
```

Where, <Switch#> is the switch model.

The following figure is an example of the show stack health command output when the stack is formed and some of the rear ports are not functioning properly.

```
#show stack health

--------------------------------------------------------------------------
UNIT#       Switch Model       Cascade Up       Cascade Down
--------------------------------------------------------------------------
1 (Base)     <Switch#>            0K            0K
2            <Switch#>            0K            0K
3            <Switch#>            0K            0K
4            <Switch#>            0K            0K
5            <Switch#>            0K            0K
6            <Switch#>            0K       UP WITH ERRORS
```

```
7            <Switch#>            UP WITH ERRORS  0K
8            <Switch#>            0K            0K

--------------------------------------------------------------------------------
Switch Units Found = 8
Stack Health Check = WARNING - NON-RESILIENT
Stack Diagnosis = Stack in non-resilient mode.
Recommend to add/replace the identified cable(s).
```

Where, <Switch#> is the switch model.

The following figure is an example of the show stack health command output when the stack is running with a temporary base

```
Switch#show stack health

--------------------------------------------------------------------------------
UNIT#             Switch Model        Cascade Up      Cascade Down
--------------------------------------------------------------------------------
1                  <Switch#>             0K           0K
2 (Temporary Base) <Switch#>             0K           0K
3                  <Switch#>             0K           0K
4                  <Switch#>             0K           0K
5                  <Switch#>             0K           0K
6                  <Switch#>             0K           0K
7                  <Switch#>             0K           0K
8                  <Switch#>             0K           0K

--------------------------------------------------------------------------------
Switch Units Found = 8
Stack Health Check = OK - RESILIENT
Stack Diagnosis = Stack in full resilient mode.
```

Where, <Switch#> is the switch model.

Viewing Stack Port Counters

About this task

Use this procedure to configure the stack port counters.

Important:

The stack counters measure the size of packets received on HiGig ports. The size of these packets is greater than the size of the packets received on front panel ports since ASIC HiGig+ header is added to each of them. The size of this header is 12 bytes, therefore another range of stack counters is incremented when sending packets having length close to the stack counters upper intervals limit.

Important:

The number of received/transmitted packets can be greater than the number of packets transmitted on front panel ports since there are different stack management packets transmitted/received.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display stacking statistics:

```
show stack port-statistics [unit <1-8>]
```

Example

```
Switch>enable
Switch#show stack port-statistics
                            UP                    DOWN
--------------------------  --------------------  --------------------
Received
    Packets:                0                     0
    Multicasts:             0                     0
    Broadcasts:             0                     0
    Total Octets:           0                     0
    Packets 64 bytes:       0                     0
            65-127 bytes:   0                     0
            128-255 bytes:  0                     0
            256-511 bytes:  0                     0
            512-1023 bytes: 0                     0
            1024-1518 bytes: 0                    0
            1519-2047 bytes: 0                    0
            2048-4095 bytes: 0                    0
            4096-9216 bytes: 0                    0
                 Jumbo     : 0                    0
    Control Packets:        0                     0
    FCS Errors:             0                     0
    Undersized Packets:     0                     0
    Oversized Packets:      0                     0
    Filtered Packets:       0                     0
    Pause Frames:           0                     0
    PFC Frames:             0                     0
Transmitted
```

Variable Definitions

Use the data in the following table to use the **`show stack port-statistics`** command.

Variable	Value
unit <1-8>	Specifies the unit in the stack.

Clearing stack port counters

About this task

Use the following procedure to clear the stack port counters.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Clear stacking statistics:

```
clear stack port-statistics [unit <1-8>]
```

Variable Definitions

Use the data in the following table to use the **clear stack port-statistics** command.

Variable	Value
unit <1-8>	Specifies the unit in the stack.

Using the stack loopback test

About this task

Use this procedure to complete a stack loopback test.

Note:

Stack Reboot on Failure must be disabled before running this test. To disable rebooting of stack units on failure see Disabling stack reboot on failure on page 64

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Launch the internal loopback test for the stack ports:

   ```
   stack loopback-test internal
   ```

3. Launch the external loopback test for the stack ports:

   ```
   stack loopback-test external
   ```

Example

```
Switch>enable
Switch#stack loopback-test internal
Testing uplink port ... Ok
Testing downlink port ... Ok
Internal loopback test PASSED
```

Next steps

If a problem exists with a unit's stack port or a stack cable, an internal loopback test using the **stack loopback-test internal** command is performed. If the test displays an error then the stack port is damaged.

If the internal test passes, the external test can be run using the **stack loopback-test external** command. If the test displays an error then the stack cable is damaged.

The output of the **stack loopback-test** commands are as follows:

```
Switch#stack loopback-test internal
 Testing uplink port … ok
 Testing downlink port … ok
 Internal loopback test PASSED.
 Switch#stack loopback-test external
```

```
External loopback test PASSED.
```

If one of the stack ports is defective (for example, such as the uplink), the output of the internal loopback test is as follows:

```
Switch#stack loopback-test internal
Testing uplink port … Failed
Testing downlink port … ok
Internal loopback test FAILED.
```

If both the stack ports are functional, but the stack cable is defective, the external loopback test detects this, and the output is as follows:

```
Switch#stack loopback-test external
External loopback test FAILED. Your stack cable might be damaged.
```

If you run the command on any unit of a stack, you see the following error message:

```
Switch#stack loopback-test internal
Stack loopback test affects the functioning of the stack.
You should run this in stand-alone mode
Switch#stack loopback-test external
Stack loopback test affects the functioning of the stack.
You should run this in stand-alone mode
```

Displaying port operational status

About this task

Use this procedure to display the port operational status.

Tip:

If you use a terminal with a width of greater than 80 characters, the system displays the output in a tabular format.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Display port operational status:

   ```
   show interfaces [port list] verbose
   ```

   Tip:

   If you issue the command with no parameters, the system displays the port status for all ports.

Example

```
Switch>enable
Switch#show interface verbose
Port:  1
    Trunk:
    Admin Status:  Enable
    Oper Status:  Up
    EAP Oper Status:  Up
    VLACP Oper Status:  Down
    STP Oper Status:  Forwarding
    Link: Up
    Last Change: 4 day(s), 07h:23m:32s ago
    Link Autonegotiation:  Enabled
    Link Speed:  100Mbps
    Link Duplex:  Full-Duplex
    Flow Control:  Disable
    Energy Saver: Disabled
    Energy Saver Oper Status: No Power Saving
    BPDU-guard (BPDU Filtering): Disabled
    BPDU-guard (BPDU Filtering) Oper Status: N/A
    SLPP-guard: Disabled
    SLPP-guard Oper Status: N/A
Port:  2
    Trunk:
    Admin Status:  Enable
    Oper Status:  Down
----More (q=Quit, space/return=Continue)----
```

Validating port operational status

Before you begin

- Using ACI, configure the EAP status for some ports as unauthorized
- Configure VLACP on port 1 from one switch and on port 2 on another switch. Create a link between these 2 ports.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Verify EAP port operational status:

   ```
   show interfaces
   ```

3. Verify VLACP port operational status:

   ```
   show interfaces
   ```

   The VLACP status is UP for the port where you entered the command. When you disconnect the link from the other switch, the system displays the VLACP status as Down.

4. After the switch boots, verify STP port operational status:

   ```
   show interfaces
   ```

The system displays STP Status as Listening. After a brief interval, the system displays the STP status as Learning. After the forward delay interval elapses, enter `show interfaces`. The system displays the STP status as Forwarding.

Example

```
Switch#show interfaces
              Status           Auto                          Flow
Port Trunk Admin   Oper  Link Negotiation  Speed   Duplex Control
---- ----- ------- ---- ---- ----------- -------- ------ -------
1          Enable  Up   Up   Enabled     100Mbps  Full   Disable
2          Enable  Down Down Enabled
3          Enable  Down Down Enabled
4          Enable  Down Down Enabled
5          Enable  Down Down Enabled
6          Enable  Down Down Enabled
7          Enable  Down Down Enabled
8          Enable  Down Down Enabled
9          Enable  Down Down Enabled
10         Enable  Down Down Enabled
11         Enable  Down Down Enabled
12         Enable  Down Down Enabled
13         Enable  Down Down Enabled
14         Enable  Down Down Enabled
15         Enable  Down Down Enabled
16         Enable  Down Down Enabled
17         Enable  Down Down Enabled
18         Enable  Down Down Enabled
19         Enable  Down Down Enabled
----More (q=Quit, space/return=Continue)----
```

Showing port information

About this task

Display configuration information for a specific port in one command. The config keyword displays information specific to the port configuration.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Display port information:

   ```
   show interfaces <portlist> config
   ```

Example

```
Switch>enable
Switch#show interfaces all config
Port:  1
    Trunk:
    Admin Status:  Enable
    Oper Status:  Up
    EAP Oper Status:  Up
    VLACP Oper Status:  Down
    STP Oper Status:  Forwarding
    Link: Up
    Last Change: 4 day(s), 07h:23m:23s ago
```

```
    Link Autonegotiation:  Enabled
    Link Speed:  100Mbps
    Link Duplex:  Full-Duplex
    Flow Control:  Disable
    Energy Saver: Disabled
    Energy Saver Oper Status: No Power Saving
    BPDU-guard (BPDU Filtering): Disabled
    BPDU-guard (BPDU Filtering) Oper Status: N/A
    SLPP-guard: Disabled
    SLPP-guard Oper Status: N/A
Port:  2
    Trunk:
    Admin Status:  Enable
    Oper Status:  Down
----More (q=Quit, space/return=Continue)----
```

Table 8: VLAN interfaces configuration

Unit/Port	Filter Untagged Frames	Filter Unregistered Frames	PVID	PRI	Tagging	Name
1/1	No	Yes	256	0	UntagAll	Unit 1, Port 1
1/2	No	Yes	2	0	UntagAll	Unit 1, Port 2

Table 9: VLAN ID port member configuration

Unit/Port	VLAN	VLAN Name	VLAN	VLAN Name	VLAN	VLAN Name
1/1	256	VLAN #256				
1/2	2	VLAN-2				

Table 10: Spanning-tree port configurations

Unit	Port	Trunk	Participation	Priority	Path	Cost	State
1	1		Disabled				
1	2		Normal	Learning	128	20000	Forwarding

Enabling DDI logging

About this task

Enable DDI logging on ports.

Procedure

1. Enter Ethernet Interface Configuration mode:

   ```
   enable
   configure terminal
   interface Ethernet <port>
   ```

2. At the command prompt, enter the following command:

```
ddi-logging enable [port <port>]
```

Note:

By default, logging is disabled for all ports.

Variable definitions

The following table describes the parameters for the **ddi-logging** command.

Variable	Value
port <port>	Specifies the port in one of the following formats: a single port (3), a range of ports (3-4), or a series of ports (3,5,6).

Viewing DDI logging status

About this task

Display DDI logging port status.

Procedure

1. Log on to CLI to enter User EXEC mode.
2. At the command prompt, enter the following command:

   ```
   show ddi-logging
   ```

Example

The following example shows sample output of the **show ddi-logging** command.

```
Switch>show ddi-logging
DDI Logging enabled on ports : 1
```

Viewing SFP DDI information

Use the following procedure to view SFP DDI information.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. At the command prompt, enter the following command:

   ```
   show interfaces gbic-info
   ```

Example

The following example displays sample output from the **show interfaces gbic-info** command.

```
Switch>show interfaces gbic-info
        Port  Number     3
        GBIC   Type      SX
        Wavelength       850 nm
        Vendor Name      EXTREME NETWORKS
        Vendor OUI       00176A
        Vendor Part #    AFBR-5715PZ-NT1
        Vendor Revision  N/A
        Vendor Serial    AVAGCNAS00FV1
        HW Options       TX_DISABLE TX_FAULT RX_LOSS
        Date Code        10/01/2011
        CLEI Code        IPUIAHCWAA
        Product Code     AA1419048-E6


Digital Diagnostic Interface supported

Calibration: Internal
Rx Power Measurement: Average

--------------------------------------------------------------------------------
        LOW_ALARM    LOW_WARN    ACTUAL   HIGH_WARN   HIGH_ALARM       STATUS
        THRESHOLD    THRESHOLD   VALUE    THRESHOLD   THRESHOLD
--------------------------------------------------------------------------------
Temp(C)        -5.000     0.000      29.468     90.000      95.000      NORMAL
Voltage(V)      2.9700    3.0200      3.2934     3.5800      3.6300     NORMAL
Bias(mA)        2.000     3.000       8.216     15.000      16.000      NORMAL
TxPower(dBm)   -9.5000    -8.9997    -5.2374    -1.9997     -1.0001     NORMAL
RxPower(dBm)   -13.0102   -11.9997   -7.3376    -1.0001      0.0000     NORMAL
```

Viewing environmental status

About this task

Perform this procedure to view the environmental status of the switch or stack.

Procedure

1. Enter User EXEC mode.
2. View environmental status of the switch:

   ```
   show environmental
   ```

Example

```
Switch>show environmental
Unit# PSU1      PSU2      FAN1 FAN2 FAN3 FAN4 Temperature
----- --------- --------- ---- ---- ---- ---- -----------
1     Primary   N/A       OK   OK   OK   N/A  OK   28C


Unit# Model         Switch Capacity Saving PoE Saving
----- ------------ ---------------------- -------------------
1     <Switch#> 0.0 watts              0.0 watts
------------------------------------------------------------
```

```
TOTAL               0.0 watts               0.0 watts
============================================================
```

Where, <Switch#> is the switch model.

Displaying port-mirroring

About this task

Use this procedure to display port-mirroring settings.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Display port mirroring:

   ```
   show port-mirroring
   ```

   Note:

   An asterisk character (*) after the instance number indicates an invalid instance.

Example

The following is an output example for the `show port-mirroring` command.

```
Switch# (config)#show port-mirroring
Port mirroring instance: 1
Monitoring Mode:   Xrx ( -> Port X )
Monitor Unit/Port:  1/5
Unit/Port X:        1/6
Mirror VLAN:        10
Allow Traffic:      Disabled

Port mirroring instance: 2
Monitoring Mode:   XrxOrYtx ( -> Port X or Port Y -> )
Monitor Unit/Port: 1/10
Unit/Port X:       1/11
Unit/Port Y:       2/14
Mirror VLAN:       20
RSPAN VLAN:        200
Allow Traffic:     Disabled

Port mirroring instance: 3
Monitoring Mode:   Disabled

Port mirroring instance: 4
Monitoring Mode:   Disabled
```

Configuring port-mirroring

About this task

Use this procedure to configure port-mirroring.

Procedure

1. Enter Global Configuration mode:

   ```
   enable

   configure terminal
   ```

2. Configure port-mirroring:

   ```
   port-mirroring <1-4> [allow-traffic] mode {disable |Adst monitor-
   port <portList> mirror-MAC-A <H.H.H>|Asrc monitor-port <portList>
   mirror-MAC-A <H.H.H> mirror-MAC-B <H.H.H> | AsrcBdstOrBsrcAdst
   monitor-port <portList> mirror-MAC-A <H.H.H> mirror-MAC-B <H.H.H> |
   AsrcOrAdst monitor-port <portList> mirror-MAC-A <H.H.H> |ManyToOneRx
   monitor-port <portList> mirror-ports <portList> |ManyToOneRxTx
   monitor-port <portList> mirror-ports <portList> |ManyToOneTx
   monitor-port <portList> mirror-ports <portList> |Xrx monitor-port
   <portList> mirror-port-X <portList> |XrxOrXtx monitor-port
   <portList> mirror-port-X <portList> |XrxOrYtx monitor-port
   <portList> mirror-port-X <portList> mirror-port-Y <portList>> |
   XrxYtx monitor-port <portList> mirror-port-X <portList> mirror-port-
   Y <portList> |XrxYtxOrYrxXtx monitor-port <portList> mirror-port-X
   <portList> mirror-port-Y <portList> |Xtx monitor-port <portList>
   mirror-port-X <portList> [mirror-vlan <VID>] [rspan-vlan <VID>]
   ```

Example

```
port-mirroring mode {Xrx|Xtx|XrxOrXtx|XrxOrYtx} monitor-port <portList> mirror-port-X
<portList> [mirror-port-Y <portList>]
```

```
(config)# port-mirroring mode Xrx monitor-port 1/5 mirror-port-X 1/6
```

```
(config)# port-mirroring mode XrxorYtx monitor-port 1/5 mirror-port-X 1/6 mirror-port-Y
2/14
```

```
(config)# port-mirroring mode Xrx monitor-port 1/5 mirror-port-X 1/6 mirror-vla
```

```
(config)# port-mirroring 2 mode XrxorYtx monitor-port 1/10 mirror-port-X 1/11 mirror-port-
Y 2/14 mirror-vlan 20 rspan-vlan 200
```

Variable definitions

Use the data in the following table to use the **port-mirroring** command.

Variable	Value
<1-4>	Port-mirroring instance number. Default is 1.
allow-traffic	Enables bi-direction Monitor Port. **Note:** You cannot use this parameter to configure RSPAN sessions, because the local monitoring

Table continues…

Variable	Value
	port allows traffic by default during an RSPAN session. If the **allow-traffic** parameter is used with the **rspan-vlan <VID>** parameter, an error message is displayed.
Adst	Mirror packets with destination MAC address A.
Asrc	Mirror packets with source MAC address A.
AsrcBdst	Mirror packet with source MAC address A and destination MAC address B.
AsrcBdstOrBsrcAdst	Mirror packets with source MAC address A and destination MAC address B or packets with source MAC address B and destination MAC address A.
AsrcOrAdst	Mirror packet with source or destination MAC address A.
ManyToOneRx	Mirror many to one port mirroring on ingress and egress packets.
ManytoOneRxTx	Mirror many to one port mirroring ingress and egress traffic.
ManyToOneTx	Mirror many to one port mirroring on egress packets.
Xrx	Mirror packets received on port X.
XrxOrXtx	Mirror packets received or transmitted on port X.
XrxOrYtx	Mirror packets received on port X or transmitted on port Y.
XrxYtx	Mirror packets received on port X and transmitted on port Y.
XrxYtxOrYrxXtx	Mirror packets received on port X and transmitted on port Y, or packets received on port Y and transmitted on port X.
Xtx	Mirror packets received on port X .
mirror-vlan <VID>	Enables VLAN port-mirroring and specifies the source/destination VLAN for mirrored traffic.
rspan-vlan <VID>	Enables remote port-mirroring and specifies the VLAN for mirrored traffic.

Disabling many-to-many port-mirroring

About this task

Use this procedure to disable many-to-many port-mirroring

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Disable a specific instance:

   ```
   port-mirroring [<1-4>] mode disable
   ```

   OR

   ```
   no port-mirroring [<1-4>]
   ```

3. Disable all instances:

   ```
   no port-mirroring
   ```

Example

Disable port-mirroring for instance 3:

```
Switch>enable
Switch#config term
Enter configuration commands, one per line.  End with CNTL/Z.
Switch(config)#no port-mirroring 3
```

```
Switch>enable
Switch#config term
Enter configuration commands, one per line.  End with CNTL/Z.
Switch(config)#no port-mirroring
```

Variable definitions

Use the data in the following table to use the **no port-mirroring** command.

Variable	Value
<1-4>	The port-mirroring instance.

Configuring an RSPAN source session

An RSPAN source session associates a port mirroring instance with an RSPAN VLAN. The output of this session is a stream of packets sent to the RSPAN VLAN.

Before you begin

Create an RSPAN VLAN and establish port membership.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Configure the RSPAN source session:

   ```
   port-mirroring [<1-4>] [allow-traffic] mode {disable | Adst monitor-
   port <portList> mirror-MAC-A <H.H.H> | Asrc monitor-port <portList>
   mirror-MAC-A <H.H.H> | AsrcBdst monitor-port <portList> mirror-MAC-A
   <H.H.H> mirror-MAC-B <H.H.H> | AsrcBdstOrBsrcAdst monitor-port
   ```

```
<portList> mirror-MAC-A <H.H.H> mirror-MAC-B <H.H.H> | AsrcOrAdst
monitor-port <portList> mirror-MAC-A <H.H.H> | ManyToOneRx monitor-
port <portList> mirror-ports <portList> | ManyToOneRxTx monitor-port
<portList> mirror-ports <portList> | ManyToOneTx monitor-port
<portList> mirror-ports <portList> | Xrx monitor-port <portList>
mirror-port-X <portList> | XrxOrXtx monitor-port <portList> mirror-
port-X <portList> | XrxOrYtx monitor-port <portList> mirror-port-X
<portList> mirror-port-Y <portList> | XrxYtx monitor-port <portList>
mirror-port-X <portList> mirror-port-Y <portList> | XrxYtxOrYrxXtx
monitor-port <portList> mirror-port-X <portList> mirror-port-Y
<portList> | Xtx monitor-port <portList> mirror-port-X <portList>
[mirror-vlan <VID>] [rspan-vlan <VID>]
```

3. Display and verify the RSPAN settings:

   ```
   show port-mirroring
   ```

Example

The following example displays sample output for configuring an RSPAN source session:

```
Switch(config)# vlan create 1009 type port remote-span
Switch(config)# vlan members add 1009 1/26
Switch(config)# port-mirroring 2 ManyToOneRx monitor-port 1/26 mirror-ports  1/1-12 rspan-
vlan 1009
```

Variable definitions

Use the data in the following table to use the **port-mirroring** command.

Variable	Value
<1-4>	Port-mirroring instance number. Up to four RSPAN destination sessions can be configured with ID from 1 to 4. Default is 1.
allow-traffic	Enables bi-direction Monitor Port.
Adst	Mirror packets with destination MAC address A.
Asrc	Mirror packets with source MAC address A.
AsrcBdst	Mirror packet with source MAC address A and destination MAC address B.
AsrcBdstOrBsrcAdst	Mirror packets with source MAC address A and destination MAC address B or packets with source MAC address B and destination MAC address A.
AsrcOrAdst	Mirror packet with source or destination MAC address A.
ManyToOneRx	Mirror many to one port mirroring on ingress and egress packets.

Table continues…

Variable	Value
ManyToOneTx	Mirror many to one port mirroring on egress packets.
Xrx	Mirror packets received on port X.
XrxOrXtx	Mirror packets received or transmitted on port X.
XrxOrYtx	Mirror packets received on port X or transmitted on port Y.
XrxYtx	Mirror packets received on port X and transmitted on port Y.
XrxYtxOrYrxXtx	Mirror packets received on port X and transmitted on port Y, or packets received on port Y and transmitted on port X.
Xtx	Mirror packets received on port X .
mirror-vlan <VID>	Enables VLAN port-mirroring and specifies the source/destination VLAN for mirrored traffic.
rspan-vlan <VID>	Enables remote port-mirroring and specifies the VLAN for mirrored traffic.

Configuring an RSPAN destination session

Use this procedure to configure an RSPAN destination session.

Before you begin

Create an RSPAN VLAN.

About this task

An RSPAN destination session associates the destination port with an RSPAN VLAN. The destination session collects all RSPAN VLAN traffic and sends it out the designated RSPAN destination port.

Procedure

1. Enter Global Configuration mode:

   ```
   enable

   configure terminal
   ```

2. Use the following command to configure an RSPAN destination session:

   ```
   [no] port-mirroring rspan <1-4> [destination-port <port>] [vlan
   <VID>]
   ```

Example

The following example displays sample output for configuring an RSPAN destination session:

```
Switch> enable
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
```

```
Switch(config)# vlan create 1009 type port remote-span
Switch(config)# vlan members add 1009 1/2,26
Switch(config)# port-mirroring rspan 2 destination-port 1/26 vlan 1009
Switch(config)#show port-mirroring rspan
```

Variable definitions

Use the data in the following table to use the **port-mirroring rspan** command.

Variable	Value
[no]	Disables the destination RSPAN session.
<1-4>	RSPAN destination session number. Up to four RSPAN destination sessions can be configured with ID from 1 to 4 Default is 1.
[destination-port <port>]	Specifies the port to be used as the destination port.
[vlan <VID>]	Specifies the RSPAN VLAN to be associated with the destination port.

Displaying RSPAN information

Use this procedure to display RSPAN information.

Procedure

1. Enter Privileged EXEC mode:

   `enable`

2. Display RSPAN information:

   `show port-mirroring rspan`

   Note:

   An asterisk character (*) after the instance number indicates an invalid instance.

Example

The following is an output example for the `show port-mirroring rspan` command.

