NORTEL

Nortel Ethernet Routing Switch 5500 Series

Configuration - System Monitoring

Trademarks
*Nortel, Nortel Networks, the Nortel logo and the Globemark are trademarks of Nortel Networks.

The asterisk (*) after a name denotes a trademarked item.

Restricted rights legend
Use, duplication, or disclosure by the United States Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013.

Notwithstanding any other license agreement that may pertain to, or accompany the delivery of, this computer software, the rights of the United States Government regarding its use, reproduction, and disclosure are as set forth in the Commercial Computer Software-Restricted Rights clause at FAR 52.227-19.

Statement of conditions
In the interest of improving internal design, operational function, and/or reliability, Nortel Networks reserves the right to make changes to the products described in this document without notice.

Nortel Networks does not assume any liability that may occur due to the use or application of the product(s) or circuit layout(s) described herein.

Nortel Networks software license agreement
This Software License Agreement ("License Agreement") is between you, the end-user ("Customer") and Nortel Networks Corporation and its subsidiaries and affiliates ("Nortel Networks"). PLEASE READ THE FOLLOWING CAREFULLY. YOU MUST ACCEPT THESE LICENSE TERMS IN ORDER TO DOWNLOAD AND/OR USE THE SOFTWARE. USE OF THE SOFTWARE CONSTITUTES YOUR ACCEPTANCE OF THIS LICENSE AGREEMENT. If you do not accept these terms and conditions, return the Software, unused and in the original shipping container, within 30 days of purchase to obtain a credit for the full purchase price.

"Software" is owned or licensed by Nortel Networks, its parent or one of its subsidiaries or affiliates, and is copyrighted and licensed, not sold. Software consists of machine-readable instructions, its components, data, audio-visual content (such as images, text, recordings or pictures) and related licensed materials including all whole or partial copies. Nortel Networks grants you a license to use the Software only in the country where you acquired the Software. You obtain no rights other than those granted to you under this License Agreement. You are responsible for the selection of the Software and for the installation of, use of, and results obtained from the Software.

1. Licensed Use of Software. Nortel Networks grants Customer a nonexclusive license to use a copy of the Software on only one machine at any one time or to the extent of the activation or authorized usage level, whichever is applicable. To the extent Software is furnished for use with designated hardware or Customer furnished equipment ("CFE"), Customer is granted a nonexclusive license to use Software only on such hardware or CFE, as applicable. Software contains trade secrets and Customer agrees to treat Software as confidential information using the same care and discretion Customer uses with its own similar information that it does not wish to disclose, publish or disseminate. Customer will ensure that anyone who uses the Software does so only in compliance with the terms of this Agreement. Customer shall not a) use, copy, modify, transfer or distribute the Software except as expressly

authorized; b) reverse assemble, reverse compile, reverse engineer or otherwise translate the Software; c) create derivative works or modifications unless expressly authorized; or d) sublicense, rent or lease the Software. Licensors of intellectual property to Nortel Networks are beneficiaries of this provision. Upon termination or breach of the license by Customer or in the event designated hardware or CFE is no longer in use, Customer will promptly return the Software to Nortel Networks or certify its destruction. Nortel Networks may audit by remote polling or other reasonable means to determine Customer's Software activation or usage levels. If suppliers of third party software included in Software require Nortel Networks to include additional or different terms, Customer agrees to abide by such terms provided by Nortel Networks with respect to such third party software.

2. Warranty. Except as may be otherwise expressly agreed to in writing between Nortel Networks and Customer, Software is provided "AS IS" without any warranties (conditions) of any kind. NORTEL NETWORKS DISCLAIMS ALL WARRANTIES (CONDITIONS) FOR THE SOFTWARE, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO THE IMPLIED WARRANTIES OF MERCHANTABLITITY AND FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OF NON-INFRINGEMENT. Nortel Networks is not obligated to provide support of any kind for the Software. Some jurisdictions do not allow exclusion of implied warranties, and, in such event, the above exclusions may not apply.

3. Limitation of Remedies. IN NO EVENT SHALL NORTEL NETWORKS OR ITS AGENTS OR SUPPLIERS BE LIABLE FOR ANY OF THE FOLLOWING: a) DAMAGES BASED ON ANY THIRD PARTY CLAIM; b) LOSS OF, OR DAMAGE TO, CUSTOMER'S RECORDS, FILES OR DATA; OR c) DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES (INCLUDING LOST PROFITS OR SAVINGS), WHETHER IN CONTRACT, TORT OR OTHERWISE (INCLUDING NEGLIGENCE) ARISING OUT OF YOUR USE OF THE SOFTWARE, EVEN IF NORTEL NETWORKS, ITS AGENTS OR SUPPLIERS HAVE BEEN ADVISED OF THEIR POSSIBILITY. The foregoing limitations of remedies also apply to any developer and/or supplier of the Software. Such developer and/or supplier is an intended beneficiary of this Section. Some jurisdictions do not allow these limitations or exclusions and, in such event, they may not apply.

4. General

a) If Customer is the United States Government, the following paragraph shall apply: All Nortel Networks Software available under this License Agreement is commercial computer software and commercial computer software documentation and, in the event Software is licensed for or on behalf of the United States Government, the respective rights to the software and software documentation are governed by Nortel Networks standard commercial license in accordance with U.S. Federal Regulations at 48 C.F.R. Sections 12.212 (for non-DoD entities) and 48 C.F.R. 227.7202 (for DoD entities).

b) Customer may terminate the license at any time. Nortel Networks may terminate the license if Customer fails to comply with the terms and conditions of this license. In either event, upon termination, Customer must either return the Software to Nortel Networks or certify its destruction.

c) Customer is responsible for payment of any taxes, including personal property taxes, resulting from Customer's use of the Software. Customer agrees to comply with all applicable laws including all applicable export and import laws and regulations.

d) Neither party may bring an action, regardless of form, more than two years after the cause of the action arose.

e) The terms and conditions of this License Agreement form the complete and exclusive agreement between Customer and Nortel Networks.

f) This License Agreement is governed by the laws of the country in which Customer acquires the Software. If the Software is acquired in the United States, then this License Agreement is governed by the laws of the state of New York.

Revision History

Date Revised	Version	Reason for revision
July 2005	1.00	New document for Software Release 4.2.
July 2006	2.00	Document updated for Software Release 5.0.
August 2007	3.01	Document updated for Software Release 5.1.

Contents

Nortel Ethernet Routing Switch 5500 Series
Configuration - System Monitoring
NN47200-505   03.01   Standard
5.1   27 August 2007

Preface

This guide provides information and instructions on the configuration and usage of system monitoring tools on the 5500 Series Nortel Ethernet Routing Switch. Please consult any documentation included with the switch and the product release notes (see ""Related publications" (page 12)") for any errata before beginning the configuration process.

Nortel Ethernet Routing Switch 5500 Series

" 5500 Series Switch Platforms" (page 11)outlines the switches that are part of the 5500 Series of Nortel Ethernet Routing Switches.

5500 Series Switch Platforms

5500 Series Switch Model	Key Features
Nortel Ethernet Routing Switch 5510-24T	A 24 port, 10/100/1GBase-T, Layer 4, diffserv-capable, stackable Ethernet switch. This switch contains two shared SFP ports.
Nortel Ethernet Routing Switch 5510-48T	A 48 port, 10/100/1GBase-T, Layer 4, diffserv-capable, stackable Ethernet switch. This switch contains two shared SFP ports.
Nortel Ethernet Routing Switch 5520-24T-PWR	A 24 port, 10/100/1GBase-T, Layer 4, diffserv-capable, stackable Ethernet switch with full Power over Ethernet (PoE) capability on all copper ports. This switch contains four shared SFP ports.
Nortel Ethernet Routing Switch 5520-48T-PWR	A 48 port, 10/100/1GBase-T, Layer 4, diffserv-capable, stackable Ethernet switch with full Power over Ethernet (PoE) capability on all copper ports. This switch contains four shared SFP ports.
Nortel Ethernet Routing Switch 5530-24TFD	A 24 port, 10/100/1GBase-T, Layer 4, diffserv-capable, stackable Ethernet switch. This switch contains twelve shared SFP ports and two XFP ports.

Related publications

For more information about the management, configuration, and usage of the Nortel Ethernet Routing Switch 5500 Series, refer to the publications listed in " Nortel Ethernet Routing Switch 5500 Series Documentation" (page 12).

Nortel Ethernet Routing Switch 5500 Series Documentation

Title	Description	Part Number
Nortel Ethernet Routing Switch 5500 Series Installation	Instructions for the installation of a switch in the Nortel Ethernet Routing Switch 5500 Series. It also provides an overview of hardware key to the installation, configuration, and maintenance of the switch.	NN47200-300
Nortel Ethernet Routing Switch 5500 Series Configuration - System	Instructions for the general configuration of switches in the 5500 Series that are not covered by the other documentation.	NN47200-500
Nortel Ethernet Routing Switch 5500 Series Configuration - Security	Instructions for the configuration and management of security for switches in the 5500 Series.	NN47200-501
Nortel Ethernet Routing Switch 5500 Series Configuration - VLANs, Spanning Tree, and Link Aggregation	Instructions for the configuration of spanning and trunking protocols on 5500 Series switches	NN47200-502
Nortel Ethernet Routing Switch 5500 Series Configuration - IP Routing Protocols	Instructions on the configuration of IP routing protocols on 5500 Series switches.	NN47200-503
Nortel Ethernet Routing Switch 5500 Series Configuration - Quality of Service	Instructions on the configuration and implementation of QoS and filtering on 5500 Series switches.	NN47200-504
Nortel Ethernet Routing Switch 5500 Series Configuration - System Monitoriing	Instructions on the configuration, implementation, and usage of system monitoring on 5500 Series switches.	NN47200-505
Nortel Ethernet Routing Switch 5500 Series Release Notes - Software Release 5.1	Provides an overview of new features, fixes, and limitations of the 5500 Series switches. Also included are any supplementary documentation and document errata.	NN47200-400

Title	Description	Part Number
Installing the Nortel Ethernet Redundant Power Supply 15	Instructions for the installation and usage of the Nortel Ethernet RPS 15.	217070-A
DC-DC Converter Module for the Baystack 5000 Series Switch	Instructions for the installation and usage of the DC-DC power converter.	215081-A
Nortel Ethernet Routing Switch 5500 Series Installation - SFP	Instructions for the installation and usage of SFP and XFP transceivers and GBICs.	318034-C

Finding the latest updates on the Nortel web site

The content of this documentation was current at the time of release. To check for updates to the documentation and software for the Nortel Ethernet Routing Switch 5500 Series, use the links provided in the following table.

Software	Nortel Ethernet Routing Switch 5500 Series Software
Documentation	Nortel Ethernet Routing Switch 5500 Series Documentation

How to get help

If a service contract for the Nortel product has been purchased from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance.

If a Nortel service program was purchased, contact Nortel Technical Support.

The following information is available online:

- contact information for Nortel Technical Support
- information about the Nortel Technical Solutions Centers
- information about the Express Routing Code (ERC) for your product

An ERC is available for many Nortel products and services. When an ERC is used, the call is routed to technical support personnel who specialize in supporting the service or product. The ERC for a particular product or service is available online.

The main Nortel support portal is available at http://www.nortel.com/support.

Chapter 1
General System Monitoring Considerations

System monitoring is an important aspect of switch operation. The Nortel Ethernet Routing Switch 5500 Series provides a wide range of system monitoring options that allow the administrator to closely follow the operation of a switch or stack.

This chapter notes two general system monitoring considerations, system logging and port mirroring, that must be taken into account when using the Nortel Ethernet Routing Switch 5500 Series. Subsequent chapters provide information on specific system monitoring tools and their use.

System logging

The Nortel Ethernet Routing Switch 5500 Series supports system logging (syslog), a software tool to log system events for debugging and analysis.

To utilize the syslog, the switch applications that run in the Nortel Ethernet Routing Switch 5500 Series, such as IGMP, MLT, STP, should be registered with the syslog tool.

Any events that happen in the above-mentioned applications can be logged with the help of the syslog tool. The logged events are stored in volatile RAM, non-volatile RAM, or in a remote host. The storage location can be selected using the Command Line Interface (CLI).

Configuring the system log with the CLI

This section outlines the CLI commands used in the configuration and management of the system log.

show logging command

The `show logging` command displays the configuration, and the current contents, of the system event log.