```
Switch(config)#show port-mirroring rspan
=========================
RSPAN Source Sessions
=========================
Inst RSPAN VLAN RSPAN MTP
---- ---------- ----------

=========================
RSPAN Destination Sessions
=========================
Inst RSPAN VLAN RSPAN MTP
---- ---------- ----------
1*   300        20
```

```
* Instance is not valid.
  Configured ports may not reside on the local unit
Switch(config)#
```

Configuring RSPAN over MLT

This procedure enables the monitor port to be part of a trunk.

Before you begin

Configure an MLT.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure an RSPAN source session:

   ```
   port-mirroring <1-4> [allow-traffic] mode {disable |Adst monitor-
   port <portList> mirror-MAC-A <H.H.H>|Asrc monitor-port <portList>
   mirror-MAC-A <H.H.H> mirror-MAC-B <H.H.H> | AsrcBdstOrBsrcAdst
   monitor-port <portList> mirror-MAC-A <H.H.H> mirror-MAC-B <H.H.H> |
   AsrcOrAdst monitor-port <portList> mirror-MAC-A <H.H.H> |ManyToOneRx
   monitor-port <portList> mirror-ports <portList> |ManyToOneRxTx
   monitor-port <portList> mirror-ports <portList> |ManyToOneTx
   monitor-port <portList> mirror-ports <portList> |Xrx monitor-port
   <portList> mirror-port-X <portList> |XrxOrXtx monitor-port
   <portList> mirror-port-X <portList> |XrxOrYtx monitor-port
   <portList> mirror-port-X <portList> mirror-port-Y <portList>> |
   XrxYtx monitor-port <portList> mirror-port-X <portList> mirror-port-
   Y <portList> |XrxYtxOrYrxXtx monitor-port <portList> mirror-port-X
   <portList> mirror-port-Y <portList> |Xtx monitor-port <portList>
   mirror-port-X <portList> [mirror-vlan <VID>] [rspan-vlan <VID>]
   ```

3. Verify settings:

   ```
   show port-mirroring rspan
   ```

Example

The following is an example of RSPAN over MLT configuration.

```
Switch(config)# mlt create 10 enable members 11,12,13
Switch(config)# port-mirroring 1 mode ManyToOneRx monitor-port 11 mirror-ports
20,21,22,23 rspan-vlan 100

Switch(config)# show port-mirroring rspan
============================
RSPAN Source Sessions
============================
Inst RSPAN VLAN RSPAN MTRUNK
```

```
---- ---------- ------------
1    100        10
============================
RSPAN Destination Sessions
============================
Inst RSPAN VLAN RSPAN MTP
---- ---------- ------------

(config)#show port-mirroring
Port mirroring instance: 1
Monitoring Mode: ManyToOne
Monitor Unit/Port: 11(Trunk 10)
Mirrored Ports:    20-23
RSPAN VLAN: 100
Allow Traffic: Disabled
```

On a neighboring DUT, configure a similar MLT to connect the ports from the monitor trunk.

Configuring RSPAN over LACP

This procedure enables the monitor port to be part of a trunk.

Before you begin

- Enable LACP aggregation.
- Configure administrative LACP key.
- Configuring LACP mode of operation.

Procedure

1. Enter Global Configuration mode:

   ```
   enable

   configure terminal
   ```

2. Configure an RSPAN source session:

   ```
   port-mirroring <1-4> [allow-traffic] mode {disable |Adst monitor-
   port <portList> mirror-MAC-A <H.H.H>|Asrc monitor-port <portList>
   mirror-MAC-A <H.H.H> mirror-MAC-B <H.H.H> | AsrcBdstOrBsrcAdst
   monitor-port <portList> mirror-MAC-A <H.H.H> mirror-MAC-B <H.H.H> |
   AsrcOrAdst monitor-port <portList> mirror-MAC-A <H.H.H> |ManyToOneRx
   monitor-port <portList> mirror-ports <portList> |ManyToOneRxTx
   monitor-port <portList> mirror-ports <portList> |ManyToOneTx
   monitor-port <portList> mirror-ports <portList> |Xrx monitor-port
   <portList> mirror-port-X <portList> |XrxOrXtx monitor-port
   <portList> mirror-port-X <portList> |XrxOrYtx monitor-port
   <portList> mirror-port-X <portList> mirror-port-Y <portList>> |
   XrxYtx monitor-port <portList> mirror-port-X <portList> mirror-port-
   Y <portList> |XrxYtxOrYrxXtx monitor-port <portList> mirror-port-X
   <portList> mirror-port-Y <portList> |Xtx monitor-port <portList>
   mirror-port-X <portList> [mirror-vlan <VID>] [rspan-vlan <VID>]
   ```

3. Verify settings:

   ```
   show port-mirroring rspan
   ```

Example

The following is an example of RSPAN over LACP configuration.

```
Switch(config)# interface fast-ehernet 11,12,13
Switch(config)# lacp key 10
Switch(config)# lacp mode active
Switch(config)# lacp aggregation enable
```

On a neighboring DUT, configure a similar LACP so that aggregation can occur.

```
(config)#show lacp aggr
Aggr ID Trunk Status   Type Members
------- ----- -------- ---- -------------
8224    32    Enabled  LA   11,12,13


(config)#port-mirroring 1 mode ManyToOneRx monitor-port 11 mirror-ports 20,21,22,23 rspan-
vlan 100

(config)#show port-mirroring rspan
============================
RSPAN Source Sessions
============================
Inst RSPAN VLAN RSPAN MTRUNK
---- ---------- ------------
1    100        32
============================
RSPAN Destination Sessions
============================
Inst RSPAN VLAN RSPAN MTP
---- ---------- ------------

(config)#show port-mirroring
Port mirroring instance: 1
Monitoring Mode: ManyToOne
Monitor Unit/Port: 11(Trunk 32)
Mirrored Ports:    20-23
RSPAN VLAN: 100
Allow Traffic: Disabled
```

Configuring port-mirroring on EAP ports

About this task

Use this procedure to enable or disable port mirroring on EAP ports.

Procedure

1. Enter Global Configuration mode:

   ```
   enable

   configure terminal
   ```

2. To enable port mirroring on EAP ports, enter the following command:

```
eapol allow-port-mirroring
```

3. To disable port mirroring on EAP ports, enter the following command:

   ```
   no eapol allow-port-mirroring
   ```

   OR

   ```
   default eapol allow-port-mirroring
   ```

Viewing MACsec statistics

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. View MACsec statistics:

   ```
   show macsec statistics [<port>]
   ```

3. View the secure-channel inbound MACsec statistics:

   ```
   show macsec statistics <port> secure-channel inbound
   ```

4. View the secure-channel outbound MACsec statistics:

   ```
   show macsec statistics <port> secure-channel outbound
   ```

Example

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

```
Switch(config)#show macsec statistics 2/1 secure-channel inbound
2017-11-10 16:45:12 GMT+03:00    UTC time: 2017-11-10 13:45:12


================================================================================
                 MACSEC Port Inbound Secure Channel Statistics
================================================================================
            Late         Delayed        Unchecked     Octets          Octets
PortId     Packets      Packets        Packets       Validated       Decrypted
--------------------------------------------------------------------------------
2/1         0            0              0             0               0
```

```
witch(config)#show macsec statistics 2/1 secure-channel outbound
2017-11-10 16:53:56 GMT+03:00    UTC time: 2017-11-10 13:53:56