The syntax for the `show logging` command is:

```
show logging [config] [critical] [serious] [informational]
[sort-reverse]
```

The **show logging** command is executed in the Privileged EXEC command mode.

" show logging parameters" (page 16) describes the parameters for this command.

show logging parameters

Parameter	Description
config	Display configuration of event logging.
critical	Display critical log messages.
serious	Display serious log messages.
informational	Display informational log messages.
sort-reverse	Display informational log messages in reverse chronological order (beginning with most recent).

logging command

The **logging** command configures the system settings for the system event log.

The syntax for the logging command is:

```
logging [enable | disable] [level critical | serious |
informational | none] [nv-level critical | serious | none]
```

The **logging** command is executed in the Global Configuration command mode.

" logging parameters" (page 16) describes the parameters for this command.

logging parameters

Parameter	Description
enable | disable	Enables or disables the event log (default is Enabled).
level critical | serious | informational | none	Specifies the level of logging stored in DRAM.
nv-level critical | serious | none	Specifies the level of logging stored in NVRAM.

no logging command

The **no logging** command disables the system event log.

The syntax for the `no logging` command is:

`no logging`

The `no logging` command is executed in the Global Configuration command mode.

default logging command

The `default logging` command configures the system settings as the factory default settings for the system event log.

The syntax for the `default logging` command is:

`default logging`

The `default logging` command is executed in the Global Configuration command mode.

clear logging command

The `clear logging` command clears all log messages in DRAM.

The syntax for the `clear logging` command is:

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

" clear logging parameters" (page 17) outlines the parameters for this command.

clear logging parameters

Parameter	Description
non-volatile	Clears log messages from NVRAM.
nv	Clears log messages from NVRAM and DRAM.
volatile	Clears log messages from DRAM.

The `clear logging` command is executed in the Privileged EXEC command mode.

Viewing the system log in the Web-based Management Interface

The Web-based Management Interface can be used to view the System Log. To perform this action, follow this procedure:

Step	Action

1   Open the **System Log** screen by selecting **Fault > System Log** from the menu. This screen is illustrated in "System Log screen" (page 18).

System Log screen
Fault > System Log

System Log (View By)	
Display Messages From	Non Volatile
Clear Messages From	None

Submit

System Log			
Index	Time Stamp	Message Type	Message
1	0D: 0H: 0M: 2S	Serious	Audit data initialized (bad checksum)

2   In the **System Log (View By)** section, select the messages to be
displayed by selecting a value from the **Display Messages From** list.

3   Click **Submit**.

—End—

Messages can also be cleared from the log by selecting a value from the
Clear Messages From list and then clicking **Submit**. If messages are not
to be cleared, ensure that this list has **None** selected.

Configuring the system log with the Java Device Manager
The Java Device Manager (JDM) also provides functionality for managing
the system log. To configure the system log, follow this procedure:

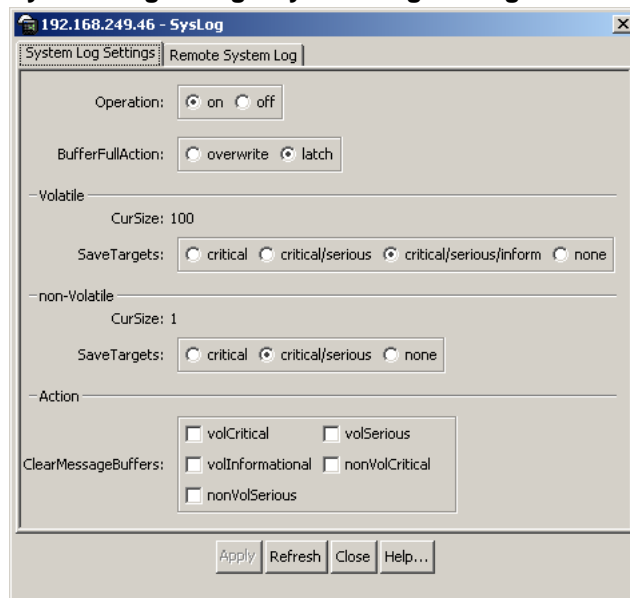**Step   Action**

---

1   Open the **System Log** screen by selecting **Edit > Diagnostics >
System Log** from the menu. Select the **System Log Settings**
tab. This screen is illustrated in "System Log dialog - System Log
Settings tab" (page 19).

System Log dialog - System Log Settings tab

2   In the fields provided, configure the system log settings. The following table outlines the fields on this screen.

System Log Settings fields

Field	Description
Operation	Turns the system log on or off.
BufferFullAction	Specifies whether the system log overwrites itself or discontinues the storage of messages when the buffer is full.
Volatile - CurSize	Shows the current number of messages stored in volatile memory.
Volatile - SaveTargets	Selects the severity of system messages to save.
non-Volatile - CurSize	Shows the current number of messages stored in non-volatile memory.
non-Volatile - SaveTargets	Selects the severity of system messages to save.
ClearMessageBuffers	Selects the sections of the system log to delete.

3   Click **Apply**.

—**End**—

Remote logging

The remote logging feature in Software Release 5.0 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.

The remote syslog server must be configured and set up on the unit to log informational messages to this remote server. The UDP packet is sent to port 514 of the configured remote syslog server.

Once 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 once the IP address of the remote server is on the system.

This feature can be configured 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.

Configuring remote logging with the CLI

Use the CLI to configure remote logging. This section discusses the commands that enable remote logging.

show logging command

The `show logging` command displays the configuration and the current contents of the system event log. Refer to ""show logging command" (page 15)" for an explanation of this command.

logging remote enable command

> *Note:* The default value for remote logging is Disabled

The `logging remote enable` command enables the use of a remote syslog server. The syntax for the `logging remote enable` command is:

`logging remote enable`

The `logging remote enable` command is executed in the Global Configuration command mode.

no logging remote enable command

The `no logging remote enable` command disables the use of a remote syslog server. The syntax for the `no logging remote enable` command is:

`no logging remote enable`

The `no logging remote enable` command is executed in the Global Configuration command mode.

logging remote address command

The `logging remote address` command sets the remote server for receiving the syslog messages. The syntax for the `logging remote address` command is:

`logging remote address <A.B.C.D>`

The `logging remote address` command is executed in the Global Configuration command mode.

"logging remote address parameters" (page 21) describes the parameters for the `logging remote address` command.

logging remote address parameters

Parameters and variables	Description
<A.B.C.D>	Specifies the IP address of the remote server in dotted-decimal notation.

The default address is 0.0.0.0.

no logging remote address command

The `no logging remote address` command clears the IP address of the remote server. The syntax for the `no logging remote address` command is:

`no logging remote address`

The `no logging remote address` command is executed in the Global Configuration command mode.

logging remote level command

The `logging remote level` command sets the severity level of the logs sent to the remote server. The syntax for the logging remote level command is:

`logging remote level {critical | informational | serious | none}`

The `logging remote level` command is executed in the Global Configuration command mode.

" logging remote level parameters" (page 22) describes the parameters for the `logging remote level` command.

logging remote level parameters

Parameters and variables	Description
{critical | serious | informational | none}	Specifies the severity level of the log messages to be sent to the remote server: • critical • informational • serious • none

no logging remote level command

The `no logging remote level` command removes any severity level setting and reverts to None. The syntax for the `no logging remote level` command is:

`no logging remote level`

The `no logging remote level` command is executed in the Global Configuration command mode.

default logging remote level command

The `default logging remote level` command sets the severity level of the logs sent to the remote server to the default value of None. The syntax for the `default logging remote level` command is:

`default logging remote level`

The `default logging remote level` command is executed in the Global Configuration command mode.

## Configuring remote logging with the Java Device Manager

The Java Device Manager (JDM) also provides functionality for managing remote logging. To configure remote logging, follow this procedure:

Step	Action
1	Open the **System Log** screen by selecting **Edit > Diagnostics > System Log** from the menu. Select the **Remote System Log** tab. This tab is illustrated below.

System Log dialog - Remote System Log tab
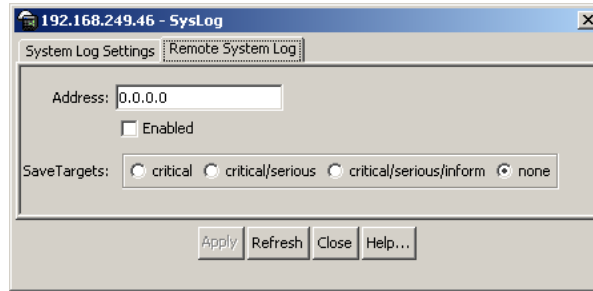
2   In the fields provided, enter the remote logging information. The following table describes the fields on this screen.

Remote System Log tab fields

Field	Description
Address	The IP address of the remote syslog server.
Enabled	Enables or disables remote logging.
SaveTargets	Sets the severity level of messages that are saved to the remote server.

3   Click **Apply**.

—End—

IGMP and the system event log

IGMP utilizes the components provided by the syslog tool. Functions such as storing messages in the NVRAM or remote host, and displaying these log messages through the CLI, console menu, 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 from the user

- Stack Join events

- IGMP messages -- Report, Leave and Query messages received by the switch

  Note: 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 in all the time. By default, logging in of such messages is disabled. This feature must be enabled through the CLI to view such messages.

In " IGMP syslog messages" (page 24):

- %d represents a decimal value for the parameter preceding it. For example, 5 for VLAN 5

- %x represents a hexadecimal value for the parameter preceding it. For example, 0xe0000a01 for Group 224.0.10.1

" IGMP syslog messages" (page 24) describes the IGMP syslog messages and their severity.

IGMP syslog messages

Severity	Log Messages
Informational	IGMP initialization success
Critical	IGMP initialization failed: Error code %d
Informational	IGMP policy initialized
Informational	IGMP configuration loaded successfully
Informational	IGMP configuration failed. Loaded to factory default
Informational	IGMP configuration changed: Snooping enabled on VLAN %d
Informational	IGMP configuration changed: Snooping disabled on VLAN %d
Informational	IGMP configuration changed: Proxy enabled on VLAN %d
Informational	IGMP configuration changed: 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 enabled
Informational	IGMP configuration changed: Trunk %d created for IGMP

Severity	Log Messages
Informational	IGMP configuration changed: Trunk %d removed for IGMP 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 exchange database sent by unit %d
Informational	IGMP exchange database received on unit %d from %d
Informational	IGMP exchange database done
Informational	IGMP stack join completed
Serious	IGMP not able to join stack: Error code %d
Informational	IGMP exchange group database sent by unit %d
Informational	IGMP exchange group database received on unit %d from %d
Informational	IGMP received report on VLAN %d for Group 0x%x on port %d
Informational	IGMP received leave on VLAN %d for Group 0x%x on port %d
Informational	IGMP received query on VLAN %d for Group 0x%x on port %d
Informational	IGMP dynamic router port %d added
Informational	IGMP dynamic router port %d removed

Port mirroring

A switch port can be designated to monitor traffic on any two specified switch ports (port-based) or to monitor traffic to or from any two specified addresses that the switch has learned (address-based).

> *Note:* A probe device, such as the Nortel Networks StackProbe or equivalent, must be connected to the designated monitor port to use this feature. Contact a Nortel Networks sales agent for details about the StackProbe.
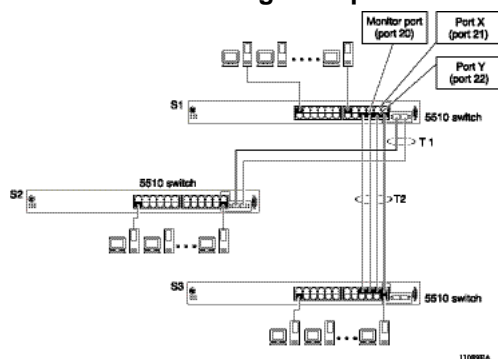
Port-based mirroring configuration

"Port-based mirroring example" (page 26) shows an example of a port-based mirroring configuration in which port 20 is designated as the monitor port for ports 21 and 22 of Switch S1. Although this example shows ports 21 and 22 monitored by the monitor port (port 20), any of the trunk members of T1 and T2 can also be monitored.

In this example, "Port-based mirroring example" (page 26) 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.

Note: Trunks cannot be monitored and trunk members cannot be configured as monitor ports.

Port-based mirroring example

In the configuration example shown in "Port-based mirroring example" (page 26), the designated monitor port (port 20) can be set 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.

- Monitor all traffic transmitted on many ports.

- Monitor all traffic received or transmitted on many ports.