================================================================================
                 MACSEC Port Outbound Secure Channel Statistics
================================================================================
          Octets          Octets
PortId    Protected       Encrypted
--------------------------------------------------------------------------------
2/1        0              0
```

Variable definitions

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

Variable	Definition
<port>	Specifies the port for which to display MACsec statistics.

System diagnostics and statistics using Enterprise Device Manager

This section provides procedures to perform system diagnostics and gather statistics using Enterprise Device Manager (EDM).

Prerequisites

- Open one of the supported browsers.
- Enter the IP address of the switch to open an EDM session.

Port Mirroring using EDM

The following sections describe Port Mirroring:

- Viewing Port Mirroring using EDM
- Configuring Port Mirroring using EDM

Viewing Port Mirroring using EDM

View Port Mirroring to troubleshoot the network.

Procedure steps

1. From the navigation tree, double-click **Edit**.
2. In the Edit tree, double-click **Diagnostics**.
3. In the Diagnostics tree, double-click **Port Mirrors**.

Variable definitions

The following table describes the Port Mirrors tab fields.

Variable	Value
Instance	Specifies the numerical assignment of the port mirroring (1-4)
Port Mode	Specifies the port monitoring mode.
Monitor Port	Identifies the monitoring port.
PortListX	Identifies the ports monitored for XrX/Xtx, and manytoOne related mode.
PortListY	Identifies the ports monitored for Yrx/Ytx related mode.
MacAddressA	Specifies the MAC address of the monitored port using Asrc/Adst related mode.
MacAddressB	Specifies the MAC address of the monitored port using Bsrc/Bdst related mode.
AllowTraffic	Indicates whether bi-directional mirroring traffic is enabled.
RspanVlan	Specifies the RspanVlan to be associated with a source port-mirroring instance.
MirrorVlan	Specifies mirroring to a destination VLAN.

Configuring Port Mirroring using EDM

Configure Port Mirroring to troubleshoot the network.

Procedure steps

1. From the navigation tree, double-click **Edit**.
2. In the Edit tree, double-click **Diagnostics**.
3. In the Diagnostics tree, double-click **Port Mirrors**.
4. In the work area, click **Insert**.
5. In the **Instance** box, type instance number.
6. In the **PortMode** section, click a mode.
7. Click the **MonitorPort** ellipsis (…).
8. In the **MonitorPort** list, click a monitor port.
9. Click **Ok**.
10. If the PortMode is Xrx, Xtx, or both, or manytoOne related modes, click the PortListX ellipsis (…).
11. In the **PortListX** list, click a port, ports, or All to add to the list.
12. Click **Ok**.
13. If the PortMode is Yrx, Ytx, or both related modes, click the **PortListY** ellipsis (…).
14. In the **PortListY**, click a port, ports, or **All** to add to the list.

15. Click **Ok**.
16. If the PortMode is Asrc, Adst, or both related modes, in the **MacAddressA**, type an address.
17. If the PortMode is Bsrc, Bdst, or both related modes, in the **MacAddressB**, type an address.
18. To enable bi-directional traffic, click the **AllowTraffic** box.
19. Click the **RspanVlan** ellipsis (…).
20. In the **RspanVlan** list, click a VLAN.
21. In the **MirrorVlan** field, enter the VLAN ID.
22. Click **Ok**.
23. Click **Insert**.

Variable definitions

The following table describes the Port Mirrors tab fields.

Variable	Value
Instance	Indicates the Port Mirroring instance number (1-4)
Port Mode	Indicates the supported Port Mirroring modes. The modes are: • Adst—Mirror packets with destination MAC address A. • Asrc—Mirror packets with source MAC address A. • AsrcBdst—Mirror packets with source MAC address A and destination MAC address B. • AsrcBdstOrBsrcAdst—Mirror packets with source MAC address A and destination MAC address B or packets with source MAC address B and destination MAC address A. • AsrcOrAdst—Mirror packets with source or destination MAC address A. • manytoOneRx—Many to one port mirroring on ingress packets. • manytoOneRxTx—Many to one port mirroring on ingress and egress traffic • manytoOneTx—Many to one port mirroring on egress packets. • Xrx—Mirror packets received on port X. • XrxOrXtx—Mirror packets received or transmitted on port X. • XrxOrYtx—Mirror packets received on port X or transmitted on port Y.

Table continues…

Variable	Value
	• XrxYtx—Mirror packets received on port X and transmitted on port Y. This mode is not recommended for mirroring broadcast and multicast traffic. • XrxYtxOrYrxXtx—Mirror packets received on port X and transmitted on port Y or packets received on port Y and transmitted on port X. • Xtx—Mirror packets transmitted on port X. The default value is Disabled.
Monitor Port	Specifies the monitoring port.
PortListX	Indicates the switch port to be monitored by the designated monitor port. This port is monitored according to the value X in the Monitoring Mode field.
PortListY	Indicates the switch port to be monitored by the designated monitor port. This port is monitored according to the value Y in the Monitoring Mode field
MacAddressA	Specifies the mirroring MAC address A.
MacAddressB	Specifies the mirroring MAC address B.
AllowTraffic	Indicates whether bi-directional mirroring traffic is enabled.
RspanVlan	Specifies the RspanVlan to be associated with a source port-mirroring instance.
MirrorVlan	Specifies mirroring to a destination VLAN.

Remote Port Mirroring using EDM

Remote Switch Port ANalyzer (RSPAN), also known as Remote Port Mirroring, enhances port mirroring by enabling mirroring traffic to be sent to one or more switches or stacks on the network using an intermediate VLAN for forwarding the mirrored traffic.

Use the following procedures to configure source and destination sessions.

Configuring an RSPAN source session using EDM

Use the following procedure to configure an RSPAN source session..

Before you begin

Create a VLAN for RSPAN traffic and enable RSPAN on this VLAN.

About this task

An RSPAN source session associates a port mirroring instance with an RSPAN VLAN. The output of this session is a stream of packets sent to the RSPAN VLAN.

Procedure

1. From the navigation tree, double-click **Edit**.
2. In the Edit tree, double-click **Diagnostics**.
3. In the Diagnostics tree, double-click **Port Mirrors**.
4. On the toolbar, click **Insert**.

   The Insert Port Mirrors dialog box appears.
5. Configure the parameters as required.
6. In the **RspanVlan** field, select the VLAN for RSPAN traffic.
7. Click **Insert**.

Variable definitions

The following table describes the fields associated with the Port Mirrors tab.

Variable	Value
Instance	Numerical assignment of the port mirroring.
Port Mode	The port monitoring mode. The following options are available: • Adst • Asrc • AsrcBdst • AsrcBdstorBsrcAdst • AsrcorAdst • manytoOneRx • manytoOneRxTx • manytoOneTx • Xrx • XrxorXtx • XrxorYtx • XrxYtx • XrxYtxOrYrxXtx • Xtx The default value is Adst.
Monitor Port	The port that is the monitoring port.
PortListX	Ports monitored for XrX/Xtx, and manytoOne related mode.
PortListY	Ports monitored for Yrx/Ytx related mode.

Table continues…

Variable	Value
MacAddressA	MAC address of the monitored port using Asrc/Adst related mode.
MacAddressB	MAC address of the monitored port using Bsrc/Bdst related mode.
Allow traffic	Allows or disallow traffic. **Note:** You cannot use the **`Allow traffic`** option with RSPAN.
RspanVlan	Specifies the RSPAN VLAN to be associated with a source port-mirroring instance.

Configuring an RSPAN destination session using EDM

Use the following procedure to configure an RSPAN destination session using EDM.

Before you begin

Create a VLAN for RSPAN traffic and enable RSPAN on this VLAN.

About this task

An RSPAN destination session associates the destination port with the RSPAN VLAN. The destination session collects all RSPAN VLAN traffic and sends it out the designated RSPAN destination port.

Procedure

1. From the navigation tree, click **Edit**.
2. In the Edit tree, click **Diagnostics**.
3. In the Diagnostics tree, click **Port Mirrors**.
4. On the work area, click the **RSPAN** tab.
5. On the toolbar, click **Insert**.

   EDM displays the Insert RSPAN window.
6. Configure the parameters as required.
7. Click **Insert**.

Variable definitions

The following table describes the fields associated with the RSPAN tab.

Variable	Value
Instance	Specifies the destination session instance number.
DestinationPort	Specifies the port to be used as a destination port
RspanVlan	Specifies the RSPAN VLAN to be associated with the destination port.

Configuring Stack Monitor using EDM

Use the following procedure to configure the Stack Monitor.

Procedure steps

1. From the navigation tree, double-click **Edit**.
2. In the Edit tree, double-click **Chassis**.
3. In the Chassis tree, double-click **Chassis**.
4. On the work area, click the **Stack Monitor** tab.
5. Select **StackErrorNotificationEnabled** to enable stack monitoring.
6. Set the stack size you want to monitor in the **ExpectedStackSize** field.
7. Sets the traps interval in the **StackErrorNotificationInterval** field.
8. Select **StackRebootUnitOnFailure** to enable rebooting of stack units on failure.
9. Set the retry count for the stack in the **StackRetryCount** field.
10. On the toolbar, click **Apply**.

Variable definitions

The following table describes the Stack Monitor tab fields.

Variable	Value
StackErrorNotificationEnabled	Enables or disables the Stack Monitoring feature.
ExpectedStackSize	Sets the size of the stack to monitor. Valid range is 2–8.
StackErrorNotificationInterval	Sets the interval between traps, in seconds. Valid range is 30 to 300 seconds.
StackRebootUnitOnFailure	Enables or disables the rebooting stack units on failure.
StackRetryCount	Sets the retry count for the stack. Valid range is 0-4294967295.

Viewing power supply information using EDM

Use this procedure to display the operating status of switch power supplies.

The power supply parameters for the PoE switches differ slightly because they support Power over Ethernet (PoE).

Procedure steps

1. From the navigation tree, double-click **Edit**.
2. In the Edit tree, double-click **Chassis**.

3. In the Chassis tree, double-click **Environment**.
4. On the work area, click the **PowerSupply** tab.

Variable definitions

Use the data in the following table to help you understand the switch power supply display.

Variable	Value
Description	Indicates the chassis number, power supply number, and the type of power supply.
OperState	Indicates the status of the applicable power supply. • **other**: some other state • **notAvail**: state not available • **removed**: component was removed • **disabled**: operation disabled • **normal**: normal operation • **resetInProg**: reset is in progress • **testing**: system is performing a self test • **warning**: system is operating at a warning level • **nonFatalErr**: system is operating at error level • **fatalErr**: a fatal error stopped operation • **notConfig**: a module needs to be configured; the allowable values are determined by the component type • **obsoleted**: obsoleted
Important: For a stack environment, this work area displays power supply information for each switch unit in the stack.	

Viewing switch fan information using EDM

Use this procedure to display information about the operating status of the switch fans.

Procedure steps

1. From the navigation tree, double-click **Edit**.
2. In the Edit tree, double-click **Chassis**.
3. In the Chassis tree, double-click **Environment**.
4. On the work area, click the **Fan** tab.

Variable definitions

The following table describes the Fan operating status.

Variable	Value
Unit 1 Fan 1	Indicates the status of Fan 1.
Unit 1 Fan 2	Indicates the status of Fan 2.
Unit 1 Fan 3	Indicates the status of Fan 3.
Unit 1 Fan 4	Indicates the status of Fan 4.
Important: For a stack environment, this work area displays similar fan information for each switch unit in the stack.	

Viewing switch temperature using EDM

Use the following procedure to display switch temperature information.

Procedure steps

1. From the navigation tree, double-click **Edit**.
2. In the Edit tree, double-click **Chassis**.
3. In the Chassis tree, double-click **Environment**.
4. In the work area, click the **Temperature** tab.
5. On the tool bar, click **Refresh** to update the data.

Variable definitions

The following table describes the Fan operating status.

Variable	Value
Unit	Indicates the switch unit number in a stack. For a standalone switch, the default value is 1.
Temperature	Indicates the switch unit operating temperature.

Chassis configuration statistics management using EDM

Use the information in this section to display and graph chassis configuration statistics.

Graphing chassis IP statistics using EDM

Perform this procedure to display and graph switch IP statistics.

Procedure steps

1. From the navigation tree, double-click **Graph**.
2. In the Graph tree, double-click **Chassis**.
3. In the work area, click the **IP** tab.
4. On the toolbar, select a **Poll Interval** from the list.
5. On the toolbar, you can click **Clear Counters** to reset the IP statistics counters.
6. To select statistics to graph, click a statistic type row under a column heading.
7. On the toolbar, click **Line Chart**, **Area Chart**, **Bar Chart**, or **Pie Chart**.

Variable definitions

Use the data in the following table to help you understand IP statistics.

Variable	Value
InReceives	The total number of input datagrams received from interfaces, including those received in error.
InHdrErrors	The number of input datagrams discarded due to errors in their IP headers, including bad checksums, version number mismatch, other format errors, time-to-live exceeded, errors discovered in processing their IP options.
InAddrErrors	The number of input datagrams discarded because the IP address in the IP header destination field was not a valid address. This count includes invalid addresses (for example, 0.0.0.0) and addresses of unsupported Classes (for example, Class E). For addresses that are not IP Gateways and therefore do not forward datagrams, this counter includes datagrams discarded because the destination address was not a local address.
ForwDatagrams	The number of input datagrams for which this entity was not their final IP destination, as a result of which an attempt was made to find a route to forward them to that final destination. For addresses that do not act as IP Gateways, this counter includes only those packets Source-Routed by way of this address with successful Source-Route option processing.
InUnknownProtos	The number of locally addressed datagrams received successfully but discarded because of an unknown or unsupported protocol.
InDiscards	The number of input IP datagrams for which no problems are encountered to prevent their continued processing, but that are discarded (for example, for lack of buffer space). This counter does not include any datagrams discarded while awaiting reassembly.
InDelivers	The total number of input datagrams successfully delivered to IP user-protocols (including ICMP).
OutRequests	The total number of IP datagrams that local IP user-protocols (including ICMP) supplied to IP in requests for transmission. This counter does not include any datagrams counted in ipForwDatagrams.
OutDiscards	The number of output IP datagrams for which no problem was encountered to prevent their transmission to their destination, but that are discarded (for

Table continues…

Variable	Value
	example, for lack of buffer space). This counter can include datagrams counted in ipForwDatagrams if any such packets met this (discretionary) discard criterion.
OutNoRoutes	The number of IP datagrams discarded because no route can be found to transmit them to their destination. This counter also includes any packets counted in ipForwDatagrams that have no route. This includes any datagrams a host cannot route because all of its default gateways are down.
FragOKs	The number of IP datagrams successfully fragmented at this entity.
FragFails	The number of IP datagrams that are discarded because they need to be fragmented at this entity but cannot be, for example, because their Don't Fragment flag was set.
FragCreates	The number of generated IP datagram fragments because of a fragmentation at this entity.
ReasmReqds	The number of IP fragments received that needed to be reassembled at this entity.
ReasmOKs	The number of IP datagrams successfully reassembled.
ReasmFails	The number of failures detected by the IP reassembly algorithm (for example, timed out, errors). This is not necessarily a count of discarded IP fragments because some algorithms (notably the algorithm in RFC815) can lose track of the number of fragments by combining them as they are received.

Graphing chassis ICMP In statistics using EDM

Use this procedure to display and graph ICMP In statistics.

Procedure steps

1. From the navigation tree, double-click **Graph**.
2. In the Graph tree, double-click **Chassis**.
3. In the work are, click the **ICMP In** tab.
4. On the toolbar, select a **Poll Interval** from the list.
5. On the toolbar, you can click **Clear Counters** to reset the IP statistics counters.
6. To select statistics to graph, click a statistic type row under a column heading.
7. On the toolbar, click **Line Chart**, **Area Chart**, **Bar Chart**, or **Pie Chart**.

Variable definitions

Use the data in the following table to help you understand ICMP In statistics.

Variable	Value
SrcQuenchs	The number of ICMP Source Quench messages received.

Table continues…

Variable	Value
Redirects	The number of ICMP Redirect messages received.
Echos	The number of ICMP Echo (request) messages received.
EchoReps	The number of ICMP Echo Reply messages received.
Timestamps	The number of ICMP Timestamp (request) messages received.
TimestampReps	The number of ICMP Timestamp Reply messages received.
AddrMasks	The number of ICMP Address Mask Request messages received.
AddrMaskReps	The number of ICMP Address Mask Reply messages received.
ParmProbs	The number of ICMP Parameter Problem messages received.
DestUnreachs	The number of ICMP Destination Unreachable messages received.
TimeExcds	The number of ICMP Time Exceeded messages received.

Graphing chassis ICMP Out statistics using EDM

Use this procedure to display and graph ICMP Out statistics.

Procedure steps

1. From the navigation tree, double-click **Graph**.
2. In the Graph tree, double-click **Chassis**.
3. In the work are, click the **ICMP Out** tab.
4. On the toolbar, select a **Poll Interval** from the list.
5. On the toolbar, you can click **Clear Counters** to reset the IP statistics counters.
6. To select statistics to graph, click a statistic type row under a column heading.
7. On the toolbar, click **Line Chart**, **Area Chart**, **Bar Chart**, or **Pie Chart**.

Variable definitions

Use the data in the following table to help you understand ICMP Out statistics.

Variable	Value
SrcQuenchs	The number of ICMP Source Quench messages sent.
Redirects	The number of ICMP Redirect messages received. For a host, this object is always zero because hosts do not send redirects.
Echos	The number of ICMP Echo (request) messages sent.
EchoReps	The number of ICMP Echo Reply messages sent.
Timestamps	The number of ICMP Timestamp (request) messages sent.
TimestampReps	The number of ICMP Timestamp Reply messages sent.
AddrMasks	The number of ICMP Address Mask Request messages sent.
AddrMaskReps	The number of ICMP Address Mask Reply messages sent.
ParmProbs	The number of ICMP Parameter Problem messages sent.

Table continues…

Variable	Value
DestUnreachs	The number of ICMP Destination Unreachable messages sent.
TimeExcds	The number of ICMP Time Exceeded messages sent.

Graphing chassis TCP statistics using EDM

Use this procedure to display and graph TCP statistics.

Procedure steps

1. From the navigation tree, double-click **Graph**.
2. In the Graph tree, double-click **Chassis**.
3. In the work area, click the **TCP** tab.
4. On the toolbar, select a **Poll Interval** from the list.
5. On the toolbar, you can click **Clear Counters** to reset the IP statistics counters.
6. To select statistics to graph, click a statistic type row under a column heading.
7. On the toolbar, click **Line Chart**, **Area Chart**, **Bar Chart**, or **Pie Chart**.

Variable definitions

Use the data in the following table to help you understand TCP statistics.

Variable	Value
ActiveOpens	The number of times TCP connections make a direct transition to the SYN-SENT state from the CLOSED state.
PassiveOpens	The number of times TCP connections make a direct transition to the SYN-RCVD state from the LISTEN state.
AttemptFails	The number of times TCP connections make 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 make a direct transition to the LISTEN state from the SYN-RCVD state.
EstabResets	The number of times TCP connections make a direct transition to the CLOSED state from either the ESTABLISHED state or the CLOSE-WAIT state.
CurrEstab	The number of TCP connections for which the current state is either ESTABLISHED or CLOSE-WAIT.
InSegs	The total number 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 number of segments retransmitted, that is, the number of TCP segments transmitted containing one or more previously transmitted octets.

Table continues…

Variable	Value
InErrs	The total number of segments received in error (for example, bad TCP checksums).
OutRsts	The number of TCP segments sent containing the RST flag.
HCInSegs	The number of segments received, including those received in error. This count includes segments received on currently established connections. This object is the 64-bit equivalent of InSegs.
HCOutSegs	The number of segments sent, including those on current connections, but excluding those containing only retransmitted octets. This object is the 64-bit equivalent of OutSegs.

Graphing chassis UDP statistics using EDM

Use this procedure to display and graph UDP statistics.

Procedure steps

1. From the navigation tree, double-click **Graph**.
2. In the Graph tree, double-click **Chassis**.
3. In the work area, click the **UDP** tab.
4. On the toolbar, select a **Poll Interval** from the list.
5. On the toolbar, you can click **Clear Counters** to reset the IP statistics counters.
6. To select statistics to graph, click a statistic type row under a column heading.
7. On the toolbar, click **Line Chart**, **Area Chart**, **Bar Chart**, or **Pie Chart**.

Variable definitions

Use the data in the following table to understand the UDP statistics.

Variable	Value
InDatagrams	The total number of UDP datagrams delivered to UDP users.
NoPorts	The total number of received UDP datagrams for which there was no application at the destination port.
InErrors	The number of received UDP datagrams that cannot be delivered for reasons other than the lack of an application at the destination port.
OutDatagrams	The total number of UDP datagrams sent from this entity.
HCInDatagrams	The number of TCP connections for which the current state is either ESTABLISHED or CLOSE-WAIT.
HCOutDatagrams	The number of UDP datagrams sent from this entity, for devices that can transmit more than 1 million UDP datagrams for each second. Discontinuities in the value of this counter can occur at reinitialization of the management system, and at other times as indicated by discontinuities in the value of sysUpTime.

Graphing chassis Fabric Attach statistics using EDM

About this task

Use this procedure to display and graph Fabric Attach statistics.

Procedure

1. From the navigation tree, double-click **Graph**.
2. In the Graph tree, double-click **Chassis**.
3. In the work area, click the **Fabric Attach** tab.
4. The table data refreshes automatically based on the value selected in the **Poll Interval** field.
5. To select statistics to graph, click a row under a column heading.
6. On the toolbar, click **Line Chart**, **Area Chart**, **Bar Chart**, or **Pie Chart**.
7. Click **Reset Counters** to reset the values of the counters.
8. Click **Clear Counters** to clear the counters and start over at zero.

Variable definitions

Use the data in the following table to help you understand Fabric Attach statistics.

Name	Description
DiscElemReceived	Indicates the number of FA Element TLVs received on the identified port.
DiscElemExpired	Indicates the number of discovered FA elements from received FA Element TLVs that have expired on the identified port. This counter is not incremented when elements are deleted for reasons other than expiration.
DiscElemDeleted	Indicates the number of discovered FA elements from received FA Element TLVs that have been deleted on the identified port. This counter is only incremented when elements are deleted for reasons other than expiration.
DiscAuthFailed	Indicates the number of received FA Element TLVs for which authentication was attempted and failed on the identified port.
AsgnReceived	Indicates the number of I-SID/VLAN bindings received in FA I-SID/VLAN Assignment TLVs on the identified port.
AsgnAccepted	Indicates the number of I-SID/VLAN bindings from received FA I-SID/VLAN Assignment TLVs that are accepted (activated) on the identified port. This counter is incremented when the binding transitions

Table continues…

Name	Description
	from a non-accepted state such as 'pending'or 'rejected' to the accepted state.
AsgnRejected	Indicates the number of I-SID/VLAN bindings from received FA I-SID/VLAN Assignment TLVs that are rejected on the identified port. This counter is incremented when the binding transitions from a non-rejected state such as 'pending' or 'accepted' to the rejected state.
AsgnExpired	Indicates the number of I-SID/VLAN bindings from received FA I-SID/VLAN Assignment TLVs that have expired on the identified port. This counter is not incremented when bindings are deleted for reasons other than expiration.
AsgnDeleted	Indicates the number of I-SID/VLAN bindings from received FA I-SID/VLAN Assignment TLVs that have been deleted on the identified port. This counter is only incremented when bindings are deleted for reasons other than expiration.
AsgnAuthFailed	Indicates the number of received FA I-SID/VLAN Assignment TLVs for which authentication was attempted and failed on the identified port.

Port configuration statistics management using EDM

Use the information in this section to display and graph port configuration statistics.

Graphing port interface statistics using EDM

Use this procedure to display and graph interface parameters for a port.

Procedure steps

1. On the Device Physical View, click a port.
2. From the navigation tree, double-click **Graph**.
3. In the Graph tree, double-click **Port**.
4. In the work area, click the **Interface** tab.
5. On the toolbar, select a **Poll Interval** from the list.
6. On the toolbar, you can click **Clear Counters** to reset the IP statistics counters.
7. To select statistics to graph, click a statistic type row under a column heading.
8. On the toolbar, click **Line Chart**, **Area Chart**, **Bar Chart**, or **Pie Chart**.

Variable definitions

Use the data in the following table to help you understand interface statistics.

Variable	Value
InOctets	The total number of octets received on the interface, including framing characters.
OutOctets	The total number of octets transmitted out of the interface, including framing characters.
InUcastPkts	The number of packets delivered by this sublayer to a higher sublayer that are not addressed to a multicast or broadcast address at this sublayer.
OutNUcastPkts	The total number of packets that higher-level protocols requested be transmitted, and that are addressed to a multicast or broadcast address at this sublayer, including those that are discarded or not sent.
InMulticastPkts	The number of packets delivered by this sublayer to a higher sublayer that were addressed to a multicast address at this sublayer. For a MAC layer protocol, this number includes both group and functional addresses.
OutMulticastPkts	The number of packets that higher-level protocols requested be transmitted, and that are addressed to a multicast address at this sublayer, including those that were discarded or not sent. For a MAC layer protocol, this number includes both group and functional addresses.
InBroadcastPkts	The number of packets delivered by this sublayer to a higher sublayer that are addressed to a broadcast address at this sublayer.
OutBroadcastPkts	The number of packets that higher-level protocols requested be transmitted, and that were addressed to a broadcast address at this sublayer, including those that were discarded or not sent.
InDiscards	The number of inbound packets chosen to be discarded even though no errors were detected to prevent their being deliverable to a higher-layer protocol. One possible reason for discarding such a packet can be to free up buffer space.
OutDiscards	The number of outbound packets chosen to be discarded even though no errors were detected to prevent their being transmitted. One possible reason for discarding such a packet can be to free up buffer space.
InErrors	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.
OutErrors	For packet-oriented interfaces, the number of outbound packets that cannot be transmitted because of errors. For character-oriented or fixed-length interfaces, the number of outbound transmission units that cannot be transmitted because of errors.
InUnknownProtos	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 zero.

Graphing port Ethernet error statistics using EDM

Use this procedure to display and graph Ethernet error statistics.

Procedure steps

1. On the Device Physical View, click a port.
2. From the navigation tree, double-click **Graph**.
3. In the Graph tree, double-click **Port**.
4. In the work area, click the **Ethernet Errors** tab.
5. On the toolbar, select a **Poll Interval** from the list.
6. On the toolbar, you can click **Clear Counters** to reset the IP statistics counters.
7. To select statistics to graph, click a statistic type row under a column heading.
8. On the toolbar, click **Line Chart**, **Area Chart**, **Bar Chart**, or **Pie Chart**.

Variable definitions

Use the data in the following table to help you understand the Ethernet error statistics.

Variable	Value
AlignmentErrors	A count of frames received on a particular interface that are not an integral number of octets in length and do not pass the FCS check. The count represented by an instance of this object is incremented when the AlignmentError status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions occur are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.
FCSErrors	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 is incremented when the FCSErrors 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.
InternalMacTransmitErrors	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.
InternalMacReceiveErrors	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.

Table continues…

Variable	Value
	The precise meaning of the count represented by an instance of this object is implementation specific. In particular, an instance of this object can represent a count of receive errors on a particular interface that are not otherwise counted.
CarrierSenseErrors	The number of times that the carrier sense condition was lost or never asserted when attempting to transmit a frame on a particular interface. The count represented by an instance of this object is incremented at most once for each transmission attempt, even if the carrier sense condition fluctuates during a transmission attempt.
FrameTooLongs	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 is incremented when the FrameTooLongs 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.
SQETestErrors	A count of times that the SQE Test Errors message is generated by the PLS sublayer for a particular interface. The SQE TEST ERROR message is defined in section 7.2.2.2.4 of ANSI/IEEE 802.3-1985 and its generation is described in section 7.2.4.6 of the same document.
DeferredTransmissions	A count of frames for which the first transmission attempt on a particular interface is delayed because the medium is busy. The count represented by an instance of this object does not include frames involved in collisions.
SingleCollisionFrames	A count of successfully transmitted frames on a particular interface for which transmission is inhibited by exactly one collision. A frame that is counted by an instance of this object is also counted by the corresponding instance of either the ifOutUcastPkts, ifOutMulticastPkts, or ifOutBroadcastPkts, and is not counted by the corresponding instance of the MultipleCollisionFrames object.
MultipleCollisionFrames	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 ifOutUcastPkts, ifOutMulticastPkts, or ifOutBroadcastPkts, and is not counted by the corresponding instance of the SingleCollisionFrames object.
LateCollisions	The number of times that a collision is detected on a particular interface later than 512 bit-times into the transmission of a packet. Five hundred and twelve 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.
ExcessiveCollisions	A count of frames for which transmission on a particular interface fails due to excessive collisions.

Graphing port RMON statistics using EDM

Use this procedure to display and graph RMON Ethernet statistics.

Procedure steps

1. On the Device Physical View, click a port.
2. From the navigation tree, double-click **Graph**.
3. In the Graph tree, double-click **Port**.
4. Click the **Rmon** tab.
5. On the toolbar, select a **Poll Interval** from the list.
6. On the toolbar, you can click **Clear Counters** to reset the IP statistics counters.
7. To select statistics to graph, click a statistic type row under a column heading.
8. On the toolbar, click **Line Chart**, **Area Chart**, **Bar Chart**, or **Pie Chart**.

Variable definitions

Use the data in the following table understand RMON Ethernet statistics.

Variable	Value
Octets	The total 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. For greater precision, sample the etherStatsPkts and etherStatsOctets objects before and after a common interval.
Pkts	The total number of packets (including bad packets, broadcast packets, and multicast packets) received.
BroadcastPkts	The total number of good packets received that are directed to the broadcast address. This does not include multicast packets.
MulticastPkts	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.
CRCAlignErrors	The total number of packets received with a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, but with either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).
UndersizePkts	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.
OversizePkts (>1518)	The total number of packets received that are longer than 1518 octets (excluding framing bits but including FCS octets) and were otherwise well formed.
Fragments	The total number of packets received that are less than 64 octets in length (excluding framing bits but including FCS octets) and with either a bad FCS

Table continues…

Variable	Value
	with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). For etherStatsFragments to increment is normal because it counts both runts (which are normal occurrences due to collisions) and noise hits.
Collisions	The best estimate of the total number of collisions on this Ethernet segment.
Jabbers	The total number of packets received that are longer than 1518 octets (excluding framing bits, but including FCS octets), with either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error). Jabber is defined as the condition where any packet exceeds 20 ms. The allowed range to detect jabber is between 20 ms and 150 ms.
1..64	The total number of packets (including bad packets) received that are less than or equal to 64 octets in length (excluding framing bits but including FCS octets).
65 ..127	The total number of packets (including bad packets) received that are greater than 64 octets in length (excluding framing bits but including FCS octets).
128 ..255	The total number of packets (including bad packets) received that are greater than 127 octets in length (excluding framing bits but including FCS octets).
256..511	The total number of packets (including bad packets) received that are greater than 255 octets in length (excluding framing bits but including FCS octets).
512..1023	The total number of packets (including bad packets) received that are greater than 511 octets in length (excluding framing bits but including FCS octets).
1024..1518	The total number of packets (including bad packets) received that are greater than 1023 octets in length (excluding framing bits but including FCS octets).

Graphing miscellaneous port statistics using EDM

Use this procedure to display and graph miscellaneous statistics for a switch port.

Procedure steps

1. On the Device Physical View, click a port.
2. From the navigation tree, double-click **Graph**.
3. In the Graph tree, double-click **Port**.
4. In the work area, click the **Misc.** tab.
5. On the toolbar, select a **Poll Interval** from the list.
6. On the toolbar, you can click **Clear Counters** to reset the IP statistics counters.
7. To select statistics to graph, click a statistic type row under a column heading.

8. On the toolbar, click **Line Chart**, **Area Chart**, **Bar Chart**, or **Pie Chart**.

Variable definitions

Use the data in the following table to help you understand miscellaneous port statistics.

Variable	Value
NoResourcesPktsDropped	The number of packets dropped due to switch memory shortage.

Viewing SFP DDI information

About this task

Use the following procedure to view SFP DDI information.

Procedure

1. From the navigation pane, double-click **Edit**.
2. In the Edit tree, double-click **DDI**.
3. In the work area, click Digital Diagnostic Interface.

   Each row represents a different supported SFP, SFP+, or GBIC transceiver connected to the switch.

Digital Diagnostic Interface field descriptions

Use the data in the following table to use the Digital Diagnostic Interface tab.

Name	Description
IfIndex	Indicates the interface index.
Calibration	Indicates if the calibration is internal or external.
RXPowerMeasurement	Indicates Rx power measurement as average or OMA.

Viewing SFP DDI information details

About this task

Use the following procedure to view detailed SFP DDI information.

Procedure

1. From the navigation pane, double-click **Edit**.
2. In the Edit tree, double click **DDI**.
3. In the work area, click Digital Diagnostic Interface.

Each row represents a different supported SFP, SFP+, or GBIC transceiver connected to the switch.

4. To select a device for which to view detailed DDI information, click on a device row.
5. On the toolbar, click the **Detail** button.

DDI Details field descriptions

Use the data in the following table to use the **DDI Detail**.

Name	Description
Temp (C)	Indicates the Temperature information in degrees Celsius of the SFP or SFP+. The values are: • HighAlarmThreshold— Indicates the high alarm threshold in degrees Celsius. • LowAlarmThreshold— Indicates the low alarm threshold in degrees Celsius. • HighWarnThreshold— Indicates the high warning threshold in degrees Celsius. • LowWarnThreshold— Indicates the low warning threshold in degrees Celsius. • Value— Indicates the current temperature in degrees Celsius of the SFP or SFP+.
Voltage (V)	Indicates the Voltage information in volts of the SFP or SFP+. The values are: • HighAlarmThreshold— Indicates the high alarm threshold in volts. • LowAlarmThreshold— Indicates the low alarm threshold in volts. • HighWarnThreshold— Indicates the high warning threshold in volts. • LowWarnThreshold— Indicates the low warning threshold in volts. • Value— Indicates the current voltage in volts of the SFP or SFP+.
Bias (mA)	Indicates the Bias information in mA of the SFP or SFP+. The values are: • HighAlarmThreshold— Indicates the high alarm threshold in mA.

Table continues…

Name	Description
	• LowAlarmThreshold— Indicates the low alarm threshold in mA. • HighWarnThreshold— Indicates the high warning threshold in mA. • LowWarnThreshold— Indicates the low warning threshold in mAs. • Value— Indicates the current Bias mA of the SFP or SFP+.
TX Power (dBm)	Indicates the TX Power in dBm of the SFP or SFP+. The values are: • HighAlarmThreshold Indicates the high alarm threshold in dBm. • LowAlarmThreshold Indicates the low alarm threshold in dBm. • HighWarnThreshold Indicates the high warning threshold in dBm. • LowWarnThreshold Indicates the low warning threshold in dBm. • Value Indicates the current TX Power in dBm of the SFP or SFP+.
RX Power (dBm)	Indicates the RX Power in dBm of the SFP or SFP+. The values are: • HighAlarmThreshold Indicates the high alarm threshold in dBm. • LowAlarmThreshold Indicates the low alarm threshold in dBm. • HighWarnThreshold Indicates the high warning threshold in dBm. • LowWarnThreshold Indicates the low warning threshold dBm. • Value Indicates the current RX Power in dBm of the SFP or SFP+.

Viewing MACsec interface statistics

Procedure

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

2. In the navigation tree, click **Edit > Chassis**.
3. Click the **MACsec Interface Stats** tab.

Note:

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

Field definitions

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

Name	Definition
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 not 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), or invalid Integrity Check Value (ICV).
RxUnknownSCIPkts	Specifies the number of packets received with an unknown secure channel identifier (SCI), and with MACsec not 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.

Viewing secure channel (SC) inbound statistics

Procedure

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

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

Note:

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

Field definitions

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

Name	Definition
UnusedSAPkts	Specifies the summary of received unencrypted packets on all SAs of this secure channel, with MACsec not in strict mode.
NoUsingSAPkts	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.
LatePkts	Specifies the number of packets received that have been discarded for this secure channel (SC) with Replay Protect enabled. **Note:** The current release does not support Replay Protect
NotValidPkts	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: • MACsec was operating in strict mode. • The packets received were encrypted but contained erroneous fields.
InvalidPkts	Specifies the summary of all packets received that were not valid for this SC, with MACsec operating in check mode.
DelayedPkts	Specifies the summary of packets for this SC, with the packet number (PN) of the packets lower than the lower bound replay protection PN. **Note:** The current release does not support Replay Protect.
UncheckedPkts	The total number of packets for this SC that: • Were encrypted and had failed the integrity check.

Table continues…

Name	Definition
	• Were not encrypted and had failed the integrity check. • Were received when MACsec validation was not enabled.
OKPkts	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.

Viewing secure channel (SC) outbound statistics

Procedure

1. In the Device Physical View tab, select the port for which you need to view the SC outbound statistics
2. In the navigation tree, expand the following folders: **Edit > Port > General**.
3. Click the **SC Outbound Stats** tab.

Note:

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

Field definitions

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

Name	Definition
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.

Chapter 5: IP Flow Information Export

This chapter provides conceptual information and procedures to configure IP Flow Information Export (IPFIX).

IP Flow Information Export

With IP Flow Information Export (IPFIX) you can monitor traffic flows by configuring observation points to collect flow statistics over a designated time period.

IPFIX supports the following external IPFIX collectors:

- NetQoS Harvester/Collector
- IP Flow Manager
- Fluke Collector

IP traffic is sampled and classified into various flows based on the following parameters:

- protocol type
- destination IP address
- source IP address
- ingress port
- type of service (TOS)

You cannot use IPFIX on secondary interfaces.

If the protocol type is TCP or UDP, a flow is defined by the following two additional parameters:

- source port
- destination port

IPFIX supports the following:

- the creation and display of sampled information
- the ability to export this sampled information

Note:

IPFIX also monitors IGMP traffic.

The IPFIX feature shares resources with QoS. If the IPFIX feature is enabled, a QoS policy precedence is used. For further information about QoS policies, see Configuring Quality of Service on Ethernet Routing Switch 4900 and 5900 Series .

IPFIX configuration

This section provides procedures to configure IP Flow Information Export (IPFIX) using CLI.

Global IPFIX management using CLI

Use the information in this section to enable or disable IPFIX globally on a switch or stack.

Enabling IPFIX globally

About this task

Use this procedure to enable IPFIX globally for a switch or stack.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enable IPFIX:

   ```
   ip ipfix enable
   ```

Disabling IPFIX globally

About this task

Use this procedure to disable IPFIX globally for a switch or stack.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Disable IPFIX:

   ```
   no ip ipfix enable
   ```

   OR

   ```
   default ip ipfix enable
   ```

Viewing the global IPFIX status

About this task

Use this procedure to display the global IPFIX operational status for a switch or stack.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Display IPFIX status:

   ```
   show ip ipfix
   ```

Example

```
Switch>enable
Switch#show ip ipfix
IPFIX   Disabled
```

IPFIX flow management

Use the information in this section to configure and manage IPFIX flow for a standalone switch or a switch in a stack.

Configuring the IPFIX aging interval

About this task

Use this procedure to configure the IPFIX flow record aging interval for a standalone switch or a switch in a stack.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure the IPFIX aging interval:

   ```
   ip ipfix slot <unit_number> aging-interval <0-2147400>
   ```

Variable definitions

Use the data in the following table to use the **`ip ipfix slot <unit_number> aging-interval`** command.

Variable	Value
`<unit_number>`	Specifies whether the switch is a standalone or part of a stack. A value of 1 indicates a standalone switch.

Table continues…

Variable	Value
<0-2147400>	Specifies the aging interval of the flow record in seconds. Values range from 0–2147400 seconds. Aging time is the period of time in which all records are verified if they are updated. If no new updates are found between two checks, the system deletes the records.

Changing the IPFIX aging interval to default

About this task

Use this procedure to change the IPFIX flow record aging interval to the default value of 30 seconds for a standalone switch or a switch in a stack.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Change the IPFIX aging interval to default:

   ```
   default ip ipfix slot <unit_number> aging-interval
   ```

Enabling the IPFIX exporter

About this task

Use this procedure to enable the IPFIX exporter for a standalone switch or a switch stack.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enable the IPFIX exporter:

   ```
   ip ipfix exporter-enable
   ```

   OR

   ```
   default ip ipfix exporter-enable
   ```

Disabling the IPFIX exporter

About this task

Use this procedure to disable the IPFIX exporter for a standalone switch or a switch stack.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Disable the IPFIX exporter:

```
no ip ipfix exporter-enable
```

Configuring the IPFIX export interval

About this task

Use this procedure to configure the IPFIX export interval for a standalone switch or a switch stack.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure the IPFIX export interval:

```
ip ipfix export-interval <10-3600>
```

Variable definitions

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

Variable	Value
`<10-3600>`	Specifies the frequency of data exports to the collector in seconds. Values range from 10 to 3600 seconds. The default is 50 seconds.

Changing the IPFIX export interval to default

About this task

Use this procedure to change the IPFIX export interval for a standalone switch or a switch stack to the default value of 50 seconds.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Change the IPFIX export interval to default:

```
default ip ipfix export-interval
```

Configuring the IPFIX refresh interval template

About this task

Use this procedure to configure the IPFIX refresh interval template for a standalone switch or a switch stack.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure the IPFIX refresh interval template:

   ```
   ip ipfix template-refresh-interval <300-3600>
   ```

Variable definitions

Use the data in the following table to use the **`ip ipfix template-refresh-interval`** command.

Variable	Value
`<300-3600>`	Specifies the refresh timeout interval template in seconds. Values range from 300 to 3600. The default is 1800 seconds. The template is sent out to the collector either at the configured interval or after the specified template packets refresh number is reached, whichever occurs first. The template is also sent out to the collector when globally enabling IPFIX.

Changing the IPFIX refresh interval template to default

About this task

Use this procedure to change the IPFIX refresh interval template for a standalone switch or a switch stack to the default value of 1800 seconds. The template is sent out to the collector either at the configured interval or after the specified template packets refresh number is reached, whichever occurs first.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Change the IPFIX refresh interval template to default:

   ```
   default ip ipfix template-refresh-interval
   ```

Configuring the IPFIX refresh packets template

About this task

Use this procedure to configure the IPFIX refresh packets template for a standalone switch or a switch stack.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure the IPFIX refresh packets template:

   ```
   ip ipfix template-refresh-packets <10000-100000>
   ```

Variable definitions

Use the data in the following table to use the **`ip ipfix template-refresh-packets`** command.

Variable	Value
`<10000-100000>`	Specifies the refresh packets template limit in numbers of packets. Values range from 10000 to 100000 packets. The default is 10000 packets. The template is sent out to the collector either after the configured template packets refresh number is reached or at the specified refresh interval, whichever occurs first. The template is also sent out to the collector when globally enabling IPFIX.

Changing the IPFIX refresh packets template to default

About this task

Use this procedure to change the IPFIX refresh packets template for a standalone switch or a switch stack to the default value of 10000 packets.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Change the IPFIX refresh packets template to default:

   ```
   default ip ipfix template-refresh-packets
   ```

Viewing IPFIX flow information

About this task

Use this procedure to display configured IPFIX flow information.

Procedure

View IPFIX flow information:

```
show ip ipfix slot <unit_number>
```

Example

The following example displays the IPFIX information for Slot 1.

```
Switch#show ip ipfix slot 1
Slot 1
-------
        Aging Interval(sec)      25
        Active Timeout(min)      30
        Export Interval(sec)     50
        Export State             enabled
        Template Refresh(sec)    1800
        Template Refresh(pkts)   10000
```

Variable definitions

The following table defines parameters that you enter with the **`show ip ipfix slot <unit_number>`** command.

Variable	Value
`slot <unit_number>`	Displays information for a switch that is a standalone or part of a stack. A value of 1 indicates a standalone switch.

IPFIX collector management using CLI

Use the information in this section to enable or disable IPFIX collectors, and to display configured IPFIX collector configuration information.

Enabling an IPFIX collector

About this task

Use this procedure to enable an IPFIX collector for a standalone switch or a switch stack.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enable an IPFIX collector:

```
ip ipfix collector <A.B.C.D> enable
```

OR

```
default ip ipfix collector <A.B.C.D> enable
```

Variable definitions

Use the data in the following table to use the **ip ipfix collector <A.B.C.D> enable** or the **default ip ipfix collector <A.B.C.D> enable** command.

Variable	Value
<A.B.C.D>	Specifies the IPFIX collector IP address.

Disabling an IPFIX collector

About this task

Use this procedure to disable an IPFIX collector for a standalone switch or a switch stack.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Disable an IPFIX collector:

   ```
   no ip ipfix collector <A.B.C.D> enable
   ```

Variable definitions

Use the data in the following table to use the **no ip ipfix collector <A.B.C.D> enable** command.

Variable	Value
<A.B.C.D>	Specifies the IPFIX collector IP address.

Viewing the IPFIX collector information

About this task

Use this procedure to display IPFIX collector configuration information for a standalone switch or a switch stack.

Procedure

1. View information for all configured IPFIX collectors:

   ```
   show ip ipfix collector
   ```

2. View information for a specific configured IPFIX collector:

   ```
   show ip ipfix collector <A.B.C.D>
   ```

Variable definitions

Use the data in the following table to use the **show ip ipfix collector <A.B.C.D>** command.

Variable	Value
<A.B.C.D>	Displays the operational status for a specific IPFIX collector IP address.

Port IPFIX management using CLI

Use the information in this section to enable or disable IPFIX for one or more switch ports on a standalone switch or a switch that is part of a stack.

Enabling port-based IPFIX for a standalone switch

About this task

Use this procedure to enable IPFIX for one or more ports on a standalone switch.

Procedure

1. Enter Interface Configuration mode:

   ```
   enable
   configure terminal
   interface ethernet <port number>
   ```

2. Enable IPFIX for the selected port or ports:

   ```
   ip ipfix enable
   ```

3. Enable IPFIX for alternate ports:

   ```
   ip ipfix port <port_list> enable
   ```

Variable definitions

Use the data in the following table to use the **ip ipfix port <port_list> enable** command.

Variable	Value
port <port_list>	Specifies an individual port or list of ports.

Disabling port-based IPFIX for a standalone switch

About this task

Use this procedure to disable IPFIX for one or more ports on a standalone switch.

Procedure

1. Enter Interface Configuration mode:

```
enable
configure terminal
interface ethernet <port number>
```

2. Disable port-based IPFIX for a standalone switch:

```
no ip ipfix [enable] [port <port_list> enable]
```

Variable definitions

Use the data in the following table to use the **`no ip ipfix [enable] [port <port_list> enable]`** command.

Variable	Value
`enable`	Disables IPFIX for the selected port or ports.
`port <port_list> enable`	Disables IPFIX for an alternate individual port or list of ports.

Changing port-based IPFIX for a standalone switch to default

About this task

Use this procedure to change the IPFIX operational status to default for one or more ports on a standalone switch.

Procedure

1. Enter Interface Configuration mode:

```
enable
configure terminal
interface ethernet <port number>
```

2. Change the IPFIX operational status for one or more ports:

```
default ip ipfix [enable] [port <port_list> enable]
```

3. Enable IPFIX for alternate ports:

```
ip ipfix port <port_list> enable
```

Variable definitions

Use the data in the following table to use the **`default ip ipfix [enable] [port <port_list> enable]`** command.

Variable	Value
`enable`	Changes the IPFIX operational status for the selected port or ports to default.
`port <port_list> enable`	Changes the IPFIX operational status for a port or list of ports to default.

Viewing the port-based IPFIX status for a standalone switch

About this task

Use this procedure to display the IPFIX operational status for one or more ports on a standalone switch.

Procedure

1. Display the IPFIX operational status for all switch ports:

   ```
   show ip ipfix interface
   ```

2. Display the IPFIX operational status for specific switch ports:

   ```
   show ip ipfix interface <port_list>
   ```

Example

```
Switch>show ip ipfix interface

Port    IPFIX
----    -----
1       Disable
2       Disable
3       Disable
4       Disable
5       Disable
6       Disable
7       Disable
8       Disable
9       Disable
10      Disable
11      Disable
12      Disable
13      Disable
14      Disable
15      Disable
16      Disable
17      Disable
18      Disable
19      Disable
----More (q=Quit, space/return=Continue)----
```

Variable definitions

Use the data in the following table to use the **`show ip ipfix interface <port_list>`** command.

Variable	Value
`<port_list>`	Specifies a specific port or list of ports for which to display the IPFIX operational mode for to default.

Enabling port-based IPFIX for a stack switch

About this task

Use this procedure to enable IPFIX for one or more ports on a switch that is part of a stack.

Procedure

1. Enter Interface Configuration mode:

   ```
   enable
   configure terminal
   interface ethernet <port number>
   ```

2. Enable IPFIX for the selected port or ports:

   ```
   ip ipfix enable
   ```

3. Enable IPFIX for alternate ports:

   ```
   ip ipfix port <unit_number/port_list> enable
   ```

Variable definitions

Use the data in the following table to use the **`ip ipfix port <unit_number/port_list> enable`** command.

Variable	Value
`port <unit_number/port_list>`	Specifies switch number in the stack and an individual port or list of ports.

Disabling port-based IPFIX for a stack switch

About this task

Use this procedure to disable IPFIX for one or more ports on a switch that is part of a stack.

Procedure

1. Enter Interface Configuration mode:

   ```
   enable
   configure terminal
   interface ethernet <port number>
   ```

2. Disable IPFIX for the selected port or ports:

   ```
   no ip ipfix enable
   ```

3. Disable IPFIX for alternate ports:

   ```
   no ip ipfix port <unit_number/port_list> enable
   ```

Variable definitions

Use the data in the following table to use the **`ip ipfix port <unit_number/port_list> enable`** command.

Variable	Value
`port <unit_number/port_list>`	Specifies a switch number in the stack and an individual port or list of ports.

Changing port-based IPFIX for a stack switch to default

About this task

Use this procedure to change the IPFIX operational status to default for one or more ports on a switch that is part of a stack.

Procedure

1. Enter Interface Configuration mode:

   ```
   enable
   configure terminal
   interface ethernet <port number>
   ```

2. Change the IPFIX operational status for the selected port or ports:

   ```
   default ip ipfix enable
   ```

3. Change the IPFIX operational status for alternate port or ports:

   ```
   default ip ipfix port <unit_number/port_list> enable
   ```

Variable definitions

Use the data in the following table to use the **`default ip ipfix port <unit_number/port_list> enable`** command.

Variable	Value
`port <unit_number/port_list>`	Specifies a switch number in the stack and an individual port or list of ports.

Viewing the port-based IPFIX status for a stack switch

About this task

Use this procedure to display the IPFIX operational status for one or more ports on a standalone switch.

Procedure

1. Display the IPFIX operational status for all ports in the stack:

   ```
   show ip ipfix interface
   ```

2. Display the IPFIX operational status for specific ports in the stack:

   ```
   show ip ipfix interface <unit_number/port_list>
   ```

Variable definitions

Use the data in the following table to use the **show ip ipfix interface <unit_number/port_list>** command.

Variable	Value
<unit_number/port_list>	Specifies a switch number in the stack and an individual port or list of ports.

Viewing the IPFIX table

About this task

Use this procedure to sort and display IPFIX statistics for a standalone switch or a switch stack.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. To view IPFIX table, enter the following command:

   ```
   show ip ipfix table <unit_number> [sort-by <sort_rule> [sort-order <sort_order>] [display <num_entries>]
   ```

3. To display information about the total number of streams learned by IPFIX on each stack unit, enter the following command:

   ```
   show ip ipfix table active-flows
   ```

Example

```
Switch>show ip ipfix table sort-by byte-count sort-order ascending display all
 Please wait while retrieving data...

   HW SA           IP SA          TOS      L4 SRC   PACKET        TIME FIRST
      DA              DA       PROTOCOL      DST    BYTE             LAST
    VLAN                      TCP FLAGS                          UNIT/PORT
------------ --------------- ----------- ------ ---------- --------------------
00001ee819fb 141.005.002.007           0      0         42     0 days, 00:53:56
0019e1533c9d   0.000.000.000         TCP      0       2688     0 days, 00:54:34
    200                              00                      1/24
Only 1 entries found
```

Variable definitions

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

Variable	Value
display <num_entries>	Specifies the number of entries to display. Values include: • all—displays all available entries

Table continues…

Variable	Value
	• top-10—displays first 10 entries • top-25—displays first 25 entries • top-50—displays first 50 entries • top-100—displays first 100 entries • top-200—displays first 200 entries
`sort-by <sort_rule>`	Specifies a rule to sort data by. Values include: • byte-count—data byte number • dest-addr—destination IP address • first-pkt-time—first packet time • last-pkt-time—last packet time • pkt-count—packet number • port—port number • protocol—protocol number • source-addr—source IP address • TCP-UDP-dest-port—TCP/UDP destination port • TCP-UDP-src-port—TCP/UDP source port • TOS—type of service
`sort-order <sort_order>`	Specifies the order in which to sort data. Values include: • ascending • descending
`<unit_number>`	Specifies whether the switch is a standalone or part of a stack. A value greater than 1 indicates the switch location in a stack. For a standalone unit, do not specify the `unit_number` parameter.
`active-flows`	Displays information about the total number of streams learned by IPFIX on each stack unit.

IPFIX configuration using Enterprise Device Manager

This section provides procedures to configure IP Flow Information Export (IPFIX) using Enterprise Device Manager (EDM).

Prerequisites

- Open one of the supported browsers.
- Enter the IP address of the switch to open an EDM session.

Configuring IPFIX globally using EDM

Use the following procedure to enable or disable IPFIX for the switch. IPFIX is disabled by default.

Prerequisites

- Open one of the supported browsers.
- Enter the IP address of the switch to open an EDM session.

Procedure steps

1. From the navigation tree, double-click **Serviceability**.
2. In the Serviceability tree, double-click **IPFIX**.
3. In the work area, click the **Global** tab.
4. In the State section, click the **enable** radio button to enable IPFIX globally.

   OR

   Click the **disable** radio button to disable IPFIX globally.
5. Click **Apply**.

Configuring IPFIX flows using EDM

Use the following procedure to configure export flow information sources.

Procedure steps

1. From the navigation tree, double-click **Serviceability**.
2. In the Serviceability tree, double-click **IPFIX**.
3. In the work area, click the **Exporters** tab.
4. To select an exporter to edit, click the exporter slot number.
5. In the exporter row, double-click the cell in the **AgingIntv** column.
6. Type a value in the dialog box.
7. In the exporter row, double-click the cell in the **ExportState** column.

8. In the exporter row, double-click the cell in the **ExportIntv** column.
9. Select a value from the list.
10. In the exporter row, double-click the cell in the **TempRefIntvSec** column.
11. Type a value in the dialog box.
12. In the exporter row, double-click the cell in the **TempRefIntvPkts** column.
13. Type a value in the dialog box.
14. Click **Apply**.

Variable definitions

Variable	Value
Slot (Unit)	Identifies the switch that is exporting IPFIX flows. This value corresponds to the unit number in a stack or is the number 1 for a stand-alone unit.
AgingIntv	Specifies the aging interval of the flow record in seconds. Values range from 0–2147400 seconds. Aging time is the period of time in which all records are verified if they are updated. The records are deleted if no new updates are found between two checks.
ActiveTimeout	Indicates the flow record active timeout value in minutes. This is a read-only cell.
ExportIntv	Specifies the frequency of data exports to the collector in seconds. Values range from 10 to 3600 seconds.
ExportState	Enables or disables the exporter.
TempRefIntvSec	Specifies the template refresh timeout in seconds. Values range from 300 to 3600. The template is sent out to the collector either at the configured interval or after the specified template packets refresh number is reached, whichever occurs first.
TempRefIntvPkts	Specifies the template refresh timeout in numbers of packets. Values range from 10000 to 100000 packets. The template is sent out to the collector either after the configured template packets refresh number is reached or at the specified refresh interval, whichever occurs first.

IPFIX collector management using EDM

Use the information in this section to display configured IPFIX collector information and to modify IPFIX collector configurations.

Viewing IPFIX collectors using EDM

Use the following procedure to display collected and analyzed data exported from an IPFIX-compliant switch.

Procedure steps

1. From the navigation tree, double-click **Serviceability**.
2. In the Serviceability tree, double-click **IPFIX**.
3. In the work area, click the **Collectors** tab.

Variable definitions

Variable	Value
Slot (Unit)	Identifies the switch that is collecting and analyzing data. This value corresponds to the unit number in a stack or is the number 1 for a stand-alone unit.
AddressType	Indicates the IP address type of the collector. Currently only IPv4 addresses are supported.
Address	Indicates the IP address of the collector.
Protocol	Indicates the protocol used to transport the IPFIX data to the collector. Currently only the UDP protocol is supported for this task.
DestPort	Indicates the port on which the collector is listening for IPFIX data. Currently only port 9995 is supported.
ProtoVer	Indicates the format in which IPFIX data is provided to the collector. Currently only Netflow version 9 formatting is supported.
Enable	Indicates the operational state of this collector.

Configuring IPFIX collectors using EDM

Use the following procedure to collect and analyze data exported from an IPFIX-compliant switch.

Procedure steps

1. From the navigation tree, double-click **Serviceability**.
2. In the Serviceability tree, double-click **IPFIX**.
3. In the work area, click the **Collectors** tab.
4. Click the **Insert**.

5. In the Slot dialog box, type a value.
6. In the Address dialog box, type an IP address.
7. Select the **Enable** check box to enable the collector.

   OR

   Clear the **Enable** check box to disable the collector.
8. Click **Apply**.

Variable definitions

Variable	Value
Slot (Unit)	Identifies the switch that is collecting and analyzing data. This value corresponds to the unit number in a stack or is the number 1 for a stand-alone unit.
AddressType	Specifies the IP address type of the collector. Currently only IPv4 addresses are supported.
Address	Specifies the IP address of the collector.
Protocol	Specifies the protocol used to transport the IPFIX data to the collector. Currently only the UDP protocol is supported for this task.
DestPort	Specifies the port on which the collector is listening for IPFIX data. Currently only port 9995 is supported.
ProtoVer	Specifies the format in which IPFIX data is provided to the collector. Currently only Netflow version 9 formatting is supported.
Enable	Enables or disables the collector.

Deleting IPFIX collectors using EDM

Use the following procedure to display collected and analyzed data exported from an IPFIX-compliant switch.

Procedure steps

1. From the navigation tree, double-click **Serviceability**.
2. In the Serviceability tree, double-click **IPFIX**.
3. In the work area, click the **Collectors** tab.
4. To select an collector to delete, click the collector slot number.
5. Click **Delete**.

IPFIX port management using EDM

Use the information in this section to view and modify IPFIX port configurations.

Viewing IPFIX port information using EDM

Use the following procedure to display IPFIX port configuration information.

Procedure steps

1. From the navigation tree, double-click **Serviceability**.
2. In the Serviceability tree, double-click **IPFIX**.
3. In the work area, click the **Ports** tab.

Variable definitions

Variable	Value
Id	Indicates the individual port on which the IPFIX parameters are being configured. Ports are itemized in the Unit/Port format.
Flush	Indicates the flushing action to take on the port. Flushing the port of data involves deleting all previously gathered information about that port. Values include: • none—the port data is not flushed. • flush—the port data is flushed, which deletes the data from switch memory. • exportAndFlush—the port data is exported to a configured collector and the data is then flushed.
AllTraffic	Indicates if IPFIX data is collected on this port. • enable—IPFIX data is collected • disable—IPFIX data is not collected

Modifying specific IPFIX port configurations using EDM

Use the following procedure to modify IPFIX configuration parameters for specific ports.

Procedure steps

1. From the navigation tree, double-click **Serviceability**.
2. In the Serviceability tree, double-click **IPFIX**.
3. In the work area, click the **Ports** tab.
4. In the port row, double-click the cell in the **Flush** column.
5. Select a value from the list.
6. In the port row, double-click the cell in the **AllTraffic** column.
7. Select a value from the list.
8. Repeat steps **4** through **8** to modify additional ports.

9. Click **Apply**.

Variable definitions

Variable	Value
Id	Specifies the individual port on which the IPFIX parameters are being configured. Ports are itemized in the Unit/Port format.
Flush	Specifies the flushing action to take on the port. Flushing the port of data involves deleting all previously gathered information about that port. Values include: • none—the port data is not flushed. • flush—the port data is flushed, which deletes the data from switch memory. • exportAndFlush—the port data is exported to a configured collector and the data is then flushed.
AllTraffic	Specifies if IPFIX data is collected on this port. • enable—IPFIX data is collected • disable—IPFIX data is not collected

Modifying IPFIX port configurations

Use the following procedure to modify the IPFIX configuration parameters for all available ports.

Procedure

1. Follow one of the following paths:
   - From the **Device Physical View**, select a port, or use Ctrl-click to select more than one port, right-click **Edit** then click the **IPFIX** tab.
   - From the **Device Physical View**, select a port, or use Ctrl-click to select more than one port, then follow the navigation tree to **Edit > Chassis > Ports > IPFIX** tab.
   - In the navigation tree, go to **Serviceability > IPFIX > Ports** tab.
2. In the port row, double-click the cell in the column to be modified and configure as required from a drop-down list or by typing a value.
3. Repeat for additional cells.
4. Repeat the above steps to configure IPFIX for additional ports.
5. Click **Apply**.
6. On the toolbar, you can click **Refresh** to update the work area data display.

Variable definitions

Variable	Value
Flush	Specifies the flushing action to take on the port. Flushing the port of data involves deleting all previously gathered information about that port. Values include: • none—the port data is not flushed. • flush—the port data is flushed, which deletes the data from switch memory. • exportAndFlush—the port data is exported to a configured collector and the data is then flushed.
AllTraffic	Specifies if IPFIX data is collected on this port. • enable—IPFIX data is collected • disable—IPFIX data is not collected

Modifying all IPFIX port configurations using EDM

Use the following procedure to modify the IPFIX configuration parameters for all available ports.

Procedure steps

1. From the navigation tree, double-click **Serviceability**.
2. In the Serviceability tree, double-click **IPFIX**.
3. In the work area, click the **Ports** tab.
4. To configure port-based IPFIX, double-click editable table cells in a particular port row as required.
5. Click **Apply**.
6. Optionally, to configure parameters for multiple ports, you can use the Make Selection section as below.
7. In the work area, in the Make Selection section of the Multiple Port Configuration pane, click the Switch/Stack/Ports ellipsis (...) to open the Port Editor dialog
8. In the Port Editor window, click the ports you want to configure.

   Note:

   If you want to configure all ports, click **All**.
9. Click **OK** to return to the Make Selection pane.

   The ports you selected appear in the Switch/Stack/Ports box.

10. To change the configuration of the selected ports, in the Multiple Port Configuration pane, double-click the cell beneath the column heading that represents the parameter you want to change and do one of the following:
    - Select a value from a drop-down list.
    - Type a value in the cell.
11. In the Make Selection pane, click **Apply Selection**.

    The changes appear in the table.
12. (Optional) Click **Clear Selection** to clear Multiple Port Configurations or click **Hide Non-Editable** to display only those parameters that are editable in the Multiple Port Configuration pane for the selected ports.

Note:

You can also display and manage IPFIX for ports from the navigation tree path **Edit > Chassis > Ports >IPFIX**, after selecting multiple ports from **Device Physical View**.

Variable definitions

Variable	Value
Flush	Specifies the flushing action to take on the port. Flushing the port of data involves deleting all previously gathered information about that port. Values include: • none—the port data is not flushed. • flush—the port data is flushed, which deletes the data from switch memory. • exportAndFlush—the port data is exported to a configured collector and the data is then flushed.
AllTraffic	Specifies if IPFIX data is collected on this port. • enable—IPFIX data is collected • disable—IPFIX data is not collected

Displaying IPFIX data information using EDM

Use this procedure to set the display criteria and display IPFIX data information.

Procedure steps

1. From the navigation tree, double-click **Serviceability**.
2. In the Serviceability tree, click **IPFIX**.
3. In the work area, click the **Data Information** tab.

Variable definition

Name	Description
Unit Number	Specifies a standalone switch or a switch that is part of a stack. For a standalone switch, use a value of 1. A value greater than 1 specifies the switch location in a stack.
Sort By	Specifies a rule to sort data by. Values include: • **Source Address** : source IP address • **Destination Address** : destination IP address • **Protocol** : protocol number • **TOS** : type of service • **Port** : port number • **TCP/UDP Src Port** : TCP/UDP source port • **TCP/UDP Dest Port** : TCP/UDP destination port • **Packet Count** : packet number • **Byte Count** : data byte number • **First Packet Time** : first packet time • **Last Packet Time** : last packet time Default: Source Address
Sort Order	Specifies the order in which to sort data. Values include: • **Ascending** • **Descending** Default: Ascending
Display	Specifies the number of entries to display. Values include: • **Top 10** : displays first 10 entries • **Top 25** : displays first 25 entries • **Top 50** : displays first 50 entries • **Top 100** : displays first 100 entries • **Top 200** : displays first 200 entries Default: Top 50

Graphing IPFIX exporter statistics for a collector using EDM

Use the following procedure to graph collected and analyzed data exported from an IPFIX-compliant switch.

Prerequisites

- Open one of the supported browsers.
- Enter the IP address of the switch to open an EDM session.

Procedure steps

1. From the navigation tree, double-click **Serviceability**.
2. In the Serviceability tree, double-click **IPFIX**.
3. In the work area, click the **Collectors** tab.
4. Click **Graph**.
5. Click the **Exporter** tab.
6. To select collector data to graph, click any column in either the **OutPkts**, **OutOctets**, or **PktsLoss** row.
7. From the **Poll Interval** list, select an interval.
8. Click a **Line Chart**, **Area Chart**, **Bar Chart**, or **Pie Chart**.

Variable definitions

Variable	Value
OutPkts	Indicates the total number of packets sent.
OutOctets	Indicates the total number of bytes sent.
PktsLoss	Indicates the total number of records lost.

Viewing the IPFIX collector clear time using EDM

Use the following procedure to display the system time after IPFIX exporter statistics were last cleared.

Prerequisites

- Open one of the supported browsers.
- Enter the IP address of the switch to open an EDM session.

Procedure steps

1. From the navigation tree, double-click **Serviceability**.

2. In the Serviceability tree, double-click **IPFIX**.
3. In the work area, click the **Collectors** tab.
4. Click **Graph**.
5. Click the **Clear Time** tab.

Chapter 6: Remote Monitoring

This chapter provides conceptual information and procedures to configure Remote Monitoring.

Remote Network Monitoring (RMON)

The Remote Network Monitoring (RMON) Management Information Base (MIB) is an interface between the RMON agent on the switch and an RMON management application, such as the Enterprise Device Manager (EDM).

RMON defines objects that are suitable for the management of any type of network, but some groups are targeted for Ethernet networks in particular.

The RMON agent continuously collects statistics and proactively monitors switch performance.

RMON has the three following major functions:

- to create and display alarms for user-defined events
- to gather cumulative statistics for Ethernet interfaces
- To track the history of statistics for Ethernet interfaces

RMON scaling

The number of RMON instances per stack is 800.

Working of RMON alarms

The alarm variable is polled, and the result is compared against upper and lower limit values you select when you create the alarm. If either limit is reached or crossed during the polling period, the alarm triggers and generates an event that you can view in the event log or the trap log.

The upper limit of the alarm is called the *rising value*, and its lower limit is called the *falling value*. RMON periodically samples the data based upon the alarm interval. During the *first* interval that the data passes above the rising value, the alarm triggers as a rising event. During the first interval that the data drops below the falling value, the alarm triggers as a falling event.

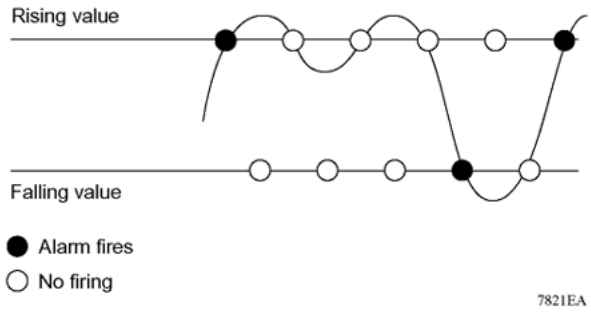

Figure 3: How alarms fire

It is important to note that the alarm triggers during the first interval that the sample goes out of range. No additional events are generated for that threshold until the opposite threshold is crossed. Therefore, it is important to carefully define the rising and falling threshold values for alarms to work as expected. Otherwise, incorrect thresholds cause an alarm to fire at every alarm interval.

A general guideline is to define one of the threshold values to an expected baseline value, and then define the opposite threshold as the out-of-bounds limit. Because of sample averaging, the value may be equal to ±1 of the baseline units. For example, assume an alarm is defined on octets going out of a port as the variable. The intent of the alarm is to provide notification to the system administrator when excessive traffic occurs on that port. If spanning tree is enabled, 52 octets are transmitted out of the port every 2 seconds, which is equivalent to baseline traffic of 260 octets every 10 seconds. This alarm provides the notification you need if the lower limit of octets going out is defined at 260 and the upper limit is defined at 320 (or at any value greater than 260 + 52 = 312).

The first time outbound traffic other than spanning tree Bridge Protocol Data Units (BPDU) occurs, the rising alarm triggers. When outbound traffic other than spanning tree ceases, the falling alarm triggers. This process provides the system administrator with time intervals of any nonbaseline outbound traffic.

You define the alarm with a falling threshold less than 260 (assuming the alarm polling interval is 10 seconds), say 250, the rising alarm can fire only once (see the following figure). For the rising alarm to fire a second time, the falling alarm (the opposite threshold) must fire. Unless the port becomes inactive or spanning tree is disabled (which causes the value for outbound octets to drop to zero), the falling alarm cannot fire because the baseline traffic is always greater than the value of the falling threshold. By definition, the failure of the falling alarm to fire prevents the rising alarm from firing a second time.

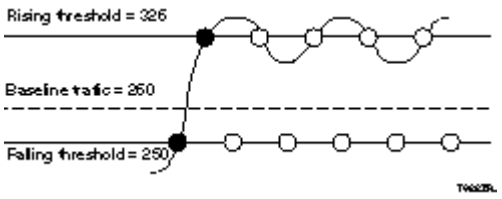

Figure 4: Alarm example - threshold less than 260

Creating alarms

When you create an alarm, select a variable from the variable list and the port, or other switch component, to which it is connected. Some variables require port IDs, card IDs, or other indices (for example, spanning tree group IDs). Then, select a rising and a falling threshold value. The rising and falling values are compared against the actual value of the variable that you choose. If the variable falls outside of the rising or falling value range, an alarm is triggered and an event is logged or trapped.

When an alarm is created, a sample type is also selected, which can be either absolute or delta. Absolute alarms are defined on the cumulative value of the alarm variable. An example of an alarm defined with absolute 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 set it an as absolute value. Therefore, an alarm can be created with a rising value of 2 and a falling value of 1 to alert a user about whether the card is up or down.

Note:

When you configure an RMON alarm with an owner, the system does not retain the owner configuration after reboot and the system displays the owner as "Entry from NVRAM".

Most alarm variables related to Ethernet traffic are set to delta value. Delta alarms are defined based on the difference in the value of the alarm variable between the start of the polling period and the end of the polling period. Delta alarms are sampled twice for each polling period. For each sample, the last two values are added together and compared to the threshold values. This process increases precision and allows for the detection of threshold crossings that span the sampling boundary. If you track the current values of a given delta-valued alarm and add them together the result is twice the actual value. (This result is not an error in the software.)

RMON events and alarms

RMON events and alarms work together to produce notification when values in the network go out of a specified range. When values pass the specified ranges, the alarm is triggered. The event specifies how the activity is recorded.

An event specifies whether a trap, a log, or a trap and a log are generated to view alarm activity. When RMON is globally enabled, two default events are generated:

- Rising Event
- Falling Event

Default events specify that when an alarm goes out of range, the firing of the alarm is tracked in both a trap and a log. For example, when an alarm triggers at the rising threshold, the rising event specifies that this information be sent to both a trap and a log. You can enable the viewing of the history of RMON fault events by using the stack. RMON Event Log window

How events work

An event specifies whether a trap, a log, or a trap and a log are generated to view alarm activity. When RMON is globally enabled, the following two default events are generated:

- RisingEvent
- FallingEvent

The default events specify that when an alarm goes out of range, the firing of the alarm is tracked in both a trap and a log. For example, when an alarm triggers at the rising threshold, the rising event specifies that this information be sent to both a trap and a log. Likewise, when an alarm passes the falling threshold, the falling event specifies that this information is sent to a trap and a log.

RMON Configuration using the CLI

This section provides procedures to configure Remote Monitoring (RMON) using the CLI.

Viewing the RMON alarms

About this task

Use this procedure to display information about RMON alarms.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Display the RMON alarms:

   ```
   show rmon alarm
   ```

Viewing the RMON events

About this task

Use this procedure to display information regarding RMON events.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   ```