Address-based mirroring configuration

"Address-based mirroring example" (page 27) shows an example of an address-based mirroring configuration in which port 20, the designated monitor port for Switch S1, is monitoring traffic occurring between address A and address B.

Address-based mirroring example
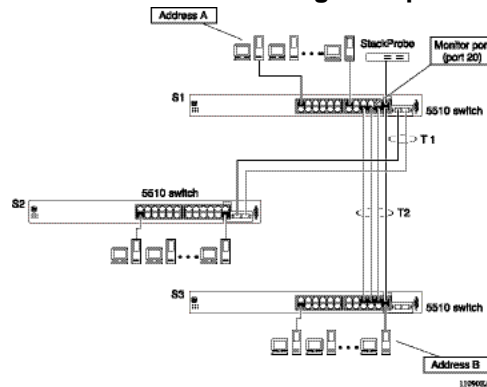
In this configuration, the designated monitor port (port 20) can be set to monitor traffic in any of the following modes:

- 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 limitations
Configuring port mirroring with the CLI

Port mirroring can be configured with the CLI commands detailed in this section.

show port-mirroring command

The `show port-mirroring` command displays the port mirroring configuration.

The syntax for the `show port-mirroring` command is:

`show port-mirroring`

The `show port-mirroring` command is executed in the Privileged EXEC command mode.

port-mirroring command

The `port-mirroring` command sets the port mirroring configuration.

The syntax for the port-mirroring command is:

```
port-mirroring mode {disable | Xrx monitor-port <portlist>
mirror-port-X <portlist> | Xtx monitor-port <portlist>
mirror-port-X <portlist> | ManytoOneRx monitor-port
<portlist> mirror-port-X <portlist> | ManytoOneTx
monitor-port <portlist> mirror-port-X <portlist> |
ManytoOneRxTx 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> | XrxYtxmonitor-port
<portlist> mirror-port-X <portlist> mirror-port-Y <portlist>
| XrxYtxOrYrxXtx monitor-port <portlist> mirror-port-X
<portlist> mirror-port-Y <portlist> | Asrc monitor-port
<portlist> mirror-MAC-A <macaddr> | Adst monitor-port
<portlist> mirror-MAC-A <macaddr> |  AsrcOrAdst monitor-port
<portlist> mirror-MAC-A <macaddr> | AsrcBdst monitor-port
<portlist> mirror-MAC-A <macaddr> mirror-MAC-B <macaddr>
| AsrcBdstOrBsrcAdst monitor-port <portlist> mirror-MAC-A
<macaddr> mirror-MAC-B <macaddr>}
```

The following table outlines the parameters for this command.

port-mirroring parameters

Parameter	Description
disable	Disables port-mirroring.
monitor-port	Specifies the monitor port.
mirror-port-X	Specifies the mirroring port X.
mirror-port-Y	Specifies the mirroring port Y.
mirror-MAC-A	Specifies the mirroring MAC address A.
mirror-MAC-B	Specifies the mirroring MAC address B.
portlist	Enter the port numbers.
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. *Note:* Do not use this mode for mirroring broadcast and multicast traffic.

Parameter	Description
XrxYtxOrXtxYrx	Mirror packets received on port X and transmitted on port Y or packets received on port Y and transmitted on port X. Note: Do not use this mode for mirroring broadcast and multicast traffic.
XrxOrYtx	Mirror packets received on port X or transmitted on port Y.
macaddr	Enter the MAC address in format H.H.H.
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.

The `port-mirroring` command is executed in the Global Configuration command mode.

no port-mirroring command
The `no port-mirroring` command disables port mirroring.

The syntax for the `no port-mirroring` command is:

`no port-mirroring`

The `no port-mirroring` command is executed in the Global Configuration command mode.

Configuring port mirroring with the Web-based Management Interface
Port mirroring can also be configured in the Web-based Management Interface.

To configure port mirroring, follow this procedure:

Step	Action

1   Open the **Port Mirroring** screen by selecting **Applications > Port Mirroring** from the menu. This screen is illustrated below.

Port Mirroring screen

Application > Port Mirroring

Port Mirroring Setting

Monitoring Mode	Disabled	
Monitor Port		
Port X		
Port Y		
Address A		(XX-XX-XX-XX-XX-XX)
Address B		(XX-XX-XX-XX-XX-XX)

Submit

Port Mirroring Active

Monitoring Mode	Disabled

2 In the **Port Mirroring Setting** section, enter the new port mirroring settings. The following table outlines the fields in this section.

Port Mirroring Setting fields

Field	Description
Monitoring Mode	Choose any one of the six port-based monitoring modes or any one of the five address-based monitoring modes. The following options are available: • Disabled • -> Port X • Port X -> • <-> Port X • -> Port X or Port Y -> • -> Port X and Port Y -> • <-> Port X and Port Y <-> • Address A -> any Address • any Address -> Address A • <-> Address A • Address A -> Address B • Address A <-> Address B The default value is Disabled.

Field	Description
Monitor Port	Select the port that will act as the monitoring port.
Port X	In port-based configurations, choose the first switch port to be monitored by the designated monitor port. This port is monitored according to the value "X" in the Monitoring Mode field.
Port Y	In port-based configurations, choose the second switch port to be monitored by the designated monitor port. This port is monitored according to the value "Y" in the Monitoring Mode field.
Address A	In address-based configurations, type the MAC address to monitor by the designated monitor port. This address is monitored according to the value "Address A" in the Monitoring Mode field.
Address B	In address-based configurations, type the MAC address to monitor by the designated monitor port. This address is monitored according to the value "Address B" in the Monitoring Mode field.

3   Click **Submit**.

The new mirroring configuration is displayed in **Port Mirroring Active** section.

—End—

Showing unit statistics

The JDM allows you to see the statistics for each unit in a stack.

To see the statistics for a unit, open the Unit Stats tab by selecting **Edit > Unit > Unit Stats**.

Unit Stats screen

Field	Description
UnitNoResourcesPktsDropped	The number of packets dropped on stack up/down ports on this unit due to a lack of resources.

Graphing switch chassis data

The JDM provides the ability to view switch statistical information in a variety of graphs.

To make use of these capabilities, open the **Graph Chassis** screen by selecting **Graph > Chassis** from the menu. This screen is illustrated below.

Graph Chassis screen

The following sections describe the informational tabs on this screen and the type of data each represents. Refer to "Creating a graph" (page 69) for the procedure to graph this data.

SNMP tab

The **SNMP** tab provides read-only statistical information about SNMP traffic.

To view the **SNMP** tab, follow this procedure:

Step	Action
1	Open the **Graph Chassis** screen by selecting **Graph > Chassis** from the menu. The **Graph Chassis** screen opens with the **SNMP** tab selected. This screen is illustrated below.

Graph Chassis dialog - SNMP tab

192.168.249.46 - Graph Chassis						
SNMP	IP	ICMP In	ICMP Out	TCP	UDP	

	AbsoluteValue	Cumulative	Average/sec	Minimum/sec	Maximum/sec	LastVal/sec
InPkts	26,341	7	0.35	0.1	0.6	0.1
OutPkts	26,340	7	0.35	0.1	0.6	0.1
InTotalReqVars	534,275	148	7.4	2.3	12.5	2.3
InTotalSetVars	6	0	0	0	0	0
InGetRequests	15,789	5	0.25	0.1	0.4	0.1
InGetNexts	74	0	0	0	0	0
InSetRequests	3	0	0	0	0	0
InGetResponses	0	0	0	0	0	0
OutTraps	0	0	0	0	0	0
OutTooBigs	0	0	0	0	0	0
OutNoSuchNames	0	0	0	0	0	0
OutBadValues	0	0	0	0	0	0
OutGenErrs	0	0	0	0	0	0
InBadVersions	0	0	0	0	0	0
InBadCommunityNames	0	0	0	0	0	0
InBadCommunityUses	0	0	0	0	0	0
InASNParseErrs	0	0	0	0	0	0
InTooBigs	0	0	0	0	0	0
InNoSuchNames	0	0	0	0	0	0
InBadValues	0	0	0	0	0	0
InReadOnlys	0	0	0	0	0	0
InGenErrs	0	0	0	0	0	0

Clear Counters | Close | Help... | Poll Interval: 10s | 0 day, 00h:00m:20s

2 The following table describes the fields on this tab.

SNMP tab fields

Field	Description
InPkts	The total number of messages delivered to the SNMP from the transport service.
OutPkts	The total number of SNMP messages passed from the SNMP protocol to the transport service.
InTotalReqVars	The total number of MIB objects retrieved successfully by the SNMP protocol as the result of receiving valid SNMP Get-Request and Get-Next PDUs.
InTotalSetVars	The total number of MIB objects altered successfully by the SNMP protocol as the result of receiving valid SNMP Set-Request PDUs.
InGetRequests	The total number of SNMP Get-Request PDUs that have been accepted and processed by the SNMP protocol.
InGetNexts	The total number of SNMP Get-Next PDUs accepted and processed by the SNMP protocol.

InSetRequests	The total number of SNMP Set-Request PDUs accepted and processed by the SNMP protocol.
InGetResponses	The total number of SNMP Get-Response PDUs accepted and processed by the SNMP protocol.
OutTraps	The total number of SNMP Trap PDUs generated by the SNMP protocol.
OutTooBigs	The total number of SNMP PDUs generated by the SNMP protocol for which the value of the error-status field is tooBig.
OutNoSuchNames	The total number of SNMP PDUs generated by the SNMP protocol for which the value of the error-status field is noSuchName.
OutBadValues	The total number of SNMP PDUs generated by the SNMP protocol for which the value of the error-status field is badValue.
OutGenErrs	The total number of SNMP PDUs generated by the SNMP protocol for which the value of the error-status field is genErr.
InBadVersions	The total number of SNMP messages delivered to the SNMP protocol for an unsupported SNMP version.
InBadCommunityNames	The total number of SNMP messages delivered to the SNMP protocol that used an unknown SNMP community name.
InBadCommunityUses	The total number of SNMP messages delivered to the SNMP protocol that represented an SNMP operation not allowed by the SNMP community named in the message.
InASNParseErrs	The total number of ASN.1 or BER errors encountered by the SNMP protocol when decoding received SNMP messages.
InTooBigs	The total number of SNMP PDUs delivered to the SNMP protocol for which the value of the error-status field is tooBig.
InNoSuchNames	The total number of SNMP PDUs delivered to the SNMP protocol for which the value of the error-status field is noSuchName.
InBadValues	The total number of SNMP PDUs delivered to the SNMP protocol for which the value of the error-status field is badValue.

InReadOnlys	The total number of SNMP PDUs delivered to the SNMP protocol for which the value of the error-status field is readOnly. It is a protocol error to generate an SNMP PDU containing the value "readOnly" in the error-status field. This object is provided to detect incorrect implementations of the SNMP.
InGenErrs	The total number of SNMP PDUs delivered to the SNMP protocol for which the value of the error-status field is genErr.

—End—

IP tab

The **IP** tab shows read-only information about the IP packets that have interfaced with the switch

To view the **IP** tab, follow this procedure:

Step	Action

1   Open the **Graph Chassis** screen by selecting **Graph > Chassis** from the menu. The **Graph Chassis** screen opens.

2   Select the **IP** tab. This tab is demonstrated in the following illustration.

Graph Chassis screen - IP tab

	AbsoluteValue	Cumulative	Average/sec	Minimum/sec	Maximum/sec	LastVal/sec
InReceives	35,335	7	0.35	0.1	0.6	0.1
InHdrErrors	0	0	0	0	0	0
InAddrErrors	0	0	0	0	0	0
ForwDatagrams	0	0	0	0	0	0
InUnknownProtos	0	0	0	0	0	0
InDiscards	0	0	0	0	0	0
InDelivers	35,335	7	0.35	0.1	0.6	0.1
OutRequests	35,477	7	0.35	0.1	0.6	0.1
OutDiscards	0	0	0	0	0	0
OutNoRoutes	0	0	0	0	0	0
FragOKs	5,246	1	0.05	0	0.1	0
FragFails	0	0	0	0	0	0
FragCreates	5,246	1	0.05	0	0.1	0
ReasmReqds	0	0	0	0	0	0
ReasmOKs	0	0	0	0	0	0
ReasmFails	0	0	0	0	0	0

Clear Counters   Close   Help...   Poll Interval: 10s

3   The following table outlines the fields on this tab.

IP tab fields

Field	Description
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 that were Source-Routed by way of this address and had successful Source-Route option processing.
InUnknownProtos	The number of locally addressed datagrams received successfully but discarded because of an unknown or unsupported protocol.
InDiscards	The number of input IP datagrams for which no problems were encountered to prevent their continued processing but that were discarded (for example, for lack of buffer space). Note that 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. Note that this counter does not include any datagrams counted in ipForwDatagrams.