```
configure terminal
```

2. View the RMON events:

```
show rmon event
```

Viewing the RMON history

About this task

Use this procedure to display information regarding the configuration of RMON history.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. View RMON history:

```
show rmon history
```

Example

```
Switch>enable
Switch#configure terminal
Switch#show interfaces
              Status           Auto                          Flow
Port Trunk Admin   Oper  Link Negotiation  Speed   Duplex Control
---- ----- ------- ---- ---- ----------- -------- ------ -------
1          Enable  Up   Up   Enabled     100Mbps  Full   Disable
2          Enable  Down Down Enabled
3          Enable  Down Down Enabled
4          Enable  Down Down Enabled
5          Enable  Down Down Enabled
6          Enable  Down Down Enabled
7          Enable  Down Down Enabled
8          Enable  Down Down Enabled
9          Enable  Down Down Enabled
10         Enable  Down Down Enabled
11         Enable  Down Down Enabled
12         Enable  Down Down Enabled
13         Enable  Down Down Enabled
14         Enable  Down Down Enabled
15         Enable  Down Down Enabled
16         Enable  Down Down Enabled
17         Enable  Down Down Enabled
18         Enable  Down Down Enabled
19         Enable  Down Down Enabled
----More (q=Quit, space/return=Continue)----
```

Viewing the RMON statistics

About this task

Use this procedure to display information regarding the configuration of RMON statistics.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. View RMON statistics:

   ```
   show rmon stats
   ```

Example

```
Switch>enable
Switch#configure terminal
Switch(config)#show rmon stats
Index Port
----- ----
1     1
2     2
3     3
4     4
5     5
6     6
7     7
8     8
9     9
10    10
11    11
12    12
13    13
14    14
15    15
16    16
17    17
18    18
19    19
20    20
----More (q=Quit, space/return=Continue)----
```

Configuring RMON alarms

About this task

Use this procedure to set RMON alarms and thresholds.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure RMON alarms:

   ```
   rmon alarm <1-65535> <WORD> <1-2147483647> {absolute |rdelta}
   rising-threshold <-2147483648-2147483647> [<1-65535>]falling-
   threshold <-2147483648-2147483647> [<1-65535>] [owner <LINE>]
   ```

Variable Definitions

Use the data in the following table to use the **rmon alarm** command.

Variable	Value
<1-65535>	Unique index for the alarm entry.
<WORD>	The MIB object to be monitored. This is an object identifier, and for most available objects. You can use an English name.
<1-2147483647>	The sampling interval, in seconds.
absolute	Use absolute values (value of the MIB object is compared directly with thresholds).
delta	Use delta values (change in the value of the MIB object between samples is compared with thresholds).
rising-threshold <-2147483648-2147483647 > [<1-65535>]	The first integer value is the rising threshold value. The optional second integer specifies the event entry to be triggered when the rising threshold is crossed. If omitted, or if an invalid event entry is referenced, no event is triggered. Unique index for the alarm entry.
falling-threshold <-2147483648-2147483647 > [<1-65535>]	The first integer value is the falling threshold value. The optional second integer specifies the event entry to be triggered when the falling threshold is crossed. If omitted, or if an invalid event entry is referenced, no event is triggered. Unique index for the alarm entry.
[owner <LINE>]	Specify an owner string to identify the alarm entry.

Deleting RMON alarms

About this task

Use this procedure to delete RMON alarm table entries.

Tip:

When you omit the variables, the system clears all entries in the table.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Delete RMON alarms:

   ```
   no rmon alarm [<1-65535>]
   ```

Variable Definitions

Use the data in the following table to use the **no rmon alarm** command.

Variable	Value
1-65535	Unique index for the event entry.

Configuring RMON events settings

About this task

Use this procedure to configure RMON event log and trap settings.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure RMON events:

   ```
   rmon event <1-65535> [log] [trap] [description <LINE>] [owner
   <LINE>]
   ```

Variable Definitions

Use the data in the following table to use the **rmon event** command.

Variable	Value
<1-65535>	Unique index for the event entry.
[log]	Records events in the log table.
[trap]	Generates SNMP trap messages for events.
[description <LINE>]	Specifies a textual description for the event.
[owner <LINE>]	Specifies an owner string to identify the event entry.

Deleting RMON events settings

About this task

Use this procedure to delete RMON event table entries.

Tip:

When you omit the variable, the system clears all entries in the table.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Delete RMON events settings:

```
no rmon alarm [<1-65535>]
```

Variable Definitions

Use the data in the following table to use the **no rmon alarm** command.

Variable	Value
1-65535	Unique index for the event entry.

Configuring RMON history settings

About this task

Use this procedure to configure RMON history settings.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure RMON history settings:

```
rmon history <1-65535> <LINE> <1-65535> <1-3600> [owner <LINE>]
```

Variable Definitions

Use the data in the following table to use the **rmon history** command.

Variable	Value
<1-65535>	Unique index for the history entry.
<LINE>	Specifies the port number to be monitored.
<1-65535>	The number of history buckets (records) to keep.
<1-3600>	The sampling rate (how often a history sample is collected).
[owner <LINE>]	Specifies an owner string to identify the history entry.

Deleting RMON history settings

About this task

Use this procedure to delete RMON history table entries. When you omit the variable, all entries in the table are cleared.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Delete RMON history:

```
no rmon history [<1-65535>]
```

Variable Definitions

Use the data in the following table to use the **no rmon history** command.

Variable	Value
1-65535	Unique index for the event entry.

Configuring RMON statistics settings

About this task

Use this procedure to configure RMON statistics settings.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure RMON statistics settings:

```
rmon stats <1-65535> <LINE> [owner <LINE>]
```

Variable Definitions

Use the data in the following table to use the **rmon status** command.

Variable	Value
<1-65535>	Unique index for the stats entry.
[owner <*LINE*>]	Specifies an owner string to identify the stats entry.

Deleting RMON statistics settings

About this task

Use this procedure to turn off RMON statistics.

Tip:

When omit the variable, the system clears all entries in the table.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Delete RMON statistics settings:

   ```
   no rmon stats [<1-65535>]
   ```

Variable Definitions

Use the data in the following table to use the **no rmon stats** command.

Variable	Value
1-65535	Unique index for the event entry.

RMON Configuration using the EDM

This provides procedures to configure Remote Monitoring (RMON) using the Enterprise Device Manager (EDM).

Prerequisites

- Open one of the supported browsers.
- Enter the IP address of the switch to open an EDM session.

RMON history management using EDM

Use the information in this section to display, create, and delete RMON history characteristics.

Viewing RMON history using EDM

Ethernet history records periodic statistical samples from a network. A sample is called a history and is gathered in time intervals referred to as buckets.

Histories establish a time-dependent method for gathering RMON statistics on a port. The default values for history are the folowing:

- Buckets are gathered at 30-minute intervals.
- Number of buckets gathered is 50.

You can configure the time interval and the number of buckets. However, when the last bucket is reached, bucket 1 is dumped and recycled to hold a new bucket of statistics. Then, bucket 2 is dumped, and so forth.

Use the following procedure to view RMON history.

Procedure steps

1. From the navigation tree, double-click **Serviceability**.
2. In the Serviceability tree, double-click **RMON**.
3. In the RMON tree, double-click **Control**.
4. On the work area, click the **History** tab to view the history.

Variable definitions

Use the data in the following table to help you create the RMON history characteristics.

Variable	Value
Index	A unique value assigned to each interface. An index identifies an entry in a table.
Port	Any Ethernet interface on the device.
BucketsRequested	Indicates the requested number of discrete time intervals over which data is to be saved in the part of the media-specific table associated with this entry.
BucketsGranted	Indicates the number of discrete sampling intervals over which data is saved in the part of the media-specific table associated with this entry. The actual number of buckets associated with this entry can be less than the value of this object. In this case, at the end of each sampling interval, a new bucket is added to the media-specific table.
Interval	Indicates the interval in seconds over which the data is sampled for each bucket in the part of the media-specific table associated with this entry. You can set this interval to any number of seconds between 1 and 3600 (1 hour). Because the counters in a bucket can overflow at their maximum value with no indication, note the possibility of overflow in any of the associated counters. Consider the minimum time in which any counter could overflow on a particular media type and set the historyControlInterval object to a value less than this interval. This minimum time is typically most important for the octets counter in any media-specific table. For example, on an Ethernet network, the etherHistoryOctets counter could overflow in about 1 hour at the maximum utilization of the Ethernet.
Owner	Indicates the network management system that created this entry.

Creating RMON history characteristics using EDM

You can use RMON to collect statistics at intervals. For example, if you want to gather RMON statistics over the weekend, you must configure enough buckets to cover two days. To do this, set the history to gather one bucket each hour, covering the 48-hour period. After you set history characteristics, you cannot modify them; you must delete the history and create another one.

Perform this procedure to establish a history for a port and set the bucket interval.

Procedure steps

1. From the navigation tree, double-click **Rmon**.

2. In the RMON tree, double-click **Control**.
3. In the work area, click **Insert** to open the Insert History dialog.
4. Type the port number or click the ellipsis to select a port from the list.
5. In the **Buckets Requested** box, type the number of buckets, or click the ellipsis to select a value from the list. The default value is 50.
6. In the **Interval** box, type the length of the interval or click the ellipsis to select a value from the list. The default value is 1800.
7. In the **Owner** box, type the owner— the network management system that created this entry.
8. Click **Insert** to add the entry to the list and return to the History tab.

   RMON collects statistics using the index, port, bucket, and interval that you specified.

Variable definitions

Use the data in the following table to help you create the RMON history characteristics.

Variable	Value
Index	A unique value assigned to each interface. An index identifies an entry in a table.
Port	Any Ethernet interface on the device.
BucketsRequested	Specifies the requested number of discrete time intervals over which data is to be saved in the part of the media-specific table associated with this entry.
BucketsGranted	Indicates the number of discrete sampling intervals over which data is saved in the part of the media-specific table associated with this entry. The actual number of buckets associated with this entry can be less than the value of this object. In this case, at the end of each sampling interval, a new bucket is added to the media-specific table.
Interval	Specifies the interval in seconds over which the data is sampled for each bucket in the part of the media-specific table associated with this entry. You can set this interval to any number of seconds between 1 and 3600 (1 hour). Because the counters in a bucket can overflow at their maximum value with no indication, note the possibility of overflow in any of the associated counters. Consider the minimum time in which any counter could overflow on a particular media type and set the historyControlInterval object to a value less than this interval. This minimum time is typically most important for the octets counter in any media-specific table. For example, on an Ethernet network, the etherHistoryOctets counter could overflow in about 1 hour at the maximum utilization of the Ethernet.
Owner	Specifies the network management system that created this entry.

Disabling RMON history using EDM

Use the following procedure to disable RMON history on a port.

Procedure steps

1. From the navigation tree, double-click **Serviceability**.
2. In the Serviceability tree, double-click **RMON**.
3. In the RMON tree, double-click **Control**.
4. On the work area, click the **History** tab to view the history.
5. In the table, select the row that you want to delete.
6. On the toolbar, click **Delete**.

Viewing RMON history statistics using EDM

Use the following procedure to display RMON history statistics:

Prerequisites

- Open one of the supported browsers.
- Enter the IP address of the switch to open an EDM session.

Procedure steps

1. From the navigation tree, double-click **Serviceability**.
2. In the Serviceability tree, double-click **RMON**.
3. In the RMON tree, double-click **Control**.
4. On the work area, click the **History** tab to view the history.
5. In the table, select a port row.
6. On the toolbar, click **Display History Data**.

Variable definitions

Use the data in the following table to help you understand the RMON history statistics display.

Variable	Value
SampleIndex	The sample number. As history samples are taken, they are assigned greater sample numbers.
Utilization	Estimate the percentage of the capacity of a link that is used during the sampling interval.
Octets	The number of octets received on the link during the sampling period.
Pkts	The number of packets received on the link during the sampling period.
BroadcastPkts	The number of packets received on the link during the sampling interval that destined for the packet address.

Table continues…

Variable	Value
MulticastPkts	The number of packets received on the link during the sampling interval that are destined for the multicast address. This does not include the broadcast packets.
DropEvents	The number of received packets that are dropped because of system resource constraints.
CRCAlignErrors	The number of packets received during a sampling interval that are between 64 and 1518 octets long. This length includes Frame Check Sequence (FCS) octets but not framing bits. The packets had a bad FCS with either an integral number of octets (FCS Error) or a nonintegral number of octets (Alignment Error).
UndersizePkts	The number of packets received during the sampling interval are less than 64 octets long (including FCS octets, but not framing bits).
OversizePkts	The number of packets received during the sampling interval are longer than 1518 octets (including FCS octets, but not framing bits, and are otherwise well formed).
Fragments	The number of packets received during the sampling interval are less than 64 octets long (including FCS octets, but not framing bits. The packets had a bad FCS with either an integral number of octects (FCS Error) or a nonintegral number of octets (Alignment Error).
Collisions	The best estimate of the number of collisions on an Ethernet segment during a sampling interval.

RMON Ethernet statistics management using EDM

Use the information in the following sections to manage RMON Ethernet statistics.

Viewing RMON Ethernet statistics using EDM

Use the following procedure to gather Ethernet statistics.

Procedure steps

1. From the navigation tree, double-click **Serviceability**.
2. In the Serviceability tree, double-click **RMON**.
3. In the RMON tree, double-click **Control**.
4. On the work area, click the **Ether Stats** tab to view the history.

Variable definitions

Use the data in the following table help you understand the RMON Ethernet statistics display.

Variable	Value
Index	A unique value assigned to each interface. An index identifies an entry in a table.
Port	A port on the device.

Table continues…

Variable	Value
DropEvents	The number of received packets that are dropped because of system resource constraints.
Octets	The total 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. For greater precision, sample the etherStatsPkts and etherStatsOctets objects before and after a common interval.
Pkts	The total number of packets (including bad packets, broadcast packets, and multicast packets) received.
BroadcastPkts	The total number of good packets received that are directed to the broadcast address. This does not include multicast packets.
MulticastPkts	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.
CRCAlignErrors	The total number of packets received with a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, but with either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).
UndersizePkts	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.
OversizePkts (>1518)	The total number of packets received that are longer than 1518 octets (excluding framing bits but including FCS octets) and were otherwise well formed.
Fragments	The total number of packets received that are less than 64 octets in length (excluding framing bits but including FCS octets) and with either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). For etherStatsFragments to increment is normal because it counts both runts (which are normal occurrences due to collisions) and noise hits.
Jabbers	The total number of packets received that are longer than 1518 octets (excluding framing bits, but including FCS octets), with either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error). Jabber is defined as the condition where any packet exceeds 20 ms. The allowed range to detect jabber is between 20 ms and 150 ms.
Collisions	The best estimate of the total number of collisions on this Ethernet segment.
1..64	The total number of packets (including bad packets) received that are less than or equal to 64 octets in length (excluding framing bits but including FCS octets).
65 ..127	The total number of packets (including bad packets) received that are greater than 64 octets in length (excluding framing bits but including FCS octets).

Table continues…

Variable	Value
128 ..255	The total number of packets (including bad packets) received that are greater than 127 octets in length (excluding framing bits but including FCS octets).
256..511	The total number of packets (including bad packets) received that are greater than 255 octets in length (excluding framing bits but including FCS octets).
512..1023	The total number of packets (including bad packets) received that are greater than 511 octets in length (excluding framing bits but including FCS octets).
1024..1518	The total number of packets (including bad packets) received that are greater than 1023 octets in length (excluding framing bits but including FCS octets).
Owner	The network management system that created this entry.

Enabling RMON Ethernet statistics gathering using EDM

Use the following procedure to gather Ethernet statistics.

Procedure steps

1. From the navigation tree, double-click **Serviceability**.
2. In the Serviceability tree, double-click **RMON**.
3. In the RMON tree, double-click **Control**.
4. On the work area, click the **Ether Stats** tab to view the history.
5. On the toolbar, click **Insert**.
6. Type an index in the **Index** field.
7. Click the Port ellipses (**...**), and select the port you want to use.
8. Type the owner name in the **Owner** field.
9. Click **Insert**.

Variable definitions

Use the data in the following table to enable RMON Ethernet statistics gathering.

Variable	Value
Index	A unique value assigned to each interface. An index identifies an entry in a table.
Port	A port on the device.
Owner	The network management system that created this entry.

Disabling RMON Ethernet statistics gathering using EDM

Use this procedure to disable Ethernet statistics.

Procedure steps

1. From the navigation tree, double-click **Serviceability**.
2. In the Serviceability tree, double-click **RMON**.
3. In the RMON tree, double-click **Control**.
4. On the work area, click the **Ether Stats** tab to view the history.
5. On the toolbar, select the port row you want to delete.
6. On the toolbar, click **Delete**.

RMON alarm management using EDM

This section describes the procedures you can use to use the alarm manager.

Viewing RMON alarm configuration information using EDM

Use the following procedure to create an alarm for receiving statistics and history using default values.

Procedure steps

1. From the navigation tree, double-click **Serviceability**.
2. In the Serviceability tree, double-click **RMON**.
3. In the RMON tree, double-click **Alarms**.
4. On the work area, click the **Alarms** tab.

Variable definitions

Use the data in the following table to help you understand the RMON alarm display.

Variable	Value
Index	Uniquely identifies an entry in the alarm table. Each such entry defines a diagnostic sample at a particular interval for an object on the device. Range is 1–65535.
Interval	Time period (in seconds) over which the data is sampled and compared with the rising and falling thresholds.
Variable	Name and type of alarm—indicated by the format: • *alarmname.x* where x=0 indicates a chassis alarm. • *alarmname*. where you must specify the index. This is a card number for module-related alarms, an STG ID for spanning tree group alarms (the default STG is 1, other STG IDs are user-configured), or the Ether Statistics Control Index for RMON Stats alarms. • *alarmname* with no dot or index is a port-related alarm and displays in the port selection tool.

Table continues…

Variable	Value
Sample Type	Specifies the sample type—absolute or delta.
Value	Indicates the value of the alarm statistic during the last sampling period, compared with the rising and falling thresholds.
StartupAlarm	Indicates the type of alarm generated at startup, based on rising and falling thresholds. Values include: • risingAlarm • risingOrFallingAlarm • fallingAlarm
RisingThreshold	When the current sampled value is greater than or equal to this threshold, and the value at the last sampling interval is less than this threshold, generates a single event.
RisingEventIndex	Index of the event entry that is used when a rising 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. (Generally, accept the default that is already filled in.)
FallingThreshold	When the current sampled value is less than or equal to this threshold, and the value at the last sampling interval is greater than this threshold, generates a single event.
FallingEventindex	Index of the event entry that is used when 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. (Generally, accept the default that is already filled in.)
Owner	Specifies the owner name.
Status	Indicates the status of the alarm entry.

Creating an RMON alarm using EDM

Use the following procedure to create an alarm for receiving statistics and history using default values.

Procedure steps

1. From the navigation tree, double-click **Serviceability**.
2. In the Serviceability tree, double-click **RMON**.
3. In the RMON tree, double-click **Alarms**.
4. On the work area, click the **Alarms** tab to view the history.
5. On the toolbar, click **Insert**.
6. Configure the parameters as required.
7. Click **Insert**.

Variable definitions

The following table describes the RMON Insert Alarm dialog box fields.

Variable	Value
Variable	Name and type of alarm—indicated by the format: • *alarmname.x* where x=0 indicates a chassis alarm. • *alarmname* . where you must specify the index. This is a card number for module-related alarms, an STG ID for spanning tree group alarms (the default STG is 1, other STG IDs are user-configured), or the Ether Statistics Control Index for RMON Stats alarms. • *alarmname* with no dot or index is a port-related alarm and displays in the port selection tool.
Sample Type	Specifies the sample type—absolute or delta.
Interval	Specifies the time period (in seconds) over which the data is sampled and compared with the rising and falling thresholds.
Index	Uniquely identifies an entry in the alarm table. Each such entry defines a diagnostic sample at a particular interval for an object on the device. Range is 1–65535.
RisingThreshold	When the current sampled value is greater than or equal to this threshold, and the value at the last sampling interval is less than this threshold, generates a single event.
RisingEventIndex	Specifies the index of the event entry that is used when a rising 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. (Generally, accept the default that is already filled in.)
FallingThreshold	When the current sampled value is less than or equal to this threshold, and the value at the last sampling interval is greater than this threshold, generates a single event.
FallingEventindexSpecifies the	Specifies the index of the event entry that is used when 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. (Generally, accept the default that is already filled in.)
Owner	Specifies the owner name.

Deleting an RMON alarm using EDM

Use this procedure to delete an alarm:

Procedure steps

1. From the navigation tree, double-click **Serviceability**.
2. In the Serviceability tree, double-click **RMON**.
3. In the RMON tree, double-click **Alarms**.
4. On the work area, click the **Alarms** tab.
5. In the table, select the alarm you want to delete.
6. On the toolbar, click **Delete**.

7. Click **Yes**.

Event management using EDM

This section describes the procedures you can use to configure RMON events and alarms work together to provide notification when values in the network are outside of a specified range. When values pass the specified ranges, the alarm is triggered. The event specifies how the activity is recorded.

Prerequisites

- Open one of the supported browsers.
- Enter the IP address of the switch to open an EDM session.

Viewing an event using EDM

Use the following procedure to view a table of events.

Procedure steps

1. From the navigation tree, double-click **Serviceability**.
2. In the Serviceability tree, double-click **RMON**.
3. In the RMON tree, double-click **Alarms**.
4. On the work area, click the **Events** tab to view the history.

Variable definitions

The following table describes the Events tab fields.

Variable	Value
Index	This index uniquely identifies an entry in the event table. Each entry defines one event that is to be generated when the appropriate conditions occur.
Description	Specifies whether the event is a rising or falling event.
Type	The type of notification that the switch provides about this event. In the case of log, an entry is made in the log table for each event. In the case of trap, an SNMP trap is sent to one or more management stations. Possible notifications follow: • none • log • trap • log-and-trap
Community	The SNMP community string acts as a password. Only those management applications with this community string can view the alarms.

Table continues…

Variable	Value
LastTimeSent	The value of sysUpTime at the time this event entry last generated an event. If this entry has not generates any events, this value is zero.
Owner	If traps are specified to be sent to the owner, this is the name of the machine that receives alarm traps.

Creating an event using EDM

Use the following procedure to create an event.

Procedure steps

1. From the navigation tree, double-click **Serviceability**.
2. In the Serviceability tree, double-click **RMON**.
3. In the RMON tree, double-click **Alarms**.
4. On the work area, click the **Events** tab to view the history.
5. On the toolbar, click **Insert**.

   The Insert Events dialog box appears.
6. Type an index in the **Index** field.
7. Type the name of the event in the **Description** field.
8. Choose the type of the event in the **Type** field.
9. Type the community information in the **Community** field.
10. Type the owner information in the **Owner** field.
11. Click **Insert**.

Deleting an event using EDM

Use this procedure to delete an event.

Procedure steps

1. From the navigation tree, double-click **Serviceability**.
2. In the Serviceability tree, double-click **RMON**.
3. In the RMON tree, double-click **Alarms**.
4. On the work area, click the **Events** tab to view the history.
5. In the table, select the event row you want to delete.
6. On the toolbar, click **Delete**.

Managing log information management using EDM

Use the information in this procedure to chronicle and describe alarm activity.

Prerequisites

- Open one of the supported browsers.
- Enter the IP address of the switch to open an EDM session.

Procedure steps

1. From the navigation tree, double-click **Serviceability**.
2. In the Serviceability tree, double-click **RMON**.
3. In the RMON tree, double-click **Alarms**.
4. On the work area, click the **Log** tab to view the history.

Variable definitions

The following table describes the Log tab fields.

Variable	Value
Time	Specifies when an event occurs that activats the log entry.
Description	Specifies whether the event is a rising or falling event.
EventIndex	Specifies the event index.

Chapter 7: sFlow

This chapter provides conceptual information and procedures to configure sFlow using Command Line Interface (CLI) and Enterprise Device Manager (EDM).

sFlow

sFlow monitors traffic in a data network. sFlow monitors routers and switches in a network and captures traffic statistics about these devices. sFlow also detects and traces unauthorized network activity.

Sampling

sFlow consists of multiple devices performing the following types of samplings:

- Packet sampling
- Periodic counter sampling

Packet 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 agent takes a sample. The default sampling rate is 0 packets. Packet sampling is performed by Application-Specific Integrated Circuit (ASIC), providing wire-speed performance.

Periodic counter sampling periodically polls and exports counters for a configured interface. The default polling interval is 0 seconds. A polling interval defines how often the network device checks various counters on sFlow enabled interfaces. Those counters are copied in a sFlow datagram. The sFlow datagram is sent to the assigned collector when its size is near to the configured max-header-size. sFlow counter sampling is more efficient than SNMP polling when monitoring a large number of interfaces.

Components

Following are the two primary sFlow components:

- sFlow agent. Exists on the network device to be monitored.
- sFlow collector. Exists on a central server.

sFlow agent combines interface counters and flow samples into sFlow datagrams that are sent across the network to an sFlow collector. sFlow collector analyzes the sFlow datagrams to produce a rich, real-time, network-wide view of traffic flows.

The following figure shows the basic elements of the sFlow system.

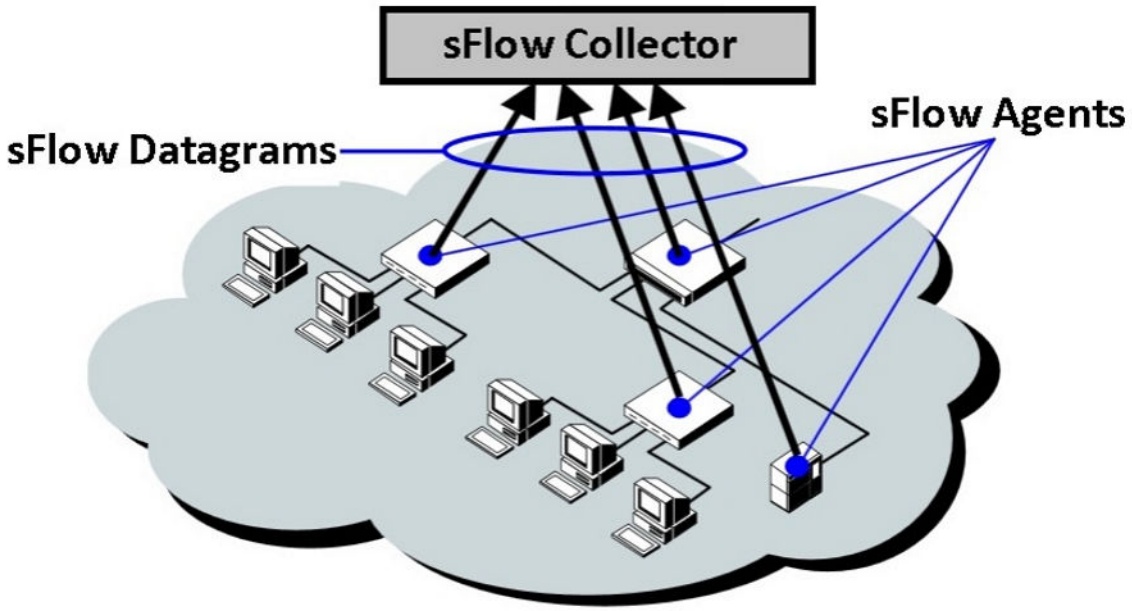

Figure 5: Basic elements of sFlow system

Note:

All sFlow samples are processed by the CPU. If you observe an unusual load on the CPU, increase the sFlow sampling rate.

sFlow configuration using CLI

This section describes the CLI commands used to configure sFlow on the switch.

Enabling sFlow globally

You must globally enable sFlow before the system can monitor and capture traffic statistics to send to an sFlow collector. By default, sFlow is globally disabled.

About this task

After you globally enable sFlow, you must configure at least one sFlow collector, and enable sFlow sampling rates on an interface port.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enable sFlow:

   ```
   sflow enable
   ```

3. Verify global sFlow state:

   ```
   show sflow
   ```

Example

```
Switch(config)#show sflow
SFLOW Administrative State: Enabled
```

Displaying sFlow interface settings

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Display sFlow per interface settings:

   ```
   show sflow interface <port>
   ```

Example

```
Switch#show sflow interface 1
Port   Packet Sample Rate  Max Hdr  Counter  Collector ID
       Ingress   Egress    Size     Polling   1  2  3  4
-----  --------- --------  -------  -------  ------------
1      0         0          128     0         0  0  0  0
```

Disabling sFlow globally

Use this procedure to disable sFlow globally.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. At the command prompt, enter the following command:

   ```
   no sflow enable
   ```

   or

   ```
   default sflow enable
   ```

Configuring an sFlow collector

Configure an sFlow collector to determine to which device the sFlow agent sends sFlow datagrams.

Before you begin

- You must globally enable sFlow.
- You must enable sFlow on a port.

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

1. Enter Global Configuration mode:

   ```
   enable

   configure terminal
   ```

2. Use the following command to configure sFlow collector:

   ```
   sflow collector <1-4> address <a.b.c.d> | <ipv6addr> [owner <word>]
   [port <1-65535>]
   ```

3. Use the following command to configure sFlow collector buffer size and timeout:

   ```
   sflow collector <1-4> [max-datagram-size <400-9216>][timeout
   <0-65535>]
   ```

4. Verify the collector configuration:

   ```
   show sflow collector  [<1-4>]
   ```

Example

```
switch(config)#show sflow collector

collector 4
------------
      IP Version               IPv4
      IP Address               192.0.2.0
      IP Destination Port      4
      Collector Buffer Size    400
      Collector Owner          3
      Collector Timeout        41
```

Variable definitions

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

Variable	Definition
sflow collector <1-4>	Specifies the Collector ID.
address <a.b.c.d> | <ipv6addr>	Specifies the IPv4 address or IPv6 address.
owner <word>	Specifies the owner that created the entry (maximum length 20).
port <1-65535>	Specifies the UDP port number. The default UDP port is 6343.
max-datagram-size <400-9216>	Specifies the maximum size of the datagram packets. The default is 1400.
timeout <0-65535>	Specifies the timeout until collector is deleted. The default is 0.

Deleting an sFlow collector

Use this procedure to delete an sFlow collector.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. At the command prompt, enter the following command:

   ```
   no sflow collector <1-4>
   ```

   or

   ```
   default sflow collector <1-4>
   ```

Enabling sFlow on a Port

Enable sFlow on a port to capture traffic statistics for the port.

Note:

You can enable traffic sampling and counter polling on a port or multiple ports by using the **`sflow`** commands in Interface Configuration mode.

Before you begin

You must globally enable sFlow.

Procedure

1. Enter Ethernet Interface Configuration mode:

   ```
   enable
   ```

```
configure terminal

interface Ethernet <port>
```

2. Set collector on port:

```
sflow [port <portlist>] collector <1-4>
```

3. Set counter-interval on port:

```
sflow [port <portlist>] counter-interval <1-3600>
```

4. Set maximum captured header size on port:

```
sflow [port <portlist>] max-header-size <64-256>
```

5. Set traffic sampling on port:

```
sflow [port <portlist>] sampling-rate <ingress | egress>
<4096-1000000>
```

6. Verify the port configuration:

```
Show sflow interface [enabled] [<port number>]
```

Example

```
Switch(config-if)#show sflow interface 1

Port   Packet Sample Rate  Max Hdr  Counter  Collector ID
       Ingress   Egress    Size     Polling   1  2  3  4
-----  --------- --------  -------  -------  -----------
1      0         0          64      3         1  0  0  4
```

Variable definitions

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

Variable	Definition
portlist	Specifies the port on the device.
collector <1-4>	Specifies the Collector ID.
counter-interval <1-3600>	Specifies the counter on the polling interval on the port.
max-header-size <64-256>	Specifies the maximum captured header size. The default is 128.
sampling-rate [egress | ingress] <4096-1000000>	Specifies the sampling rate on the port. The default is 0.

Deleting or defaulting sFlow settings on a port

Use this procedure to delete or default sFlow settings on a port.

Procedure

1. Enter Ethernet Interface Configuration mode:

```
enable
configure terminal
interface Ethernet <port>
```

2. At the command prompt, enter the following command:
   a. To delete an sFlow collector on a port:

      ```
      no sflow collector <1-4>
      ```

      or

      ```
      default sflow collector <1-4>
      ```

   b. To disable counter polling on a port:

      ```
      no sflow counter-interval
      ```

      or

      ```
      default sflow counter-interval
      ```

   c. To disable sampling on a port:

      ```
      no sflow sampling-rate [ingress | egress]
      ```

      or

      ```
      default sflow sampling-rate [ingress | egress]
      ```

   d. To default maximum captured header size on a port:

      ```
      no sflow max-header-size
      ```

      or

      ```
      default sflow max-header-size
      ```

sFlow configuration using EDM

This section describes sFlow configuration.

Enabling sFlow globally

Use the following procedure to enable sFlow globally for a switch or stack.

Procedure

1. From the navigation pane, double-click **Serviceability**.
2. In the **Serviceability** tree, click **sFlow**.
3. In the sFlow work area, click the **Global** tab.

4. In the sFlow work area, click **Enable** check box to enable sFlow globally.
5. On the tool bar, click **Apply**.
6. On the tool bar, click **Refresh** to verify sFlow global configuration.

Configuring sFlow collectors

Use this procedure to configure sFlow collectors for a switch or stack.

Note:

UDP port cannot be changed from EDM for an existing collector. If UDP port for a collector is needed, it must be configured while creating the collector.

Procedure

1. From the navigation pane, double-click **Serviceability**.
2. In the **Serviceability** tree, click **sFlow**.
3. In the sFlow work area, click the **Collectors** tab.
4. Select and configure Collectors parameters as required.
5. On the toolbar, click **Apply**.
6. On the toolbar, you can click **Refresh** to verify the sFlow configuration.

Variable definitions

Variable	Definition
Index	Specifies the collector index.
Owner	Specifies the owner that created the entry.
AddressType	Specifies the IP address type.
Address	Specifies the IP address.
Timeout	Specifies the timeout until collector is deleted.
MaximumDatagramSize	Specifies the maximum datagram size.
UDP Port	Specifies the UDP port number.
DatagramVersion	Specifies the datagram version.

Configuring sFlow interfaces

Use the following procedure to configure sFlow interfaces for a switch or stack.

Procedure

1. Follow one of the following paths:
   - From the **Device Physical View**, use Ctrl-click to select more than one port, right-click **Edit**, then click the **sFlow** tab.
   - From the **Device Physical View**, use Ctrl-click to select more than one port, then follow the navigation tree to **Edit > Chassis > Ports > sFlow** tab.
   - From the navigation tree, select **Serviceability > sFlow > Interfaces** tab.
2. Configure the parameters as required in the port row.
3. Optionally, to configure parameters for multiple ports, you can use the Multiple Port Configuration section as below.
4. In the work area, in the Make Selection section of the Multiple Port Configuration pane, click the Switch/Stack/Ports ellipsis (...) to open the Port Editor dialog. If there is no Switch/Stack/Ports selection and you have already selected ports from the **Device Physical View**, proceed to the next step.
   a. In the Port Editor window, click the ports you want to configure. If you want to configure all ports, click **All**.
   b. Click **OK** to return to the Make Selection pane.

   The ports you selected appear in the Switch/Stack/Ports box.
5. To change the configuration of the selected ports, in the Multiple Port Configuration pane, double-click the cell beneath the column heading that represents the parameter you want to change and do one of the following:
   - If applicable, select a value from a drop-down list.
   - Otherwise, type a value in the cell.
6. In the Make Selection pane, click **Apply Selection**.

   The changes appear in the table.
7. **(Optional)** Click **Clear Selection** to clear Multiple Port Configurations or click **Hide Non-Editable** to display only those parameters that are editable in the Multiple Port Configuration pane for the selected ports.

Variable definitions

Variable	Definition
DataSource	Specifies the data source.
Collectors	Specifies the destination collector defined for the selected port(s).
IngressPacketSamplingRate	Specifies the ingress packet sampling rate.
EgressPacketSamplingRate	Specifies the egress sampling rate.

Table continues…

Variable	Definition
MaximumHeaderSize	Specifies the maximum captured header size.
CounterInterval	Specifies the counter interval on the port.

Chapter 8: Application Telemetry

This chapter provides conceptual information and procedures to configure Application Telemetry using Command Line Interface (CLI) and Enterprise Device Manager (EDM).

Application Telemetry Fundamentals

Application Telemetry is an analytics solution that combines the Deep Packet Inspection capabilities of Extreme Analytics Engine with sFlow data. This feature provides granular visibility into your network and monitors application performance, users, locations, and devices without the need for expensive sensors or collectors.

Application Telemetry uses policy rules to filter packets for analysis. This methodology enables Application Telemetry to monitor *all* application-level traffic flows at wire speed on *all* specified interfaces simultaneously.

Extreme Networks offers two Analytics solutions that monitor traffic on your network:

- sFlow
- Application Telemetry

Important:

You can use either only sFlow or only Application Telemetry, or both at the same time, as they can coexist on a switch. For a full XMC experience, the use of both sFlow and Application Telemetry is recommended.

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.

For further information about sFlow, see sFlow on page 164.

For more information about Extreme Management Center, see the documentation on the Extreme Networks Documentation Portal (https://www.extremenetworks.com/support/documentation/extremeanalytics-8-1/) with special attention to the *Extreme 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 and sends packet samples. 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* specified interfaces simultaneously.

The policy rules that Application Telemetry uses are based on filters configured in a policy configuration file called `apptel_default.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.

The following figure shows the Application Telemetry agent on various routers and switches with packets being sent to the Analytics Engine.

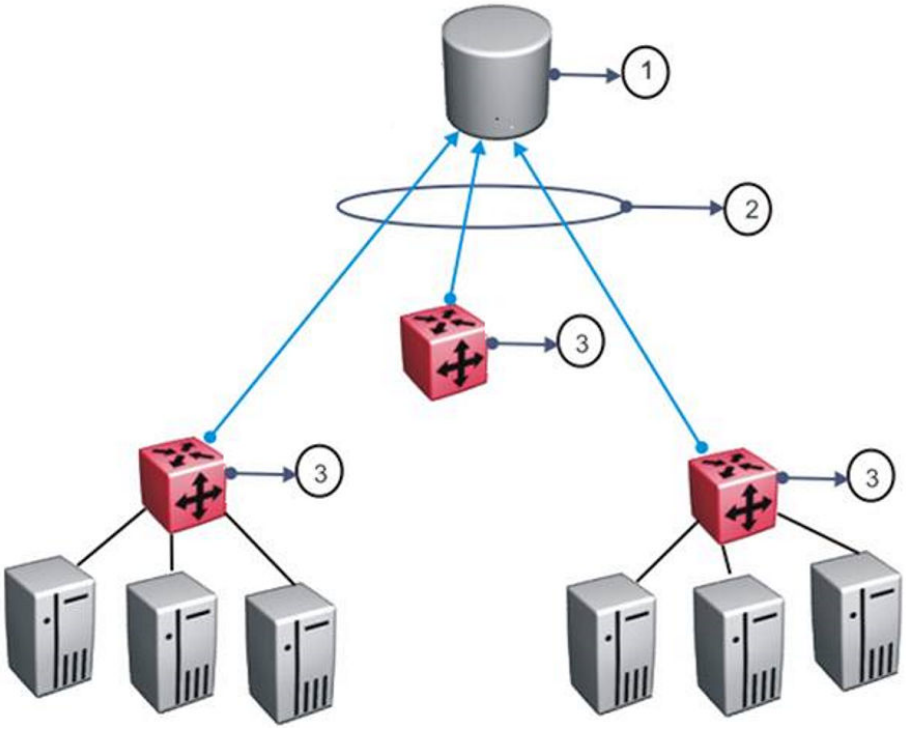

Figure 6: Application Telemetry Overview

Table 11: Application Telemetry Legend

Number	Description
1	Analytics Engine
2	GRE tunnels
3	Application Telemetry agents

Important:

When you enable Application Telemetry, the switch loads the filter rules based on the following logic:

- Application Telemetry uses the `apptel_custom.pol` or the `apptel_default.pol` file because the filter rules can exist in either file.

  The `apptel_custom.pol` file is the user-defined file. To use this file, you must upload it to `NVRAM` using either the `script run` command or `configure network` command.

  The `apptel_default.pol` file is the default file and is included with the image. This file contains the default filter rules.

- When you enable Application Telemetry, the feature uses the files in the following order:
  - If the user-defined file (`apptel_custom.pol`) exists, then the switch loads the rules from this file.

- If the `apptel_custom.pol` file does not exist or if there is a problem reading this file, then the switch uses the default `apptel_default.pol` file.

Common Elements Between sFlow and Application Telemetry

sFlow and Application Telemetry send mirrored packets from a common *source* to a common *destination*.

The tunnel source is the switch that you want to monitor:

- Application Telemetry sends packets that match its policy rules.

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.

Configuration Considerations and Restrictions

Consider the following when you configure Application Telemetry:

- The GRE tunnel source IPs cannot be configured using the CLI. The tunnel source IP is the switch/stack's inband address and only IPv4 addresses are supported. The destination IP address can be configured using the CLI.
- Application Telemetry uses one of the four Port Mirroring instances and the GRE tunnel will serve as a destination for that one mirroring instance.
- If a filter rule conflicts with another one, the rule with the better priority takes precedence.
- Application Telemetry filtering rules are not user configurable.
- Application Telemetry is not supported over Layer 2 VSNs.
- Application Telemetry supports Analytics Engine reachability over IP Shortcut Routing, and the global routing table (GRT). It does not support Analytics Engine reachability over the management VRF, VRFs or Layer 2 VSNs.
- Application Telemetry sends data only over IPv4 GRE tunnels.
- Application Telemetry filters the traffic on all front panel ports. Hence, oversubscription on the monitor port and on the stack ports is possible.
- Application Telemetry filters require up to 2 QoS precedences for installation, depending on the number of the entries in the policy file (1 precedence is required for every 128 entries). Before attempting to enable Application Telemetry, you should make sure that enough precedences are available.

- **Port Mirroring Resources:**
  - Port mirroring resources are limited to four instances.
  - Port mirroring shares these four resources with other applications such as RSPAN, Fabric Extend, Application Telemetry. Each one of these applications consumes at least one port mirroring resource. (RSPAN consumes two if you configure both Ingress and Egress modes.)

    Important:

    To enable any one of the above applications, you must have at least one free mirroring resource. If all four port mirroring resources are already in use, the switch displays a `Resource not available` error message when you try to enable the application.

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.

Defaulting Application Telemetry

Perform this procedure to default all the configured options of the Application Telemetry feature.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enter the following command to default the Application Telemetry settings:

   ```
   default app-telemetry
   ```

   Note:

   Defaulting the Application Telemetry settings will delete the uploaded user-defined policy.

Example

The following is an example for defaulting Application Telemetry:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch:1#(config)#default app-telemetry
```