Field	Description
OutDiscards	The number of output IP datagrams for which no problem was encountered to prevent their transmission to their destination, but that were discarded (for example, for lack of buffer space). Note that this counter would include datagrams counted in ipForwDatagrams if any such packets met this (discretionary) discard criterion.
OutNoRoutes	The number of IP datagrams discarded because no route could be found to transmit them to their destination. Note that this counter includes any packets counted in ipForwDatagrams that meet this no-route criterion. Note that this includes any datagrams a host cannot route because all of its default gateways are down.
FragOKs	The number of IP datagrams that have been successfully fragmented at this entity.
FragFails	The number of IP datagrams that have been discarded because they needed to be fragmented at this entity but could not be, for example, because their Don't Fragment flag was set.
FragCreates	The number of IP datagram fragments that have been generated as a result of fragmentation at this entity.
ReasmReqds	The number of IP fragments received that needed to be reassembled at this entity.
ReasmOKs	The number of IP datagrams successfully reassembled.
ReasmFails	The number of failures detected by the IP reassembly algorithm (for whatever reason: timed out, errors, and so on). Note that this is not necessarily a count of discarded IP fragments because some algorithms (notably the algorithm in RFC 815) can lose track of the number of fragments by combining them as they are received.

—End—

ICMP In tab

The **ICMP In** tab provides read-only information about inbound ICMP messages.

To view the **ICMP In** tab, follow this procedure:

Step	Action
1	Open the **Graph Chassis** screen by selecting **Graph > Chassis** from the menu. The **Graph Chassis** screen opens.
2	Select the **ICMP In** tab. This tab is illustrated below.

Graph Chassis screen - ICMP In tab

Step	Action
3	The following table describes the fields on this tab.

ICMP In tab fields

Field	Description
SrcQuenchs	The number of ICMP Source Quench messages received.
Redirects	The number of ICMP Redirect messages received.
Echos	The number of ICMP Echo (request) messages received.
EchoReps	The number of ICMP Echo Reply messages received.
Timestamps	The number of ICMP Timestamp (request) messages received.
TimestampReps	The number of ICMP Timestamp Reply messages received.
AddrMasks	The number of ICMP Address Mask Request messages received.
AddrMaskReps	The number of ICMP Address Mask Reply messages received.
ParmProbs	The number of ICMP Parameter Problem messages received.

Field	Description
DestUnreachs	The number of ICMP Destination Unreachable messages received.
TimeExcds	The number of ICMP Time Exceeded messages received.

—End—

ICMP Out tab

The **ICMP Out** tab provides read-only information about outbound ICMP messages.

To view the **ICMP Out** tab, follow this procedure:

Step	Action
1	Open the **Graph Chassis** screen by selecting **Graph > Chassis** from the menu. The **Graph Chassis** screen will open.
2	Select the **ICMP Out** tab. This tab is illustrated below.

Graph Chassis screen - ICMP Out tab

Step	Action
3	The following table describes the fields on this tab.

ICMP Out tab fields

Field	Description
SrcQuenchs	The number of ICMP Source Quench messages sent.
Redirects	The number of ICMP Redirect messages received. For a host, this object will always be zero, because hosts do not send redirects.

Field	Description
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.
DestUnreachs	The number of ICMP Destination Unreachable messages sent.
TimeExcds	The number of ICMP Time Exceeded messages sent.

—End—

TCP tab

The **TCP** tab provides read-only information about TCP activity on the switch.

To view the **TCP** tab, follow this procedure:

Step	Action

1 Open the **Graph Chassis** screen by selecting **Graph > Chassis** from the menu. The **Graph Chassis** screen will open.

2 Select the **TCP** tab. This tab is illustrated below.

Graph Chassis screen - TCP tab

3   The following table describes the fields on this tab.

TCP tab fields

Field	Description
ActiveOpens	The number of times TCP connections have made a direct transition to the SYN-SENT state from the CLOSED state.
PassiveOpens	The number of times TCP connections have made a direct transition to the SYN-RCVD state from the LISTEN state.
AttemptFails	The number of times TCP connections have made a direct transition to the CLOSED state from either the SYN-SENT state or the SYN-RCVD state, plus the number of times TCP connections have made a direct transition to the LISTEN state from the SYN-RCVD state.
EstabResets	The number of times TCP connections have made a direct transition to the CLOSED state from either the ESTABLISHED state or the CLOSE-WAIT state.
CurrEstab	The number of TCP connections for which the current state is either ESTABLISHED or CLOSE-WAIT.
InSegs	The 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.
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.

—End—

UDP tab

The **UDP** tab provides read-only information about UDP activity on the switch.

To view the **UDP** tab, follow this procedure:

Step	Action

1   Open the **Graph Chassis** screen by selecting **Graph > Chassis** from the menu. The **Graph Chassis** screen opens.

2   Select the **UDP** tab. This tab is illustrated below.

Graph Chassis screen - UDP tab

3   The following table describes the fields on this tab.

UDP tab fields

Field	Description
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 could not 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.

—End—

OSPF tab

The **OSPF** tab is used to display statistical information about OSPF operation on the switch.

To view the **OSPF** tab, use the following procedure:

Step	Action

1   Open the **Graph Chassis** screen by selecting **Graph > Chassis** from the menu. The **Graph Chassis** screen opens.

2   Select the **OSPF** tab. This tab is illustrated below.

Graph Chassis screen - OSPF tab

3   Use the provided fields to view the OSPF statistics. These fields are outlined in the following table.

OSPF tab fields

Field	Description
LsdbTblSize	Indicates the number of entries in the link state database.
TxPackets	Indicates the number of packets transmitted by OSPF.
RxPackets	Indicates the number of packets received by OSPF.
RxBadPackets	Indicates the number of bad packets received by OSPF.
SpfRuns	Indicates the total number of SPF calculations performed by OSPF.
BuffersAllocated	Indicates the total number of buffers allocated by OSPF.
BuffersFreed	Indicates the total number of buffers freed by OSPF.
BufferAllocFailures	Indicates the number of times that OSPF has failed to allocate buffers.
BufferFreeFailures	Indicates the number of times that OSPF has failed to free buffers.

—End—

VRRP tab

The **VRRP** tab is used to display statistical information about VRRP operation on the switch.

To view the **VRRP** tab, use the following procedure:

Step	Action

1   Open the **Graph Chassis** screen by selecting **Graph > Chassis** from the menu. The **Graph Chassis** screen opens.

2   Select the **VRRP** tab. This tab is illustrated below.

Graph Chassis screen - VRRP tab

3   Use the provided fields to view the VRRP statistics. These fields are outlined in the following table.

Graph Chassis screen - VRRP tab

Field	Description
RouterChecksumErrors	The total number of VRRP packets received with an invalid VRRP checksum value.
RouterVersionErrors	The total number of VRRP packets received with an unknown or unsupported version number.
RouterVrIdErrors	The total number of VRRP packets received with an invalid VRID for this virtual router."

—End—

Graphing switch port data

The Java Device Manager (JDM) provides the ability to view port statistical information in a variety of graphs.

To make use of these capabilities, open the **Graph Port** screen by selecting one or multiple ports on the **Device View** and then selecting **Graph > Port** from the menu. This screen is illustrated below.

Graph Port screen

The following sections describe the informational tabs on this screen and
what type of data each represents. Refer to "Creating a graph" (page 69) for
the procedure to graph this data.

Note: Some statistics are only available when a single port is graphed.

Interface tab

The **Interface** tab displays read-only information about the selected
interfaces.

To view the **Interface** tab, follow this procedure:

Step	Action
1	Open the **Graph Port** screen by selecting one or multiple ports on the **Device View** and then selecting **Graph > Port** from the menu.
2	Select the **Interface** tab. This tab is illustrated below.

Graph Port screen - Interface tab

3 The following table describes the fields on this tab.

Interface tab fields

Field	Description
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 were not addressed to a multicast or broadcast address at this sublayer.
OutUcastPkts	The number of packets that higher-level protocols requested be transmitted that were not addressed to a multicast address at this sublayer. This total number includes those packets discarded or unsent.
InNUcastPkts	The number of packets delivered by this sublayer to a higher (sub)layer, which were addressed to a multicast or broadcast address at this sublayer.
OutNUcastPkts	The total number of packets that higher-level protocols requested be transmitted, and which were addressed to a multicast or broadcast address at this sublayer, including those that were discarded or not sent.

Field	Description
InDiscards	The number of inbound packets that were chosen to be discarded even though no errors had been detected to prevent their being deliverable to a higher-layer protocol. One possible reason for discarding such a packet could be to free up buffer space.
OutDiscards	The number of outbound packets which were chosen to be discarded even though no errors had been detected to prevent their being transmitted. One possible reason for discarding such a packet could 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 could not be transmitted because of errors. For character-oriented or fixed-length interfaces, the number of outbound transmission units that could not be transmitted because of errors.
InUnknownProtos	For packet-oriented interfaces, the number of packets received via the interface that were 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 via the interface that were discarded because of an unknown or unsupported protocol. For any interface that does not support protocol multiplexing, this counter will always be 0.

—End—

Ethernet Errors tab

The **Ethernet Errors** tab displays read-only information about port Ethernet error statistics.

To view the **Ethernet Errors** tab, follow this procedure:

Step	Action
1	Open the **Graph Port** screen by selecting one or multiple ports on the **Device View** and then selecting **Graph > Port** from the menu. The **Graph Port** screen opens.
2	Select the **Ethernet Errors** tab. This tab is illustrated below.

Graph Port screen - Ethernet Errors tab

Step	Action
3	The following table describes the fields on this tab.

Ethernet Errors tab fields

Field	Description
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.

Field	Description
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 frameCheckError status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions occur are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.
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. The precise meaning of the count represented by an instance of this object is implementation specific. In particular, an instance of this object may 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 per transmission attempt, even if the carrier sense condition fluctuates during a transmission attempt.

Field	Description
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 frameTooLong status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions occur are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.
SQETestErrors	A count of times that the SQE TEST ERROR message is generated by the PLS sublayer for a particular interface. The SQE TEST ERROR message is defined in section 7.2.2.2.4 of ANSI/IEEE 802.3-1985 and its generation is described in section 7.2.4.6 of the same document.
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.

Field	Description
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.

—End—

Bridge tab

The **Bridge** tab displays read-only information about port frame statistics.

To view the **Bridge** tab, follow this procedure:

Step	Action

1   Open the **Graph Port** screen by selecting one or multiple ports on the **Device View** and then selecting **Graph > Port** from the menu. The **Graph Port** screen will open.

2   Select the **Bridge** tab. This tab is illustrated below.

Graph Port screen - Bridge tab

3   The following table describes the fields on this tab.

Bridge tab fields

Field	Description
DelayExceededDiscards	Number of frames discarded by the port due to excessive transit delays through the bridge. It is incremented by both transparent and source route bridges.
MtuExceededDiscards	Number of frames discarded by the port due to an excessive size. It is incremented by both transparent and source route bridges.
InFrames	The number of frames that have been received by this port from its segment.
OutFrames	The number of frames that have been received by this port from its segment.
InDiscards	Count of valid frames received which were discarded (filtered) by the Forwarding Process.

—End—

Rmon tab

The **Rmon** tab displays read-only remote monitoring statistics.

To view the **Rmon** tab, follow this procedure:

Step	Action
1	Open the **Graph Port** screen by selecting one or multiple ports on the **Device View** and then selecting **Graph > Port** from the menu. The **Graph Port** screen opens.
2	Select the **Rmon** tab. This tab is illustrated below.

Graph Port screen - Rmon tab

3   The following table describes the fields on this tab.

RMON tab fields

Field	Description
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 were directed to the broadcast address. Note that this does not include multicast packets.
MulticastPkts	The total number of good packets received that were directed to a multicast address. Note that this number does not include packets directed to the broadcast address.
CRCAlignErrors	The total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, but had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).

Field	Description
UndersizePkts	The total number of packets received that were less than 64 octets long (excluding framing bits but including FCS octets) and were otherwise well formed.
OversizePkts	The total number of packets received that were longer than 1518 octets (excluding framing bits but including FCS octets) and were otherwise well formed.
Fragments	The total number of packets received that were less than 64 octets in length (excluding framing bits but including FCS octets) and had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). It is entirely normal for etherStatsFragments to increment 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 were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a 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 were between 1 and 64 octets in length (excluding framing bits but including FCS octets).
65..127	The total number of packets (including bad packets) received that were between 65 and 127 octets in length (excluding framing bits but including FCS octets).
128..255	The total number of packets (including bad packets) received that were between 128 and 255 octets in length (excluding framing bits but including FCS octets).
256..511	The total number of packets (including bad packets) received that were between 256 and 511 octets in length (excluding framing bits but including FCS octets).

Field	Description
511..1023	The total number of packets (including bad packets) received that were between 511 and 1023 octets in length (excluding framing bits but including FCS octets).
1024..1518	The total number of packets (including bad packets) received that were between 1024 and 1518 octets in length (excluding framing bits but including FCS octets).

—End—

EAPOL Stats tab

The **EAPOL Stats** tab displays read-only EAPOL statistics.

To open the **EAPOL Stats** tab, follow this procedure:

Step	Action

1   Open the **Graph Port** screen by selecting one or multiple ports on the **Device View** and then selecting **Graph > Port** from the menu. The **Graph Port** screen opens.

2   Select the **EAPOL Stats** tab. This tab is illustrated below.

Graph Port screen - EAPOL Stats tab

3   The following table describes the fields on this tab.

EAPOL Stats tab fields

Field	Description
EapolFramesRx	The number of valid EAPOL frames of any type that have been received by this authenticator.
EapolFramesTx	The number of EAPOL frame types of any type that have been transmitted by this authenticator.
EapolStartFramesRx	The number of EAPOL start frames that have been received by this authenticator.
EapolLogoffFramesRx	The number of EAPOL Logoff frames that have been received by this authenticator.
EapolRespIdFramesRx	The number of EAPOL Resp/Id frames that have been received by this authenticator.
EapolRespFramesRx	The number of valid EAP Response frames (other than Resp/Id frames) that have been received by this authenticator.
EapolReqIdFramesTx	The number of EAPOL Req/Id frames that have been transmitted by this authenticator.
EapolReqFramesTx	The number of EAP Req/Id frames (Other than Rq/Id frames) that have been transmitted by this authenticator.
InvalidEapolFramesRx	The number of EAPOL frames that have been received by this authenticator in which the frame type is not recognized.
EapLengthErrorFramesRx	The number of EAPOL frames that have been received by this authenticator in which the packet body length field is not valid.

—End—

EAPOL Diag tab

The **EAPOL Diag** tab displays read-only EAPOL diagnostic statistics.

To open the **EAPOL Diag** tab, follow this procedure:

Step	Action
1	Open the **Graph Port** screen by selecting one or multiple ports on the **Device View** and then selecting **Graph > Port** from the menu. The **Graph Port** screen opens.

2   Select the **EAPOL Diag** tab. This tab is illustrated below.

Graph Port screen - EAPOL Diag tab

3   The following table describes the fields on this tab.

EAPOL Diag tab fields

Field	Description
EntersConnecting	Counts the number of times that the Authenticator PAE state machine transitions to the Connecting state from any other state.
EapLogoffsWhileConnecting	Counts the number of times that the Authenticator PAE state machine transitions from Connected to Disconnected as a result of receiving an EAPOL-Logoff message.
EntersAuthenticating	Counts the number of times that the Authenticator PAE state machine transitions from Connecting to Authenticating as a result of receiving an EAP-Response/Identity message from the supplicant.
AuthSuccessWhileAuthenticating	Counts the number of times that the Authenticator PAE state machine transitions from Authenticating to Authenticated as a result of the Backend authentication state machine indicating successful authentication of the supplicant.

Field	Description
AuthTimeoutsWhile Authenticating	Counts the number of times that the Authenticator PAE state machine transitions from Authenticating to Aborting as a result of the Backend authentication state machine indicating authentication timeout.
AuthFailWhileAuthenticating	Counts the number of times that the Authenticator PAE state machine transitions from Authenticating to Held as a result of the Backend authentication state machine indicating authentication failure.
AuthReauthsWhileAuthenticating	Counts the number of times that the Authenticator PAE state machine transitions from Authenticating to Aborting as a result of a reauthentication request.
AuthEapStartsWhileAuthenticating	Counts the number of times that the Authenticator PAE state machine transitions from Authenticating to Aborting as a result of an EAPOL-Start message being received from the supplicant.
AuthEapLogoffWhileAuthenticating	Counts the number of times that the Authenticator PAE state machine transitions from Authenticating to Aborting as a result of an EAPOL-Logoff message being received from the supplicant.
AuthReauthsWhileAuthenticated	Counts the number of times that the Authenticator PAE state machine transitions from Authenticated to Connecting as a result of a reauthentication request.
AuthEapStartsWhileAuthenticated	Counts the number of times that the Authenticator PAE state machine transitions from Authenticated to Connecting as a result of an EAPOL-Start message being received from the supplicant.
AuthEapLogoffWhileAuthenticated	Counts the number of times that the Authenticator PAE state machine transitions from Authenticated to Disconnected as a result of an EAPOL-Logoff message being received from the supplicant.

Field	Description
BackendResponses	Counts the number of times that the Backend Authentication state machine sends an Initial-Access request packet to the Authentication server.
BackendAccessChallenges	Counts the number of times that the Backend Authentication state machine receives an Initial-Access challenge packet from the Authentication server.
BackendOtherRequestsToSupplicants	Counts the number of times that the Backend Authentication state machine sends an EAP request packet (other than an Identity, Notification, failure, or success message) to the supplicant.
BackendNonNakResponsesFromSupplicant	Counts the number of times that the Backend Authentication state machine receives a response from the supplicant to an initial EAP request and the response is something other than EAP-NAK.
BackendAuthSuccesses	Counts the number of times that the Backend Authentication state machine receives an EAP-success message from the Authentication server.
BackendAuthFails	Counts the number of times that the Backend Authentication state machine receives an EAP-failure message from the Authentication server.

—End—

LACP tab

The LACP tab displays read-only Link Aggregation Control Protocol (LACP) diagnostic statistics.

To view the LACP tab, follow this procedure:

Step	Action

1   Open the **Graph Port** screen by selecting one or multiple ports on the **Device View** and then selecting **Graph > Port** from the menu. The **Graph Port** screen opens.

2   Select the **LACP** tab. This tab is illustrated below.

Graph Port screen - LACP tab

Note: The Marker Protocol Generator/Receiver is currently not a supported feature.

3 The following table describes the fields on this tab.

LACP tab fields

Field	Description
LACPDUsRX	Denotes the number of valid LACPDUs received on this Aggregation Port. This value is read-only.
MarkerPDUsRX	Signifies the number of valid Marker PDUs received on this Aggregation Port. This value is read-only.
MarkerResponsePDUsRX	The number of valid Marker Response PDUs received on this Aggregation Port. This value is read-only.
UnknownRX	Indicates the number of frames received that can • Carry the Slow Protocols Ethernet Type value (43B.4), but contain an unknown PDU. • Are addressed to the Slow Protocols group MAC Address (43B.3), but do not carry the Slow Protocols Ethernet Type. This value is read-only.
IllegalRX	Denotes the number of frames received that carry the Slow Protocols Ethernet Type value (43B.4), but contain a badly formed PDU or an illegal value of Protocol Subtype (43B.4). This value is read-only.
LACPDUsTX	Signifies the number of LACPDUs that are transmitted on this Aggregation Port. This value is read-only.

Field	Description
MarkerPDUsTX	Displays the number of Marker PDUs transmitted on this Aggregation Port. This value is read-only.
MarkerResponsePDUsTX	Indicates the number of Marker Response PDUs that are transmitted on this Aggregation Port. This value is read-only.

—End—

Misc tab

The **Misc** tab is used to display statistical information that does not belong grouped with the other tabs. To view the **Misc** tab, follow this procedure:

Step	Action

1 Open the **Graph Port** screen by selecting one or multiple ports on the **Device View** and then selecting **Graph > Port** from the menu. The **Graph Port** screen opens.

2 Select the **Misc** tab. This tab is illustrated below.

Graph Port dialog - Misc tab

3 Using the fields provided, view the statistical information. These fields are outlined in the following table.

Misc tab fields

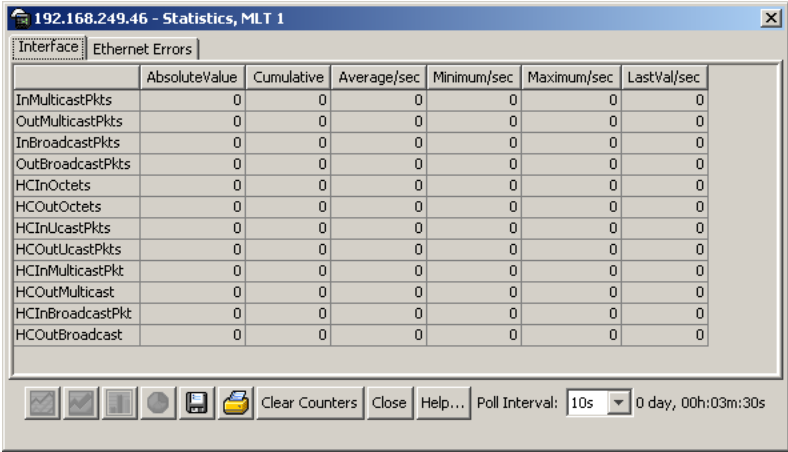| Field	Description
NoResourcesPktsDropped	The number of packets dropped due to a lack of resources.

—End—

Graphing multilink trunk statistics

The Java Device Manager (JDM) provides the ability to view Multilink Trunk (MLT) statistical information in a variety of graphs.

To make use of these capabilities, open the **MLT_LACP** screen by selecting **VLAN > MLT/LACP** from the menu. This screen opens with the **Multilink Trunks** tab selected. On this tab, select the row that represents the **MLT** to graph and click the **Graph** button. The **MLT Statistics** screen opens. This screen is illustrated in .

MLT Statistics screen

The following sections will describe the informational tabs on this screen and the type of data each represents. Refer to for the procedure to graph this data.

Interface tab

The **Interface** tab provides read-only statistical information about the selected Multilink Trunk.

To view the **Interface** tab, follow this procedure:

Step	Action
1	Open the **MLT_LACP** screen by selecting **VLAN > MLT/LACP** from the menu. This screen opens with the **Multilink Trunks** tab selected.
2	On this tab, select the row that represents the MLT to graph and click the **Graph** button. The **MLT Statistics** screen opens with the **Interface** tab selected. This screen and tab are illustrated below.

MLT Statistics screen - Interface tab

192.168.249.46 - Statistics, MLT 1

Interface | Ethernet Errors

	AbsoluteValue	Cumulative	Average/sec	Minimum/sec	Maximum/sec	LastVal/sec
InMulticastPkts	0	0	0	0	0	0
OutMulticastPkts	0	0	0	0	0	0
InBroadcastPkts	0	0	0	0	0	0
OutBroadcastPkts	0	0	0	0	0	0
HCInOctets	0	0	0	0	0	0
HCOutOctets	0	0	0	0	0	0
HCInUcastPkts	0	0	0	0	0	0
HCOutUcastPkts	0	0	0	0	0	0
HCInMulticastPkt	0	0	0	0	0	0
HCOutMulticast	0	0	0	0	0	0
HCInBroadcastPkt	0	0	0	0	0	0
HCOutBroadcast	0	0	0	0	0	0

Clear Counters | Close | Help... | Poll Interval: 10s | 0 day, 00h:03m:30s

The following table describes the fields on this tab.

Interface tab fields

Field	Description
InMulticastPkt	The number of packets delivered to this MLT that were addressed to a multicast address at this sublayer. For a MAC layer protocol, this number includes both Group and Functional addresses.
OutMulticast	The total number of packets that higher-level protocols requested be transmitted, and that were addressed to a multicast address at this MLT, including those that were discarded or not sent. For a MAC layer protocol, this number includes both Group and Functional addresses.
InBroadcastPkt	The number of packets delivered to this MLT that were addressed to a broadcast address at this sublayer.
OutBroadcast	The total number of packets that higher-level protocols requested be transmitted, and that were addressed to a broadcast address at this MLT, including those that were discarded or not sent.
HCInOctets	The total number of octets received on the MLT interface, including framing characters.
HCOutOctets	The total number of octets transmitted out of the MLT interface, including framing characters.
HCInUcastPkts	The number of packets delivered by this MLT to higher level protocols that were not addressed to a multicast or broadcast address at this sublayer.