Uploading the User-Defined Policy Configuration File

Perform this procedure to upload the user-defined policy configuration file.

Note:

If you do not upload a user-defined policy configuration file, the switch applies the default policy configuration file `apptel_default.pol`.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Upload the user-defined policy configuration file:

   ```
   configure network address <A.B.C.D> filename <WORD>
   ```

   or

   ```
   script run tftp {<hostname>|<A.B.C.D>} <WORD>
   ```

Example

The following is an example for uploading the user-defined policy configuration file:

```
Switch:1>enable
Switch:1#configure network address 10.10.10.2 filename apptel_custom.pol
or
Switch:1>enable
Switch:1#script run tftp 10.10.10.2 apptel_custom.pol
```

Variable Definitions

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

Variable	Definition
address*<A.B.C.D>*	Specifies the TFTP Server IP address (IPv4).
filename *word*	Specifies the filename of the configuration file.

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

Variable	Definition
tftp *<hostname>*|*<A.B.C.D>*	Specifies the hostname or IP address of TFTP server
word	Specifies the filename of the configuration file.

Enabling Application Telemetry

Perform this procedure to enable Application Telemetry globally or on a specific port.

Note:

By default, Application Telemetry is globally disabled.

Before you begin

You must upload the policy configuration file.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enable Application Telemetry:

   ```
   app-telemetry enable [ports <LINE>]
   ```

   Note:

   If the port details are not specified, then the Application Telemetry will be enabled on all ports.

Example

The following is an example to enable Application Telemetry globally:

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

The following is an example to enable Application Telemetry on a specific port:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch:1(config)#app-telemetry enable ports 1/5
```

Variable Definitions

Use the data in the following table to use the **`app-telemetry enable [ports]`** command.

Variable	Definition
enable	Enables Application Telemetry.
ports *<LINE>*	Specifies the port on which Application Telemetry can be enabled.

Configuring a Collector Address

Perform this procedure to configure a collector address.

Note:

Before you change or remove the Collector address, you must disable Application Telemetry.

Before you begin

You must upload the policy configuration file.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure the Application Telemetry Collector address:

   ```
   app-telemetry collector address <A.B.C.D>
   ```

Example

The following is an example to configure a collector address:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch:1(config)#app-telemetry collector address 10.10.10.2
```

Variable Definitions

Use the data in the following table to use the **`app-telemetry collector`** command.

Variable	Definition
address <*A.B.C.D.*>	Specifies the IP address of the collector. **Note:** Application Telemetry sends data only over IPv4 GRE tunnels.

Displaying Application Telemetry Counters

Perform this procedure to display 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. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Display Application Telemetry counters:

```
show app-telemetry counters {id<1-256>|name<LINE>}
```

Example

The following is an example for the **show app-telemetry counters** command output:

Displaying Application Telemetry Status

Perform this procedure to display the Application Telemetry status.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Display Application Telemetry status:

   ```
   show app-telemetry status
   ```

Example

The following is an example for the **show app-telemetry status** command output:

```
Switch:1>enable
Switch:1#show app-telemetry status
APPTEL is disabled
The collector's address is: 0.0.0.0
```

The following is an example for the **show app-telemetry status** command output, when the Application Telemetry is enabled:

```
Switch:1>enable
Switch:1#show app-telemetry status
APPTEL is enabled
Current set of ports: ALL
The collector's address is: 10.10.10.2
The collector is reachable via:
        U 0/P 1/NH 0-9-f-9-0-6/172.16.120.1
Policy file in use: apptel_default.pol
```

Clearing Application Telemetry Counters

Perform this procedure to clear the Application Telemetry status counters. You can clear all the counters or specify just the counters you want to clear by name or ID.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Clear Application Telemetry counters:

   ```
   clear app-telemetry counters {id<1-256>|name<LINE>}
   ```

3. Verify that the counters were cleared:

   ```
   show app—telemetry counters {id<1-256>|name<LINE>}
   ```

Example

The following is an example for the **`clear app-telemetry counters`** command output:

```
Switch:1>enable
Switch:1#clear app-telemetry counters
% All counters cleared!
Switch:1#show app-telemetry counters
================================================
Filter |  Filter  |    No. of     |    No. of
Number |   Name   |    Packets    |    Bytes
================================================
 1      ssh            0               0
 2      sslclient      0               0
```

Disabling Application Telemetry

Perform this procedure to disable Application Telemetry globally.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   ```

   ```
   configure terminal
   ```

2. Disable Application Telemetry:

   ```
   no app-telemetry enable
   ```

Example

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

Deleting the Collector Address

Perform this procedure to delete a collector address.

Note:

Before you change or remove the Collector address, you must disable Application Telemetry.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   ```

```
configure terminal
```

2. Configure the Application Telemetry Collector address:

```
no app-telemetry collector address
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch:1(config)#no app-telemetry collector address
```

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 Application Telemetry Globally

Perform 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

If you want to use a custom policy file, you must make sure that your custom policy file is uploaded.

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. In the **AddressType** field, select **ipv4** button.

   Note:

   Application Telemetry sends data only over IPv4 GRE tunnels.
6. In the **Address** field, enter the IPv4 address.
7. In the **PortList** box, click the ellipsis (...) button, and select the port(s).
8. Click **Apply**.

Globals Field Descriptions

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

Name	Description
AdminEnable	Shows whether Application Telemetry is enabled. By default, the check box is not enabled.
ClearCounterStats	Clears the Application Telemetry status counters.
AddressType	Specifies the collector IP address type. Only IPv4 collector addresses are supported.
Address	Specifies the IP address.
PortList	Specifies the port or ports on which Application Telemetry's filtering and mirroring policy is applied.
PolicyFileInUse	Shows the policy configuration file which is currently in use.

Viewing Application Telemetry Counters

Perform this 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.

Counter Field Descriptions

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

Name	Description
CounterId	Specifies the counter ID.
CounterName	Specifies the counter name.
CounterPkts	Specifies the number of counter packets received.
CounterBytes	Specifies the number of counter bytes used.

Clearing Application Telemetry Counters

Perform this procedure to clear the Application Telemetry status counters. You can clear all 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 **Clear Counter**.
4. Click **Apply**.

Viewing Application Telemetry Status

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.

Variable	Value
IsReachable	Shows whether the Application Telemetry collector is reachable.
NextHop	Shows the name or address of the next hop through which the collector is reachable.

Chapter 9: Service Level Agreement Monitor

This chapter provides conceptual information and procedures to configure Service Level Agreement Monitor (SLA Mon) using Command Line Interface (CLI) and Enterprise Device Manager (EDM).

SLA Mon Fundamentals

The switch supports the SLA Mon agent as part of the SLA Mon solution. You must have an Avaya Diagnostic Server with SLA Mon technology in your network to use the SLA Mon feature. Most of the SLA Mon configuration occurs on the server; configuration on the SLA Mon agent is minimal.

SLA Mon uses a server and agent relationship to perform end-to-end network Quality of Service (QoS) validation. You can use the test results to target under-performing areas of the network for deeper analysis.

SLA Mon Server and Agent

SLA Monitor agent performs QoS tests after it receives a request from the SLA Monitor server. The tests can be performed even if the server is not available.

The SLA Monitor server initiates the SLA Monitor functions on two or more agents. The agents run specific QoS tests at the request of the server. Agents exchange packets between one another to conduct the QoS tests. The test schedule and the exact nature and intensity of each test depends on the parameters that are configured on the server. The server stores the data it collects from the agents about the network. SLA Monitor can monitor a number of key items, including the following:

- network paths
- Differentiated Services Code Point (DSCP) markings
- loss
- jitter
- delay

The following figure illustrates an SLA Monitor implementation:

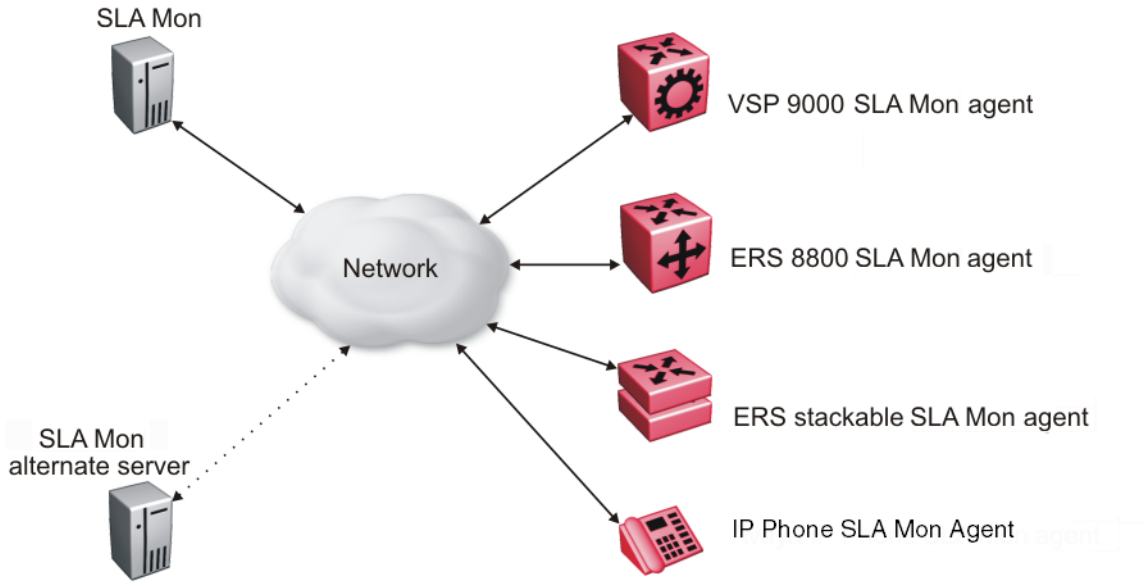

An SLA Monitor agent remains dormant until it receives a User Datagram Protocol (UDP) discovery packet from the server. The agent accepts the discovery packet to register with an SLA Monitor server. If the registration process fails, the agent remains dormant until it receives another discovery packet.

An agent can attempt to register with a server once every 60 seconds. After a successful registration, the agent will reregister with the server every 6 hours to exchange a new encryption key, if encryption is supported.

An agent only accepts commands from the server to which it is registered. An agent can use alternate servers to provide backup for time-out and communication issues with the primary server.

Secure agent-server communication

The secure SLA Monitor agent-server communication feature supports certificate-based authentication and encrypted agent-server communication. The communication mode is based on the ERS image. Secure images use authentication/encryption and non-secure images use clear text communication. Mocana security libraries are used for authentication and encryption. During registration, an X.509 certificate is retrieved from the server and then validated against the stored CA certificate. If the received certificate is trusted, a secure channel is established. A symmetric encryption key is exchanged and used for all subsequent agent server communication.

Note:

The certificate-based authentication and encrypted agent-server communication is automatically enabled on secure ERS images. This feature cannot be configured by the user.

QoS tests

SLA Monitor uses two types of tests to determine QoS benchmarks:

- Real Time Protocol (RTP)

  This test measures network performance, for example, jitter, delay, and loss, by injecting a short stream of UDP packets from source to destination (an SLA Monitor agent).

- New Trace Route (NTR)

  This test is similar to traceroute but also includes DSCP values at each hop in the path from the source to the destination. The destination does not need to be an SLA Monitor agent.

You can use NTR and RTP to perform the following tests in the absence of an SLA Monitor server:

- You can access the SLA Monitor CLI through the SLAMon Agent Address and SLAMon Agent Port. By default, access to the SLA Monitor CLI interface is disabled. If access is enabled, the SLA Monitor CLI interface becomes available when the SLA Monitor agent is enabled. Tests are run serially and only one type of test can be run at a time. Established sessions time-out after a specified interval. The time interval can be 60 seconds to 600 seconds. By default, the interval is 60 seconds. You can disable the SLA Monitor CLI interface if the functionality is not required.
- You can run the NTR and RTP tests through the CLI using the Application Configuration mode. The SLA Monitor agent must be enabled. Tests are run serially and only one type of test can be run at a time.

Note:

Server bypass must be enabled on the agents that are not registered with the server but are target agents for the RTP tests.

The error message “Unable to initiate test - agent busy” or “Reported Issue: test request denied by remote agent” appears if any tests are executed during the same time when the tests initiated by the server are executed. The server initiated tests typically takes priority. Do any one of the following if the error message appears:

- Stop the server.
- Enable SLAMon Agent Refuse Server Tests on the remote agent.

Note:

Command execution fails if you disable the SLA Monitor agent.

Limitations

SLA Monitor agent communications are IPv4–based. Agent communications do not currently support IPv6.

SLA Monitor configuration using CLI

Use the procedures in this section to configure the SLA Monitor agent.

Displaying SLA Monitor agent settings

Use this procedure to view the global SLA Monitor agent settings.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Display SLA Monitor agent settings:

   ```
   show application slamon agent
   ```

Example

```
Switch>enable
Switch#show application slamon agent

SLAMon Operational Mode: Disabled
SLAMon Agent Encryption: Supported
SLAMon Agent Address: 0.0.0.0
SLAMon Agent Port: 50011
SLAMon Agent Registration Status: Not Registered
SLAMon Registered Server Address: 0.0.0.0
SLAMon Registered Server Port: 0
SLAMon Server Registration Time: 0
SLAMon CLI Mode: Disabled
SLAMon CLI Timeout Mode: Enabled
SLAMon CLI Timeout: 60 seconds
SLAMon Configured Agent Address: 0.0.0.0
SLAMon Configured Agent Port: 0
SLAMon Configured Server Address: 0.0.0.0 0.0.0.0
SLAMon Configured Server Port: 0
SLAMon Agent-To-Agent Communication Port: 50012
SLAMon Configured Agent-To-Agent Communication Port: 0
SLAMon Agent Server Bypass: Disabled
SLAMon Agent Refuse Server Tests: Allow Tests
```

Configuring the SLA Monitor

Use this procedure to configure the SLA Monitor agent to communicate with an SLA Monitor server to perform Quality of Service (QoS) tests of the network.

Before you begin

To take full advantage of the SLA Monitor agent, you must have an SLA Monitor server in your network. The Quality of Service (QoS) tests can be performed without a server.

About this task

To configure the agent, you must enable the agent and assign an IP address. By default, the agent uses the switch/stack IP address if a specific agent address is not configured. Remaining agent parameters are optional and you can operate the agent using the default values.

Procedure

1. Enter Application Configuration mode:

   ```
   enable
   configure terminal
   application
   ```

2. Configure the agent IP address:

   ```
   slamon agent ip address {A.B.C.D}
   ```

3. Configure the agent IP address to its default value:

   ```
   default slamon agent ip address
   ```

4. Configure the UDP port:

   ```
   slamon agent port <0, 1024-65535>
   ```

5. Configure the agent UDP port to its default value:

   ```
   default slamon agent port
   ```

6. Enable the agent:

   ```
   slamon oper-mode enable
   ```

7. Disable the agent:

   ```
   no slamon oper-mode [enable]
   ```

   OR

   ```
   default slamon oper-mode
   ```

8. Configure the agent-to-agent communication port:

   ```
   slamon agent—comm—port <0, 1024-65535>
   ```

9. Configure the agent-to-agent communication port to its default value:

   ```
   default slamon agent—comm—port
   ```

10. Enable the SLA Monitor agent CLI support:

    ```
    slamon cli enable
    ```

    Note:

    The CLI commands from step 10 to 14 affect only the SLA Monitor (SLM) CLI commands and not the standard platform CLI commands.

11. Disable the SLA Monitor agent CLI support:

```
no slamon cli [enable]
```

OR

```
default slamon cli
```

12. Configure the agent automatic CLI session timeout value:

```
[default] slamon cli-timeout <60-600>
```

13. Enable the agent automatic CLI session timeout:

```
slamon cli-timeout—mode enable
```

OR

```
default slamon cli-timeout—mode
```

14. Disable the agent automatic CLI session timeout:

```
no slamon cli-timeout—mode [enable]
```

15. Configure the agent server IP address:

```
slamon server ip address {A.B.C.D} [{A.B.C.D}]
```

16. Configure the agent server IP address to its default value:

```
default slamon server ip address
```

17. Configure the server TCP registration port:

```
slamon server port <0-65535>
```

18. Configure the server TCP registration port to its default value:

```
default slamon server port
```

19. Enable the agent refuse server test mode:

```
slamon refuse-server-tests [enable]
```

20. Disable the agent refuse server test mode:

```
no slamon refuse-server-tests [enable]
```

OR

```
default slamon refuse-server-tests
```

21. Enable the agent server bypass mode:

```
slamon server-bypass [enable]
```

22. Disable the agent server bypass mode:

```
no slamon server-bypass [enable]
```

OR

```
default slamon server-bypass
```

23. Display the SLA monitor configuration:

```
show application slamon agent
```

Example

```
Switch>enable
Switch#configure terminal
Switch(config)#application
Switch(config-app)#slamon oper-mode enable
Switch(config-app)#show application slamon agent

SLAMon Operational Mode: Enabled
SLAMon Agent Encryption: Not Supported
SLAMon Agent Address: 192.0.2.1
SLAMon Agent Port: 50011
SLAMon Agent Registration Status: Not Registered
SLAMon Registered Server Address: 0.0.0.0
SLAMon Registered Server Port: 0
SLAMon Server Registration Time: 0
SLAMon CLI Mode: Disabled
SLAMon CLI Timeout Mode: Enabled
SLAMon CLI Timeout: 60 seconds
SLAMon Configured Agent Address: 0.0.0.0
SLAMon Configured Agent Port: 0
SLAMon Configured Server Address: 0.0.0.0 0.0.0.0
SLAMon Configured Server Port: 0
SLAMon Agent-To-Agent Communication Port: 50012
SLAMon Configured Agent-To-Agent Communication Port: 0
SLAMon Agent Server Bypass: Disabled
SLAMon Agent Refuse Server Tests: Allow Tests
```

Next steps

If you have configured SLA Monitor yet the agent is not functioning as expected, perform typical troubleshooting steps to verify agent accessibility:

- Verify IP address assignment and port use.
- Verify that the SLA Monitor agent is enabled.
- Ping the server IP address.
- Verify the server configuration.

If the agent is still not functioning, reset the system to ensure that the agent has started.

Variable definitions

The following table describes the parameters for the **`slamon`** command.

Variable	Value
agent	Configures the SLA Monitor agent.
agent-comm-port *<0, 1024-65535>*	Configures the SLA Monitor agent-to-agent communication UDP port.
agent ip address *<A.B.C.D>*	Configures the agent IP address. If no IP address is specified, the default value is 0.0.0.0, which causes the agent to use the switch/stack IP address.

Table continues…

Variable	Value
	Note: If you specify an IP address, ensure the address is a valid Layer 3 IPv4 address that is already configured for use by the switch.
agent port *<0, 1024-65535>*	Configures the UDP port for agent-server communication. The agent receives discovery packets on this port. The default is port 50011. The server must use the same port.
cli	Configures the SLA Monitor agent CLI interface.
cli-timeout *<60–600>*	Configures the CLI timeout value in seconds. The default is 60 seconds. **Note:** The CLI commands only impact the SLA Monitor CLI and not the standard platform CLI.
ntr	Initiates the SLA Monitor NTR test.
oper-mode	You can enable or disable the SLA Monitor agent. By default, SLA Monitor agent is disabled. If you disable the agent, it does not respond to discover packets from a server. If you disable the agent because of resource concerns, consider changing the server configuration instead, to alter the test frequency or duration, or the number of targets.
server ip address *{A.B.C.D} [{A.B.C.D}]*	Restricts the agent to use of this server IP address only. The default is 0.0.0.0, which means the agent can register with any server. You can specify a secondary server as well.
server port *<0–65535>*	Restricts the agent to use of this registration port only. The default is 0, which means the agent disregards the source port information in server traffic. The server must use the same port.
rtp	Initiates the SLA Monitor RTP test.
refuse-server-tests	Agent rejects NTR and RTP test requests from the server when this mode is enabled. If you disable this mode, the agent accepts test requests from the server with which it is registered. Test requests originating from platform, SLM CLI interfaces, and SNMP are not affected.