Field	Description
HCOutUcastPkts	The number of packets that higher-level protocols requested be transmitted that were not addressed to a multicast address at this MLT. This total number includes those packets discarded or unsent.
HCInMulticastPkt	The number of packets delivered to this MLT that were addressed to a multicast address at this sublayer. For a MAC layer protocol, this number includes both Group and Functional addresses.
HCOutMulticast	The total number of packets that higher-level protocols requested be transmitted, and that were addressed to a multicast address at this MLT, including those that were discarded or not sent. For a MAC layer protocol, this number includes both Group and Functional addresses.
HCInBroadcastPkt	The number of packets delivered to this MLT that were addressed to a broadcast address at this sublayer.
HCOutBroadcast	The total number of packets that higher-level protocols requested be transmitted, and that were addressed to a broadcast address at this MLT, including those that were discarded or not sent.

—End—

Ethernet Errors tab

The **Ethernet Errors** tab provides read-only statistical information about Ethernet errors that have occurred on the selected Multilink Trunk.

To view the **Ethernet Errors** tab, follow this procedure:

Step	Action
1	Open the **MLT_LACP** screen by selecting **VLAN > MLT/LACP** from the menu. This screen will open with the **Multilink Trunks** tab selected.
2	On this tab, select the row that represents the MLT to graph and click the **Graph** button. The **MLT Statistics** screen will open.
3	Select the **Ethernet Errors** tab. This tab is illustrated below.

MLT Statistics screen - Ethernet Errors tab

The following table describes the fields on this tab.

Ethernet Errors tab fields

Field	Description
AlignmentErrors	A count of frames received on a particular MLT 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 an MLT 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 frameCheckError status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions occur are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.

Field	Description
IMacTransmitError	A count of frames for which transmission on a particular MLT fails due to an internal MAC sublayer transmit error. A frame is only counted by an instance of this object if it is not counted by the corresponding instance of either the LateCollisions object, the ExcessiveCollisions object, or the CarrierSenseErrors object.
IMacReceiveError	A count of frames for which reception on a particular MLT fails due to an internal MAC sublayer receive error. A frame is only counted by an instance of this object if it is not counted by the corresponding instance of either the FrameTooLongs object, the AlignmentErrors object, or the FCSErrors object. The precise meaning of the count represented by an instance of this object is implementation specific. In particular, an instance of this object may 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 MLT. The count represented by an instance of this object is incremented at most once per transmission attempt, even if the carrier sense condition fluctuates during a transmission attempt.
FrameTooLong	A count of frames received on a particular MLT that exceed the maximum permitted frame size. The count represented by an instance of this object is incremented when the frameTooLong status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions occur are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.

Field	Description
SQETestError	A count of times that the SQE TEST ERROR message is generated by the PLS sublayer for a particular MLT. The SQE TEST ERROR message is defined in section 7.2.2.2.4 of ANSI/IEEE 802.3-1985 and its generation is described in section 7.2.4.6 of the same document.
DeferredTransmiss	A count of frames for which the first transmission attempt on a particular MLT is delayed because the medium is busy. The count represented by an instance of this object does not include frames involved in collisions.
SingleCollFrames	A count of successfully transmitted frames on a particular MLT for which transmission is inhibited by exactly one collision. A frame that is counted by an instance of this object is also counted by the corresponding instance of either the ifOutUcastPkts, ifOutMulticastPkts, or ifOutBroadcastPkts, and is not counted by the corresponding instance of the MultipleCollisionFrames object.
MultipleCollFrames	A count of successfully transmitted frames on a particular MLT for which transmission is inhibited by more than one collision. A frame that is counted by an instance of this object is also counted by the corresponding instance of either the ifOutUcastPkts, 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 MLT 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.
ExcessiveColls	A count of frames for which transmission on a particular MLT fails due to excessive collisions.

—End—

Graphing VLAN DHCP statistics

To create a graph of VLAN DHCP configuration, follow this procedure:

Step   Action

---

1   Open the VLANs screen by selecting **VLAN > VLANs** from the menu.

2   Select the desired VLAN.

3   Click **IP**. The **IP VLAN** screen opens with the **IP Address** tab selected.

4   Click the **DHCP** tab.

5   Click **Graph**. The **DHCP Stats** screen opens. This screen is illustrated in below.

DHCP Stats screen

6   Highlight the required data and click on the type of graph to produce. For a detailed explanation of graph creation, refer to .

The following table explains the fields found on this screen.

DHCP Stats screen fields

Field	Description
NumRequests	The number of DHCP requests handled.
NumReplies	The number of DHCP replies handled.

—End—

Creating a graph

Several screens in the Java Device Manager (JDM) provide a means to view and make use of statistical information gathered by the switch. To turn this statistical information in either a bar, line, area, or pie graph, follow this procedure:

Step	Action

1   After opening a screen that provides graphing capabilities and selecting the desired tab, select the information to graph in one of the following ways:

    a. Click and drag the mouse across the rows and columns of data to graph.

    b. Hold the **Control (CTRL)** key and click on the cells of data to graph.

    c. Hold the **Shift** key and click a range of data to graph.

2   Press the graph button that corresponds to the type of graph to be created. These graph buttons are illustrated below.

Graph buttons
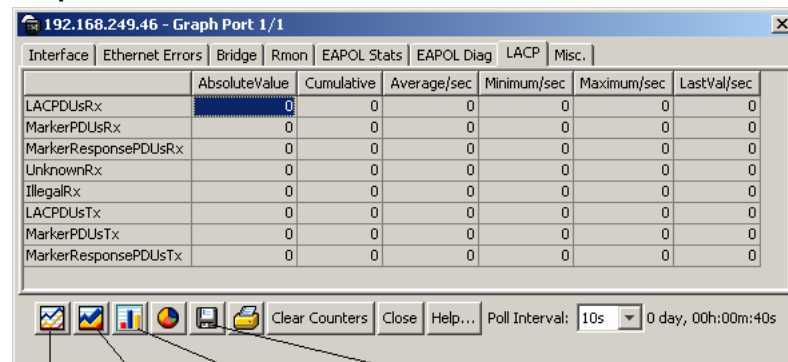
—**End**—

Chapter 2
System Diagnostics and Statistics

This chapter outlines the system diagnostic and statistical information and tools available in the Nortel Ethernet Routing Switch 5500 Series.

Use commands in this chapter to get statistics and diagnose problems. Some are available in both the CLI and in the Web-based Management interface.

Diagnostic information in the CLI

The CLI commands detailed in this section are used to display diagnostic and statistical information from the switch.

Port statistics (CLI)

Use the CLI commands in this section to gather port statistics from the switch.

show port-statistics command

The `show port-statistics` command displays the statistics for the port on both received and transmitted traffic, including the 10 GByte ports (ports 25 and 26) on the 5530.

The syntax for the `show port-statistics` command is:

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

The `show port-statistics` command is executed in the Privileged EXEC command mode.

"show port-statistics parameters" (page 72) outlines the parameters for this command.

show port-statistics parameters

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

show port-statistics examples
`show port-statistics 12`

clear-stats command

The `clear-stats` command clears all statistical information for the specified port. All counters are set to zero (0).

The syntax for the `clear-stats` command is:

`clear-stats [port <portlist>]`

"clear-stats parameters" (page 72) outlines the parameters for this command.

clear-stats parameters

Parameter	Description
port <portlist>	The port(s) to clear statistics for. If no port is specified, the system defaults to the port entered when entering the Interface Configuration command mode.

show stack port-statistics command

The `show stack port-statistics` command displays port counters including congestion and dropped packets on stacks of ERS 5500 series switches.

The syntax for the `show stack port-statistics` command is:

```
show stack port-statistics {down-stack|up-stack} [unit<1-8>]
```

where `down-stack` displays the statistics from the down stack port from the requested unit, and `up-stack` displays the statistics from the up stack port from the requested unit.

This `show stack port-statistics` command is available in the PrivExec mode.

clear stack port-statistics command

The `clear stack port-statistics` command clears the statistics counters.

The syntax for the `clear stack port-statistics` command is:

```
clear stack port-statistics {down-stack|up-stack} [unit
<1-8>]
```

where `down-stack` clears the statistics from the down stack port from the requested unit, and `up-stack` clears the statistics from the up stack port from the requested unit.

The `clear stack port-statistics` command is available in the PrivExec mode.

Stack loopback tests

You can quickly test your stack ports and stack cable using the stack loopback test. This is very useful when you need to determine whether the source of the problem is a bad stack cable or a damaged stack port, and can help prevent unnecessarily sending switches for service.

There are two types of loopback tests. The internal loopback test verifies that the stack ports are functional.

The external loopback test checks the stack cable to find out if it is the source of the problem. The external loopback test is done by connecting the stack uplink port with the stack downlink port, sending a packet from the uplink port and verifying that the packet is received on the downlink port.

The internal test should always be run first, because the cable tests won't be conclusive until the stack ports are proven to be working correctly.

Testing the stack ports

Verify the operation of the stack ports using the internal loopback test.

Prerequisites for testing the stack ports

*   Any previous tests have finished running.

Step	Action
1	Remove the unit you want to test from the stack.
2	Disconnect the stacking cables and any active cables connected to the ports on the unit.
3	Go to NNCLI via serial interface.
4	Enable action mode. **enable**
5	Run the internal loopback test. **stack loopback-test internal**
6	Check the results of the command. The "stack loopback-test internal output" (page 74) table shows the output from a pass and from a failure.

stack loopback-test internal output

Pass	Fail

`Testing uplink port ... ok`
`Testing downlink port ... ok`
`Internal loopback test PASSED`
	`Testing uplink port ... Failed`
	`Testing downlink port ... ok`
	`Internal loopback test FAILED`

—End—

Testing the stack cables
Verify the integrity of the stack cables using the external loopback test.

Prerequisites for testing stack cables

- The internal loopback test passed. See "Testing the stack ports" (page 73).

- Any previous tests have finished running.

Step	Action
1	Remove the suspect unit from the stack.
2	Power up the suspect unit.
3	Connect a stacking cable from the stack uplink port to the stack downlink port.

4   Go to NNCLI via serial interface.

5   Enable action mode.

enable

6   Run the external loopback test.

stack loopback-test external

7   Check the results of the command. shows the output from a pass and from a failure.

stack loopback-test external output

Pass	Fail
External loopback test PASSED.	External loopback test FAILED. Your stack cable might be damaged.

—End—

Stack monitor

The Stack Monitor uses a set of control values to enable its operation, to set the expected stack size, and to control the frequency of trap sending. The stack monitor, if enabled, detects problems with the units in the stack and sends a trap.

The stack monitor sends a trap for the following events.

•   The number of units in stack changes.

•   The trap sending timer expires.

Each time the number of units in a stack changes, the trap sending timer resets and the stack monitor compares the current number of stack units with the configured number of stack units. If it is not equal, the switch sends a trap and logs a message to syslog. The stack monitor sends traps from a standalone unit or the base unit of the stack.

When the trap sending timer reaches the configured number of seconds at which traps are sent, the switch sends a trap and logs a message to syslog and restarts the trap sending timer. The syslog message is not repeated unless the stack configuration changes. This prevents the log being filled with stack configuration messages.

When you enable the stack monitor on a stack, the stack monitor captures current stack size and uses it as the expected stack size. You can choose a different value and set it after enabling the feature.

Displaying stack monitor configuration parameters values

Use the `show stack-monitor` command to display the current configuration values for the stack monitor.

Parameter	Description
Status	Enabled or disabled.
Stack size	Number of units in the stack <2-8>. Default value is 2.
Trap interval	Number of seconds between traps <3-300>. Default value is 60 seconds if there are 2 units in the stack.

Configuring stack monitor parameters

Use the `config stack-monitor` command to configure the stack monitor.

Configuring stack monitor parameters

Parameter	Description
enable/disable	Enable or disable stack monitoring
stack-size	Set stack size to be monitored <2-8>
trap-interval	Set interval between traps (seconds) <3-300>

Configuring stack monitor parameters with the JDM

To configure stack monitor parameters with the JDM, click **Edit > Chassis > Stack Monitor**.

Configuring the stack monitor with the JDM

Configuring stack monitor parameters

Parameter	Description
StackErrorNotificationEnabled	Enable or disable stack monitoring
ExpectedStackSize	Set stack size to be monitored <2-8>
StackErrorNotificationInterval	Set interval between traps (seconds) <3-300>

Click **Apply** to apply new settings.

Diagnostic information in the Web-based Management Interface

The procedures detailed in this section enable the viewing of diagnostic and statistical information through the Web-based Management Interface.

Viewing port statistics (Web)

To view statistical data about a selected port, perform the following task:

Step	Action
1	Open the **Port Statistics** screen by selecting **Statistics > Port** from the menu. This screen is illustrated below.

Port Statistics screen
Statistics > Port

Port Statistics (View By)
Port   1

Submit

Port Statistics Table			
Received		**Transmitted**	
Packets	78143	Packets	9583
Multicasts	58130	Multicasts	6938
Broadcasts	17683	Broadcasts	91
Total Octets	8495670	Total Octets	1928558
Pause Frames	0	Pause Frames	0
FCS/Frame Errors	0	Collisions	0
Undersized Packets	0	Single Collisions	0
Oversized Packets	0	Multiple Collisions	0
Filtered Packets	23	Excessive Collisions	0
		Deferred Packets	0
		Late Collisions	0
Packets Received and Transmitted			
64 bytes	50312	256-511 bytes	2030
65-127 bytes	25990	512-1023 bytes	1723
128-255 bytes	7314	1024-1518 bytes	357
		1522-9216 bytes (Jumbo)	0

Update   Zero Port   Zero All Ports

Step	Action
2	Select a port from the **Port** list in the **Port Statistics (View By)** section.
3	Click **Submit**.

—End—

Port statistics are displayed in the **Port Statistics Table** section. The following table describes the fields in this section.

Port Statistics Table fields

Field	Description
Packets	The number of packets received/transmitted on this port, including bad packets, broadcast packets, and multicast packets.
Multicast	The number of good multicast packets received/transmitted on this port, excluding broadcast packets.
Broadcasts	The number of good broadcast packets received/transmitted on this port.
Total Octets	The number of octets of data received/transmitted on this port, including data in bad packets and FCS octets, and framing bits.
Pause Frames	The number of pause frames received/transmitted on this port.
FCS-Frame Errors	The number of valid-size packets received on this port with proper framing but discarded because of FCS or frame errors.
Undersized Packets	The number of packets received on this port with fewer than 64 bytes and with proper CRC and framing (also known as short frames or runts).
Oversized Packets	The number of packets that were received on this port with proper CRC and framing that meet the following requirements: 1518 bytes if no VLAN tag exists 1522 bytes if a VLAN tag exists
Filtered Packets	The number of packets that were received on this port and discarded because of the specific configuration.This counter does not count the FCS/Frames error packets; they are counted in that counter. This counter counts packets discarded because STP is not set to forwarding, the frame setting in VLAN directs discarding, or a mismatch in ingress/egress port speeds.
Collisions	The number of collisions detected on this port.
Single Collisions	The number of packets that were transmitted successfully on this port after a single collision.

Field	Description
Multiple Collisions	The number of packets that were transmitted successfully on this port after more than one collision.
Excessive Collisions	The number of packets lost on this port due to excessive collisions.
Late Collisions	The number of packets collisions that occurred after a total length of time that exceeded 512 bit-times of packet transmission.
Deferred Packets	The number of packets that were received on this port that were delayed on the first transmission attempt, but never incurred a collision.
Packets 64 bytes 65-127 bytes 128-255 bytes 256-511 bytes 512-1023 bytes 1024-1518 bytes 1522-9216 bytes	The number of packets received/transmitted on the port.

Three further options are available on this screen:

1. Click **Update** to refresh the statistical information.

2. Click **Zero Port** to reset the counters for the selected port.

3. Click **Zero All Ports** to reset the counters for all ports.

Viewing all port errors

To view a summary of the port errors, follow this procedure:

Step	Action
1	Open the **Port Error Summary** screen by selecting **Statistics > Port Error Summary** from the menu. This screen is illustrated below.

Port Error Summary screen

Statistics > Port Error Summary

Port Error Summary Table								
Port	Status	Link	Speed/Duplex	FCS/Frame Errors	Collisions	Single Collisions	Multiple Collisions	Excessive Collisions	Late Collisions

Update

The following table describes the fields on this screen.

Port Error Summary fields

Field	Description
Unit	Displays the unit number in the stack.
Port	Displays the port number of the unit.
Status	Displays the status of the port (Enabled/Disabled).
Link	Displays the link status of the port (Up/Down).
Speed/Duplex	Displays the speed at which the port is operating, as well as whether it is in half- or full-duplex mode.
FCS/Frame Errors	Displays the number of frame check sequence (FCS) and frame errors received on this port.
Collisions	Displays the number of collisions errors received on this port.
Single Collisions	Displays the number of single collisions errors received on this port.
Multiple Collisions	Displays the number of multiple collisions errors received on this port.
Excessive Collisions	Displays the number of excessive collisions errors received on this port.
Late Collisions	Displays the number of late collisions errors received on this port.

—End—

Click **Update** to refresh the statistical information.

Viewing interface statistics

To view statistical information for an interface, follow this procedure:

Step	Action
1	Open the **Interface Statistics** screen by selecting **Statistics > Interface** from the menu. This screen is illustrated in the following table.

Interface Statistics screen

Statistics > Interface

Port	In Octets	Out Octets	In Unicast	Out Unicast	In Non-Unicast	Out Non-Unicast	In Discards	Out Discards	In Errors	Out Errors	In Unknown Protos
1	8563956	2064248	2501	2757	76122	7057	23	0	0	0	0
2	0	0	0	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0	0	0	0
4	0	0	0	0	0	0	0	0	0	0	0
5	0	0	0	0	0	0	0	0	0	0	0
6	0	0	0	0	0	0	0	0	0	0	0
7	0	0	0	0	0	0	0	0	0	0	0
8	0	0	0	0	0	0	0	0	0	0	0
9	0	0	0	0	0	0	0	0	0	0	0
10	0	0	0	0	0	0	0	0	0	0	0
11	0	0	0	0	0	0	0	0	0	0	0
12	0	0	0	0	0	0	0	0	0	0	0
13	0	0	0	0	0	0	0	0	0	0	0
14	0	0	0	0	0	0	0	0	0	0	0
15	0	0	0	0	0	0	0	0	0	0	0
16	0	0	0	0	0	0	0	0	0	0	0
17	0	0	0	0	0	0	0	0	0	0	0
18	0	0	0	0	0	0	0	0	0	0	0
19	0	0	0	0	0	0	0	0	0	0	0
20	0	0	0	0	0	0	0	0	0	0	0
21	0	0	0	0	0	0	0	0	0	0	0
22	0	0	0	0	0	0	0	0	0	0	0
23	0	0	0	0	0	0	0	0	0	0	0
24	0	0	0	0	0	0	0	0	0	0	0
25	0	0	0	0	0	0	0	0	0	0	0
26	0	0	0	0	0	0	0	0	0	0	0

[Update]

The following table describes the fields on this screen.

Interface Statistics screen

Field	Description
Port	The port number corresponding to the selected switch.
In Octets	The number of octets received on the interface, including framing characters.
Out Octets	The number of octets transmitted out of the interface, including framing characters.
In Unicast	The number of subnetwork-unicast packets delivered to a higher-layer protocol.
Out Unicast	The number of packets that higher-layer protocols requested be transmitted to a subnetwork-unicast address, including those that were discarded or not sent.

Field	Description
In Non-Unicast	The number of non-unicast packets, for example, subnetwork-broadcast or subnetwork-multicast packets, delivered to a higher protocol.
Out Non-Unicast	The number of packets that higher-level protocols requested be transmitted to a non-unicast address. For example, a subnetwork-broadcast or a subnetwork multicast address, including those that were discarded or not sent.
In Discards	The number of inbound packets which were selected to be discarded even though no errors were detected to prevent their being delivered to a higher-layer protocol. Packet discarding is not arbitrary. One reason for discarding packets is to free buffer space.
Out Discards	The number of outbound packets which were selected to be discarded even though no errors were detected to prevent their being transmitted. Packet discarding is not arbitrary. One reason for discarding packets is to free buffer space.
In Errors	The number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol.
Out Errors	The number of outbound packets that could not be transmitted because of errors.
In Unknown Protos	The number of packets received through the interface that were discarded because of an unknown or unsupported protocol.

—End—

Click **Update** to refresh the statistical information.

Viewing Ethernet error statistics

To view Ethernet error statistics, follow this procedure:

Step	Action
1	Open the **Ethernet Errors** screen by selecting **Statistics > Ethernet Errors** from the menu. This screen is illustrated in the following table.

Ethernet Errors screen

Statistics > Ethernet Errors

Ethernet Errors Statistics Table

Port	FCS/Frame Errors	Internal MAC Transmit Errors	Internal MAC Receive Errors	Carrier Sense Errors	SQE Test Errors	Deferred Transmissions	Single Collisions Frames	Multiple Collisions Frames	Late Collisions	Excessive Collisions
1	0	0	0	0	0	0	0	0	0	0
2	0	0	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0	0	0
4	0	0	0	0	0	0	0	0	0	0
5	0	0	0	0	0	0	0	0	0	0
6	0	0	0	0	0	0	0	0	0	0
7	0	0	0	0	0	0	0	0	0	0
8	0	0	0	0	0	0	0	0	0	0
9	0	0	0	0	0	0	0	0	0	0
10	0	0	0	0	0	0	0	0	0	0
11	0	0	0	0	0	0	0	0	0	0
12	0	0	0	0	0	0	0	0	0	0
13	0	0	0	0	0	0	0	0	0	0
14	0	0	0	0	0	0	0	0	0	0
15	0	0	0	0	0	0	0	0	0	0
16	0	0	0	0	0	0	0	0	0	0
17	0	0	0	0	0	0	0	0	0	0
18	0	0	0	0	0	0	0	0	0	0
19	0	0	0	0	0	0	0	0	0	0
20	0	0	0	0	0	0	0	0	0	0
21	0	0	0	0	0	0	0	0	0	0
22	0	0	0	0	0	0	0	0	0	0
23	0	0	0	0	0	0	0	0	0	0
24	0	0	0	0	0	0	0	0	0	0
25	0	0	0	0	0	0	0	0	0	0
26	0	0	0	0	0	0	0	0	0	0

Update

The following table outlines the fields on this screen.

Ethernet Error fields

Field	Description
Port	The port number corresponding to the selected switch.
FCS/Frame Errors	The number of frames received on a particular interface that are an integral number of octets in length, but do not pass the FCS check or have frame errors.

Field	Description
Internal MAC Transmit Errors	The number of frames for which transmission on a particular interface fails because of an internal MAC sublayer transmit error. A frame only is counted by an instance of this object if it is not counted by the corresponding instance of either the dot3StatsLateCollisions object, the dot3StatsExcessiveCollisions object, or the dot3StatsCarrierSenseErrors object.
Internal MAC Receive Errors	The number of frames for which reception on a particular interface fails because of an internal MAC sublayer transmit error. A frame only is counted by an instance of this object if it is not counted by the corresponding instance of either the dot3StatsLateCollisions object, the dot3StatsExcessiveCollisions object, or the dot3StatsCarrierSenseErrors object.
Carrier Sense Errors	The number of times that the carrier sense conditions was lost or never asserted when attempting to transmit a frame on a particular interface.
SQE Test Errors	The number of times that the SQE TEST ERROR message is generated by the PLS sublayer for a particular interface. The SQE TEST ERROR 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.
Deferred Transmissions	The number of frames for which the first transmission attempt on a particular interface is delayed because the medium is busy.
Single Collision Frames	The number of successfully transmitted frames on a particular interface for which transmission is inhibited by more than one collision.
Multiple Collision Frames	The number of successfully transmitted frames on a particular interface for which transmission is inhibited by a single collision.

Field	Description
Late Collisions	The number of times a collision is detected on a particular interface later than 512 bit-times into the transmission of a packet.
Excessive Collisions	The number of frames for which transmission on a particular interface fails due to excessive collisions.

—End—

Click **Update** to refresh the statistical information.

Viewing transparent bridging statistics

To view transparent bridging statistics, follow this procedure:

Step	Action
1	Open the **Transparent Bridging** screen by selecting **Statistics > Transparent Bridging** from the menu. This screen is illustrated below.