Table continues…

Variable	Value
server	Configures the SLA Monitor server.
server-bypass	You can enable or disable the SLA Monitor agent server-bypass mode. Allows an enabled agent to always accept agent-to-agent traffic. When enabled a small number of network ports remain open to process network traffic. You must take this into account if security concerns are high.

Executing NTR test using CLI

Use this procedure to execute a new trace route (NTR) test on the network to establish the QoS benchmark.

Before you begin

To execute the NTR test, you must enable the agent and assign an IP address.

Procedure

1. Enter Application Configuration mode:

   ```
   enable

   configure terminal

   application
   ```

2. Execute the NTR test:

   ```
   slamon ntr {A.B.C.D} <0-63>
   ```

Example

```
Switch>enable
Switch#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch(config)#application
Switch(config-app)#slamon oper-mode enable
Switch(config-app)#slamon ntr 192.0.2.1 46

-------------------------------------------------
SLAMon Network Trace Report
-------------------------------------------------
Source IP/Port: 192.0.2.2:50013
Source DSCP Marking: 46
Destination IP/Port: 192.0.2.1:33434
Maximum TTL: 1
Request Result: OK (Port unreachable)
                  Ingress  Egress
IP Address         DSCP     DSCP      RTT (ms)
-------------------------------------------------
192.0.2.2           46        0       0.000
192.0.2.1           0         0       1.240
```

Variable definitions

The following table describes the parameters for the **slamon ntr** command.

Variable	Value
IPv4 Address *<A.B.C.D>*	Specifies the destination IP address. If no IP address is specified, the test execution fails.
DSCP *<0–63>*	Specifies the Differential Services Code Point (DSCP) value for use in packets that are generated by the NTR test.
attempts *<1–10>*	Specifies the number of attempts generated by the NTR test. The default value is 2.
period *<10000–200000>*	Specifies the interval between packets in microseconds, generated by the NTR test. The default interval is 20000 microseconds.

Executing RTP test using CLI

Use this procedure to execute a real time protocol (RTP) test on the network to establish the QoS benchmark.

Before you begin

To execute the RTP test, you must enable the agent and assign an IP address.

Note:

You must enable the SLA Monitor agent ServerBypass mode for the RTP test to complete successfully.

Procedure

1. Enter Application Configuration mode:

   ```
   enable

   configure terminal

   application
   ```

2. Execute the RPT test:

   ```
   slamon RTP {A.B.C.D} <0-63>
   ```

Example

```
Switch>enable
Switch#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch(config)#application
Switch(config-app)#slamon oper-mode enable
Switch(config-app)#slamon rtp 192.0.2.1 46

-------------------------------------------------
SLAMon Real Time Protocol Network Report
```

```
-------------------------------------------------
Source IP/Port: 192.0.2.2:50012
Source DSCP Marking: 46
Destination IP/Port: 192.0.2.1:50012

Delay (RTT): average 1.824 (ms)  median 1.701 (ms)
Packet Loss: 0

Out-of-Order Arrivals:0

-------------------------------------------------
          Network Jitter - Quartiles (ms)
    0         1         2         3         4
-------------------------------------------------
    0.007     0.173     0.208     0.224     1.343
```

Variable definitions

The following table describes the parameters for the **`slamon rtp`** command.

Variable	Value
IPv4 Address *<A.B.C.D>*	Specifies the destination IP address. If no IP address is specified, the test execution fails.
DSCP *<0–63>*	Specifies the DSCP value for use in packets that are generated by the RTP test.
npack *<10–100>*	Specifies the number of test packets generated by the RTP test. Test packets are used to determine jitter. The value ranges from 10 to 100. The default value is 50.
nsync *<10–100>*	Specifies the number of synchronization packets generated by the RTP test. Synchronization packets are used to determine network delay. The value ranges from 10 to 100. The default value is 10.
period *<10000–200000>*	Specifies the interval between packets in microseconds, generated by the RTP test. The default interval is 20000 microseconds.

SLA Monitor Configuration using Enterprise Device Manager

A server is required to fully utilize the capabilities of the SLA Monitor agent. The agent can be used without a server.

The SLA Monitor agent must be enabled to run specific QoS tests in the absence of an SLA Monitor server. Agents exchange packets between one another to conduct the QoS tests. SLA Monitor uses Real Time Protocol (RTP) and New Trace Route (NTR) tests to determine QoS benchmarks.

Note:

SLA Monitor agent communications are IPv4-based. Agent communications do not currently support IPv6.

Use the following procedures to configure SLA Monitor using EDM

Configuring SLA Monitor using EDM

Use this procedure to configure SLA Monitor.

Procedure

1. In the navigation tree, double-click **Serviceability**.
2. In the Serviceability tree, click **SLA Monitor**.
3. In the **SLA Monitor** tab, configure parameters as required.
4. On the toolbar, click **Apply**.

Variable definition

Name	Description
Status	Enables or disables the SLA Monitor agent. The default is disabled. • **enabled**: enables the SLA Monitor agent • **disabled**: disables the SLA Monitor agent If you disable the agent, it does not respond to discover packets from a server. If you disable the agent because of resource concerns, consider changing the server configuration instead, to alter the test frequency or duration, or the number of targets.
ServerBypass	Enables or disables the SLA Monitor agent server bypass mode. • **enabled**: enables the SLA Monitor agent server bypass mode. • **disabled**: disables the SLA Monitor agent server bypass mode.
RefuseServerTests	Enables or disables the NTR and RTP test requests from the server. • **enabled**: the SLA Monitor agent rejects test requests from the server with which it is registered.

Table continues…

Name	Description
	• **disabled**: the SLA Monitor agent server accepts test requests from the server with which it is registered. Test requests originating from platform, SLM CLI interfaces, and SNMP are not affected.
ConfiguredAgentToAgentPort	Specifies the UDP port utilized by the SLA Monitor agent for agent-agent communication. If the value of this attribute is zero, the SLA Monitor agent utilizes a default port value for the base agent-agent UDP communication port.
ConfiguredAgentAddrType	Indicates IPv4–based communications.
ConfiguredAgentAddr	Specifies the agent IP address. The default value is 0.0.0.0, which causes the agent to use the switch/stack IP address. **Note:** If you specify an IP address, ensure the address is a valid Layer 3 IPv4 address that is already configured for use by the switch.
ConfiguredAgentPort	Specifies the UDP port for agent-server communication. The agent receives discovery packets on this port. The default is port 50011. The server must use the same port.
CliAvailable	Specifies whether SLA Monitor agent CLI is available or not available.
CliTimeout	Specifies the maximum amount of time, in seconds, until the CLI session is automatically terminated. The value of this attribute is pertinent only if CLI timeouts are enabled. The default is 60 seconds.
CliTimeoutMode	Configures whether the agent automatic CLI session timeout is enabled or disabled.
ConfiguredServerAddrType	Indicates IPv4–based communications.
ConfiguredServerAddr	Specifies the server IP address. If an IP address is specified, the agent is restricted to use this server IP address. The default is 0.0.0.0, which allows the agent to register with any server.
ConfiguredServerPort	Specifies the server port. The default is 0, which allows the agent to disregard the source port information in server traffic. The server must use the same port.
ConfiguredAltServerAddrType	Indicates IPv4–based communications.
ConfiguredAltServerAddr	Specifies a secondary server IP address.

Table continues…

Name	Description
SupportApps	Indicates SLA Monitor supported applications. This is a read-only field.
AgentAddressType	Indicates IPv4–based communications. This is a read-only field.
AgentAddress	Indicates the agent IP address. This is a read-only field.
AgentPort	Indicates the agent port. This is a read-only field.
RegisteredWithServer	Indicates whether the agent is registered with a server. This is a read-only field.
RegisteredServerAddrType	Indicates IPv4–based communications. This is a read-only field.
RegisteredServerAddr	Indicates IP address of the SLA Monitor server with which the agent is registered. This is a read-only field.
RegisteredServerPort	Indicates the TCP port used by the SLA Monitor server with which the agent is registered. This is a read-only field.
RegistrationTime	Indicates the time in seconds since the agent is registered with the server. This is a read-only field.
AgentToAgentPort	Indicates the base UDP port used by the SLA Monitor agent for agent-to-agent communication. The base UDP port is used to derive multiple agent communication ports. This is a read-only field.
EncryptionSupport	Indicates if encrypted agent-server communication is supported.

Executing NTR test using EDM

Use this procedure to execute NTR test on the network to establish QoS benchmark.

Important:

When executing the script using EDM, do not run other commands while the script is in progress, because this slows down the execution. EDM can time-out while waiting for a response and even when a time-out occurs, the script execution continues on EDM.

Procedure steps

1. From the navigation tree, double-click **Serviceability**.
2. In the Serviceability tree, double-click **SLA Monitor**.
3. In the SLA Monitor work area, click **NTR**.
4. In the NTR work area, click **Insert** to enter parameters for the new test.

5. In the **OwnerId** dialog box, type the owner id.
6. In the **TestName** dialog box, type the test name.
7. In the **TargetAddress** dialog box, type the target IP address.
8. In the **Dscp** dialog box, type the dscp value.
9. In the **Attempts** dialog box, type the number of attempts.
10. In the **Period** dialog box, type the duration in microseconds.
11. In the **Label** dialog box, type the label.
12. Click **enabled** to enable the administrator status.
13. Click **Insert** to initiate the NTR test.
14. In the NTR work area, click **Results** to view the test results.

Variable definition

Variable	Value
OwnerId	Specifies the owner of an NTR test.
TestName	Specifies the name of an NTR test.
TargetAddress	Specifies the target IP address for the NTR test.
Dscp	Specifies the Differential Services Code Point (DSCP) value for use in packets that are generated by the NTR test. The value ranges from 0 to 63.
Attempts	Specifies the number of attempts generated by the NTR test. The value ranges from 1 to 10. The default value is 2.
Period	Specifies the interval between packets in microseconds, generated by the NTR test. The value ranges from 10000 to 200000. The default interval is 20000 microseconds.
AdminStatus	Specifies the administrator status. You must enable the administrator status to initiate the NTR test. The administrator status is disabled by default.
Label	Specifies the text label used to reference the NTR control entry.

NTR field descriptions

Name	Description
OwnerId	Specifies the owner of an NTR test.
TestName	Specifies the name of an NTR test.
TargetAddress	Specifies the target IP address for the NTR test.
Dscp	Specifies the Differential Services Code Point (DSCP) value for use in packets that are generated by the NTR test.
Attempts	Specifies the number of attempts generated by the NTR test. The default value is 2.

Table continues…

Name	Description
Period	Specifies the interval between packets in microseconds, generated by the NTR test. The default interval is 20000 microseconds.
Label	Specifies the text label used to reference the NTR control entry.
AdminStatus	Specifies the administrator status. You must enable the administrator status to initiate the NTR test. The administrator status is disabled by default.

Viewing NTR test results

Use this procedure to view the NTR test results.

Before you begin

You must execute the NTR test before you view the results.

Procedure

1. In the navigation tree, double-click **Serviceability** .
2. In the Serviceability tree, click **SLA Monitor** .
3. In the SLA Montior work area, click **NTR** .
4. In the NTR work area, click to select the saved test and then click **Results** .
5. In the results work area, click **NTR Results** to view the NTR test results.

Variable definition

Name	Description
HopIndex	Indicates the hop index for an NTR test hop.
TgtAddress	Indicates the IP address associated with the NTR test hop.
Rtt	Indicates the round-trip-time of an NTR test in milliseconds.
IngressDscp	Indicates the DSCP value in the NTR test packet received by the end station for the specified hop.
EgressDscp	Indicates the DSCP value in the NTR test packet received by the SLA Monitor agent for the specified hop.

Executing RTP test using EDM

Use this procedure to execute RTP test on the network to establish QoS benchmark.

Important:

When executing the script using EDM, do not run other commands while the script is in progress, because this slows down the execution. EDM can time-out while waiting for a response and even when a time-out occurs, the script execution continues on EDM.

Note:

The ServerBypass must be enabled on the target to complete the test successfully.

Procedure steps

1. From the navigation tree, double-click **Serviceability**.
2. In the Serviceability tree, double-click **SLA Monitor**.
3. In the SLA Monitor work area, click **RTP**.
4. In the RTP work area, click **Insert** to enter parameters for the new test.
5. In the **OwnerId** dialog box, type the owner id.
6. In the **TestName** dialog box, enter the test name.
7. In the **TargetAddress** dialog box, type the target IP address.
8. In the **Dscp** dialog box, type the dscp value.
9. In the **TestPackets** dialog box, type the number of test packets.
10. In the **SyncPackets** dialog box, type the number of synchronization packets.
11. In the **Period** dialog box, type the duration in microseconds.
12. Click **enabled** to enable the administrator status.
13. In the **Label** dialog box, type the label.
14. Click **Insert** to initiate the RTP test.
15. In the RTP work area, click **Results** to view the test results.

Variable Definition

Variable	Value
OwnerId	Specifies the owner of an RTP test.
TestName	Specifies the name of an RTP test.
TargetAddress	Specifies the target IP address for the RTP test.
Dscp	Specifies the Differential Services Code Point (DSCP) value for use in packets that are generated by the RTP test. The value ranges from 0 to 63.
TestPackets	Specifies the number of test packets generated by the RTP test. Test packets are used to determine jitter. The value ranges from 10 to 100.
SyncPackets	Specifies the number of synchronization packets generated by the RTP test. Synchronization packets are used to determine network delay. The value ranges from 10 to 100.

Table continues…

Variable	Value
Period	Specifies the interval between packets in microseconds, generated by the RTP test. The value ranges from 10000 to 200000. The default interval is 20000 microseconds.
AdminStatus	Specifies the administrator status. You must enable the administrator status to initiate the RTP test. The administrator status is disabled by default.
Label	Specifies the text label used to reference the RTP control entry.

RTP field descriptions

Name	Description
OwnerId	Specifies the owner of an RTP test.
TestName	Specifies the name of an RTP test.
TargetAddress	Specifies the target IP address for the RTP test.
Dscp	Specifies the Differential Services Code Point (DSCP) value for use in packets that are generated by the RTP test.
TestPackets	Specifies the number of test packets generated by the RTP test. Test packets are used to determine jitter.
SyncPackets	Specifies the number of synchronization packets generated by the RTP test. Synchronization packets are used to determine network delay.
Period	Specifies the interval between packets in microseconds, generated by the RTP test. The default interval is 20000 microseconds.
Label	Specifies the text label used to reference the RTP control entry.
AdminStatus	Specifies the administrator status. You must enable the administrator status to initiate the RTP test. The administrator status is disabled by default.

Viewing real time protocol test results

Use this procedure to view the RTP test results.

Before you begin

You must execute the RTP test before you view the results.

Procedure

1. In the navigation tree, double-click **Serviceability** .
2. In the Serviceability tree, click **SLA Monitor** .

3. In the SLA Montior work area, click **RTP** .
4. In the RTP work area, click to select the saved test and then click **Results** to view the RTP test results.

Variable definitions

Name	Description
OperStatus	Indicates the status of an RTP test. • **inProgess** indicates that an RTP test is in progress. • **aborted** indicates that an RTP test is aborted. • **completed** indicates that an RTP test is completed.
SrcAddress	Indicates the source IP address used for the RTP test.
SrcPort	Indicates the port used for the RTP test.
DstAddress	Indicates the destination IP address used for the RTP test.
DstPort	Indicates the destination port used for the RTP test.
Dscp	Specifies the Differential Services Code Point (DSCP) value for use in packets that are generated by the RTP test.
AverageDelay	Indicates the average network delay (RTT) experienced during the RTP test execution in microseconds.
MedianDelay	Indicates the median network delay (RTT) experienced during the RTP test execution in microseconds.
PacketLoss	Indicates the count of packets lost during an RTP test execution.
OutOfOrderArrivals	Indicates the count of packets arriving out-of-order during an RTP test execution.
JitterQuartile0 – JitterQuartile5	Indicates the resulting quartile boundaries after sorting the network jitter values of all test packets during the RTP test execution. The value is represented in microseconds.
AbortData	Indicates the details of the RTP test that was aborted.

Glossary

Application-specific Integrated Circuit (ASIC)	An application-specific integrated circuit developed to perform more quickly and efficiently than a generic processor.
base unit (BU)	When you connect multiple switches into a stack, one unit, and only one unit, must be designated as a base unit to perform stack configuration tasks. The position of the unit select switch, on the back of the switch, determines base unit designation.
Bridge Protocol Data Unit (BPDU)	A data frame used to exchange information among the bridges in local or wide area networks for network topology maintenance.
cascade down	Refers to the stack configuration. The system automatically numbers the physical units based on the designated base unit, which is Unit 1. In the cascade down configuration, the base unit is physically located on the top of the stack and stacking cables are connected in the appropriate order.
cascade up	Refers to the stack configuration. The system automatically numbers the physical units based on the designated base unit, which is Unit 1. In the cascade up configuration, the base unit is physically located on the bottom of the stack and stacking cables are connected in the appropriate order.
CLI	Command Line Interface (CLI) is a text-based, common command line interface used for device configuration and management across Extreme Networks products.
Distributed MultiLink Trunking (DMLT)	A point-to-point connection that aggregates similar ports from different modules to logically act like a single port, but with the aggregated bandwidth.
Enterprise Device Manager (EDM)	A web-based embedded management system to support single-element management. EDM provides complete configuration management functionality for the supported devices and is supplied to the customer as embedded software in the device.
Frame Check Sequence (FCS)	Frames are used to send upper-layer data and ultimately the user application data from a source to a destination.

Internet Control Message Protocol (ICMP)	A collection of error conditions and control messages exchanged by IP modules in both hosts and gateways.
Internet Group Management Protocol (IGMP)	IGMP is a host membership protocol used to arbitrate membership in multicast services. IP multicast routers use IGMP to learn the existence of host group members on their directly attached subnets.
Internet Protocol Flow Information eXport (IPFIX)	An IETF standard that improves the Netflow V9 protocol. IPFIX monitors IP flows.
Internet Protocol version 4 (IPv4)	The protocol used to format packets for the Internet and many enterprise networks. IPv4 provides packet routing and reassembly.
Internet Protocol version 6 (IPv6)	An improved version of the IP protocol, IPv6 improves the IPv4 limitations of security and user address numbers.
light emitting diode (LED)	A semiconductor diode that emits light when a current passes through it.
Light Emitting Diode (LED)	The switch displays diagnostic and operational information through the LEDs on the unit. For detailed information regarding the interpretation of the LEDs, see Installing Ethernet Routing Switch 5900 Series.
Link Aggregation	Provides the mechanism to create and manage trunk groups automatically using Link Aggregation Control Protocol (LACP).
Link Aggregation Control Protocol (LACP)	A network handshaking protocol that provides a means to aggregate multiple links between appropriately configured devices.
Link Aggregation Control Protocol Data Units (LACPDU)	Link aggregation control protocol data unit (LACPDU) is used for exchanging information among LACP-enabled devices.
Link Aggregation Group (LAG)	A group that increases the link speed beyond the limits of one single cable or port, and increases the redundancy for higher availability.
Logical Link Control (LLC)	A protocol used in LANs to transmit protocol data units between two end stations. This LLC layer addresses and arbitrates data exchange between two endpoints.
management information base (MIB)	The MIB defines system operations and parameters used for the Simple Network Management Protocol (SNMP).
mask	A bit string that the device uses along with an IP address to indicate the number of leading bits in the address that correspond with the network part.

media	A substance that transmits data between ports; usually fiber optic cables or category 5 unshielded twisted pair (UTP) copper wires.
Media Access Control (MAC)	Arbitrates access to and from a shared medium.
mirrored port	The port to mirror. The port is also called the source port.
mirroring port	The port to which the system mirrors all traffic, also referred to as the destination port.
MultiLink Trunking (MLT)	A method of link aggregation that uses multiple Ethernet trunks aggregated to provide a single logical trunk. A multilink trunk provides the combined bandwidth of multiple links and the physical layer protection against the failure of a single link.
multiplexing	Carriage of multiple channels over a single transmission medium; a process where a dedicated circuit is shared by multiple users. Typically, data streams intersperse on a bit or byte basis (time division), or separate by different carrier frequencies (frequency division).
NonVolatile Random Access Memory (NVRAM)	Random Access Memory that retains its contents after electrical power turns off.
port	A physical interface that transmits and receives data.
port mirroring	A feature that sends received or transmitted traffic to a second destination.
port VLAN ID	Used to coordinate VLANs across multiple switches. When you create a port-based VLAN on a switch, assign a VLAN identification number (VLAN ID) and specify the ports that belong to the VLAN.
Power over Ethernet (PoE)	The capacity of a switch to power network devices, according to the 802.3af standard, over an Ethernet cable. Devices include IP phones, Wireless LAN Access Points (WLAN AP), security cameras, and access control points.
Protocol Data Units (PDUs)	A unit of data that is specified in a protocol of a specific layer and that consists of protocol-control information of the specific layer and possibly user data of that layer.
quality of service (QoS)	QoS features reserve resources in a congested network, allowing you to configure a higher priority for certain devices. For example, you can configure a higher priority for IP deskphones, which need a fixed bit rate and split the remaining bandwidth between data connections if calls in the network are more important than the file transfers.

Remote Network Monitoring (RMON)	Creates and displays alarms for user-defined events, gathers cumulative statistics for Ethernet interfaces, and tracks statistical history for Ethernet interfaces.
routing switch	Virtualizes the physical router interfaces to switches. A virtual router port, or interface, acts as a router port to consolidate switching and routing functions in the broadcast domain, or between broadcast domains, and enable IP routing for higher traffic volumes.
spanning tree	A simple, fully-connected active topology formed from the arbitrary physical topology of connected bridged Local Area Network components by relaying frames through selected bridge ports. The protocol parameters and states that are used and exchanged to facilitate the calculation of the active topology and to control the bridge relay function.
Spanning Tree Group (STG)	A collection of ports in one spanning-tree instance.
Spanning Tree Protocol (STP)	MAC bridges use the STP to exchange information across Local Area Networks to compute the active topology of a bridged Local Area Network in accordance with the Spanning Tree Protocol algorithm.
stack	Stackable Extreme Networks Ethernet Routing Switch can be connected in a stack configuration of two or more units, up to eight units maximum. A switch stack operates and is managed as a single virtual switch.
stand-alone	Refers to a single Extreme Networks Ethernet Routing Switch operating outside a stack.
Temporary Base Unit (TBU)	If an assigned base unit in a stack fails, the next unit in the stack automatically becomes the temporary base unit (TBU). The TBU maintains stack operations until the stack is restarted or the TBU fails. If the old base unit rejoins the stack, it does not take over from the TBU until the stack is reset.
time-to-live (TTL)	The field in a packet used to determine the valid duration for the packet. The TTL determines the packet lifetime. The system discards a packet with a TTL of zero.
Transmission Control Protocol (TCP)	Provides flow control and sequencing for transmitted data over an end-to-end connection.
trunk	A logical group of ports that behaves like a single large port.
type of service (TOS)	A field in the IPv4 header that determines the Class of Service prior to the standardization of Differentiated Services.

User Datagram Protocol (UDP)	In TCP/IP, a packet-level protocol built directly on the Internet Protocol layer. TCP/IP host systems use UDP for application-to-application programs.
Virtual Link Aggregation Control Protocol (VLACP)	Virtual Link Aggregation Control Protocol (VLACP) is a Layer 2 handshaking protocol that can detect end-to-end failure between two physical Ethernet interfaces.
Virtual Local Area Network (VLAN)	A Virtual Local Area Network is a group of hosts that communicate as if they are attached to the same broadcast domain regardless of their physical location. VLANs are layer 2 constructs.
Voice over IP (VOIP)	The technology that delivers voice information in digital form in discrete packets using the Internet Protocol (IP) rather than the traditional circuit-committed protocols of the public switched telephone network (PSTN).
wiring closet	A central termination area for telephone or network cabling or both.