Transparent Bridging screen

Statistics > Transparent Bridging

Transparent Bridging Statistics Table			
Port	In Frames	Out Frames	In Discards
1	80673	10439	23
2	0	0	0
3	0	0	0
4	0	0	0
5	0	0	0
6	0	0	0
7	0	0	0
8	0	0	0
9	0	0	0
10	0	0	0
11	0	0	0
12	0	0	0
13	0	0	0
14	0	0	0
15	0	0	0
16	0	0	0
17	0	0	0
18	0	0	0
19	0	0	0
20	0	0	0
21	0	0	0
22	0	0	0
23	0	0	0
24	0	0	0
25	0	0	0
26	0	0	0

Update

The following table describes the fields on this screen.

Transparent Bridging screen

Field	Description
Port	The port number that corresponds to the selected switch.
In Frames (dot1dTpPortInFrames)	The number of frames that have been received by this port from its segment. A frame received on the interface corresponding to this port is counted only if it is for a protocol being processed by the local bridging function, including bridge management errors.

Field	Description
Out Frames (dot1dTpPortOutFrames)	The number of frames that have been transmitted by this port from its segment. A frame received on the interface corresponding to this port is counted only if it is for a protocol being processed by the local bridging function, including bridge management errors.
In Discards (dot1dTpPortInDiscards)	The number of valid frames received which were discarded by the forwarding process.

—End—

Click **Update** to refresh the statistical information.

Monitoring MLT traffic

Bandwidth usage can be monitored for the Multilink Trunk (MLT) member ports within each trunk in a configuration by selecting the traffic type to monitor.

To monitor MultiLink Trunk traffic, follow this procedure:

Step	Action

1   Open the **MLT Utilization** screen by selecting **Application > MultiLink Trunk > Utilization** from the menu. This screen is illustrated below.

MLT Utilization screen

2   In the **MultiLink Trunk Utilization Selection (View By)** section, select a trunk to monitor in the **Trunk** list and a type of traffic in the **Traffic Type** list.

3   Click **Submit**.

—End—

The **MultiLink Trunk Utilization Table** section will be populated with information. The following table describes the fields in this table.

MultiLink Trunk Utilization Table fields

Field	Description
Unit/Port	A list of the trunk member switch ports that correspond to the trunk specified in the Trunk column.
Last 5 Minutes	The percentage of packets (of the type specified in the Traffic Type field) used by the port in the last five minutes. This field provides a running average of network activity, and is updated every 15 seconds.
Last 30 Minutes	The percentage of packets (of the type specified in the Traffic Type field) used by the port in the last 30 minutes. This field provides a running average of network activity, and is updated every 15 seconds.
Last Hour	The percentage of packets (of the type specified in the Traffic Type field) used by the port in the last 60 minutes. This field provides a running average of network activity, and is updated every 15 seconds.

Chapter 3
Configuring Remote Network Monitoring (RMON)

The Remote Network Monitoring (RMON) Management Information Base (MIB) is an interface between the RMON agent on the Nortel Ethernet Routing Switch 5500 Series and an RMON management application, such as the Java Device Manager.

It 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 monitors switch performance.

RMON has three major functions:

*   Creating and displaying alarms for user-defined events

*   Gathering cumulative statistics for Ethernet interfaces

*   Tracking a history of statistics for Ethernet interfaces

Configuring RMON with the CLI

This section describes the CLI commands used to configure and manage RMON.

show rmon alarm

The **show rmon alarm** command displays information on RMON alarms.

The syntax for the **show rmon alarm** command is:

show rmon alarm

The **show rmon alarm** command is executed in the Privileged EXEC mode.

show rmon event

The `show rmon event` command displays information regarding RMON events.

The syntax for the `show rmon event` command is:

`show rmon event`

The `show rmon event` command is executed in the Privileged EXEC command mode.

show rmon history

The `show rmon history` command displays information regarding the configuration of RMON history.

The syntax for the `show rmon history` command is:

`show rmon history`

The `show rmon history` command is executed in the Privileged EXEC command mode.

show rmon stats

The `show rmon stats` command displays information regarding the configuration of RMON statistics.

The syntax for the `show rmon stats` command is:

`show rmon stats`

The `show rmon stats` command is executed in the Privileged EXEC command mode.

rmon alarm

The `rmon alarm` command allows you to set RMON alarms and thresholds.

The syntax for the `rmon alarm` command is:

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

outlines the parameters for this command.

rmon alarm parameters

Parameter	Description
<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, an English name may be used.
<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 will be triggered.
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 will be triggered.
[owner <LINE>]	Specify an owner string to identify the alarm entry.

The `rmon alarm` command is executed in the Global Configuration command mode.

no rmon alarm

The `no rmon alarm` command deletes RMON alarm table entries. When the variable is omitted, all entries in the table are cleared.

The syntax for the `no rmon alarm` command is:

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

Substitute `<1-65535>` above with the unique ID of the alarm entry.

The `no rmon alarm` command is executed in the Global Configuration command mode.

rmon event

The `rmon event` configures RMON event log and trap settings.

The syntax for the `rmon event` command is:

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

"rmon event parameters" (page 92) outlines the parameters for this
command.

rmon event parameters

Parameter	Description
<1-65535>	Unique index for the event entry.
[log]	Record events in the log table.
[trap]	Generate SNMP trap messages for events.
[description <LINE>]	Specify a textual description for the event.
[owner <LINE>]	Specify an owner string to identify the event entry.

The `rmon event` command is executed in the Global Configuration
command mode.

no rmon event

The `no rmon event` deletes RMON event table entries. When the variable
is omitted, all entries in the table are cleared.

The syntax for the `no rmon event` command is:

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

Substitute `<1-65535>` above with the unique ID of the event to be deleted.

The `no rmon event` command is executed in the Global Configuration
command mode.

rmon history

The `rmon history` configures RMON history settings.

The syntax for the `rmon history` command is:

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

"rmon history parameters" (page 92) outlines the parameters for this
command.

rmon history parameters

Parameter	Description
<1-65535>	Unique index for the history entry.
<LINE>	Specify the port number to be monitored.

Parameter	Description
<1-65535>	The number of history buckets (records) to keep.
<1-3600>	The sampling rate (how often a history sample is collected).
[owner <LINE>]	Specify an owner string to identify the history entry.

The **rmon history** command is executed in the Global Configuration command mode.

no rmon history

The **no rmon history** deletes RMON history table entries. When the variable is omitted, all entries in the table are cleared.

The syntax for the **no rmon history** command is:

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

Substitute **<1-65535>** above with the unique ID of the history entry.

The **no rmon history** command is executed in the Global Configuration command mode.

rmon stats

The **rmon stats** command configures RMON statistics settings.

The syntax for the **rmon stats** command is:

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

"rmon stats parameters" (page 93) outlines the parameters for this command.

rmon stats parameters

Parameter	Description
<1-65535>	Unique index for the stats entry.
[owner <*LINE*>]	Specify an owner string to identify the stats entry.

The **rmon stats** command is executed in the Global Configuration command mode.

no rmon stats

The **no rmon stats** turns off RMON statistics. When the variable is omitted, all entries in the table are cleared.

The syntax for the **no rmon stats** command is:

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

Substitute `<1-65535>` above with the unique ID of the stats entry.

The `no rmon stats` command is executed in the Global Configuration command mode.

Configuring RMON with the Web-based Management Interface

This section discusses the configuration and management of RMON using the Web-based Management Interface.

Configuring RMON fault threshold parameters

Alarms are used to alert a system administrator when the value of a variable goes out of range. RMON alarms can be defined on any MIB variable that resolves to an integer value. String variables (such as system description) cannot be used as alarm variables.

Creating an RMON fault threshold

To configure an RMON fault threshold, follow this procedure:

Step	Action
1	Open the **RMON Threshold** screen by selecting **Fault > RMON Threshold** from the menu. This screen is illustrated below.

RMON Threshold screen
Fault > RMON Threshold

2   In the fields provided in the **RMON Threshold Creation** section, enter the information for the new threshold. The following tables outlines the fields in this section.

RMON Threshold Creation fields

Field	Description
Alarm Index	Type the unique number to identify the alarm entry.
Port	Choose the port on which to set an alarm.
Parameter	Choose the sampled statistic.
Rising Level	Type the event entry to be used when a rising threshold is crossed.
Falling Level	Type the event entry to be used when a falling threshold is crossed.
Rising Action	Choose the type of notification for the event. Selecting Log generates an entry in the RMON Event Log table for each event. Selecting SNMP Trap sends an SNMP trap to one or more management stations.
Interval	Type the time period (in seconds) to sample data and compare the data to the rising and falling thresholds.
Alarm Sample	Choose the sampling method: Absolute: *Absolute* alarms are defined on the current 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 for absolute value. Therefore, an alarm could be created with a rising value of 2 and a falling value of 1 to alert a user to whether the card is up or down. Delta: 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 per 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. Therefore, if you keep track of 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.)

3   Click **Submit**.

—End—

The new RMON threshold is displayed in the **RMON Threshold Table** section.

Deleting an RMON threshold configuration

To delete an existing RMON threshold configuration, follow this procedure:

Step	Action
1	Open the **RMON Threshold** screen by selecting **Fault > RMON Threshold** from the menu. This screen is illustrated above.
2	In the **RMON Threshold Table**, click the **Delete** icon in the row of the entry to be deleted.
3	A message prompts for confirmation of the request. Click **Yes**.

—End—

Viewing the RMON fault event log

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 fires at the rising threshold, the rising event specifies that this information be sent to both a trap and a log. The **RMON Event Log** screen works in conjunction with the **RMON Threshold** screen to enable viewing the history of RMON fault events.

To view a history of RMON fault events, follow this procedure:

Step	Action
1	Open the **RMON Event Log** screen by selecting **Fault > RMON Event Log** from the menu. This screen is illustrated below.

RMON Event Log screen
Fault > RMON Event Log

RMON Event Log
Time Stamp | Description | Triggered By | ID

—End—

The RMON event log is displayed.

Configuring RMON with the Java Device Manager
This section will discuss the configuration and management of RMON using
the Java Device Manager (JDM).

See also

- "Working with RMON information" (page 97)
- "Alarms" (page 107)
- "Events" (page 114)
- "Log information" (page 118)

Working with RMON information
RMON information is viewed by looking at the graphing information
associated with the port or chassis.

See also

- "Viewing statistics" (page 97)
- "Viewing history" (page 101)
- "Enabling ethernet statistics gathering" (page 105)
- "Disabling Ethernet statistics gathering" (page 107)

Viewing statistics
The JDM gathers Ethernet statistics that can be graphed in a variety of
formats or saved to a file that can be exported to an outside presentation or
graphing application.

To view RMON ethernet statistics:

Step	Action
1	Select a port.
2	Do one of the following:

a. From the shortcut menu, choose **Graph**.

b. Select **Graph > Port** from the menu.

c. On the toolbar, click the **Graph** button.

3 The **Graph Port** screen opens. Click the **RMON** tab. This tab is illustrated below.

Graph Port screen - RMON tab

	AbsoluteValue	Cumulative	Average/sec	Minimum/sec	Maximum/sec	LastVal/sec
Octets	1,231,702	1,798	179.8	179.8	179.8	179.8
Pkts	9,805	18	1.8	1.8	1.8	1.8
BroadcastPkts	2,008	3	0.3	0.3	0.3	0.3
MulticastPkts	6,613	14	1.4	1.4	1.4	1.4
CRCAlignErrors	0	0	0	0	0	0
UndersizePkts	0	0	0	0	0	0
OversizePkts	0	0	0	0	0	0
Fragments	0	0	0	0	0	0
Collisions	0	0	0	0	0	0
Jabbers	0	0	0	0	0	0
1..64	6,421	12	1.2	1.2	1.2	1.2
65..127	3,238	7	0.7	0.7	0.7	0.7
128..255	876	0	0	0	0	0
256..511	449	2	0.2	0.2	0.2	0.2
512..1023	786	0	0	0	0	0
1024..1518	160	0	0	0	0	0

Interface | Ethernet Errors | Bridge | Rmon | EAPOL Stats | EAPOL Diag | LACP | Misc.

192.168.249.46 - Graph Port 1/1

Clear Counters | Close | Help... | Poll Interval: 10s

—End—

The following table describes the fields on the RMON tab.

Graph Port screen - RMON tab

Field	Descriptions
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 were directed to the broadcast address. Note that this does not include multicast packets.

Field	Descriptions
MulticastPkts	The total number of good packets received that were directed to a multicast address. Note that this number does not include packets directed to the broadcast address.
CRCAlignErrors	The total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, but had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).
UndersizePkts	The total number of packets received that were less than 64 octets long (excluding framing bits but including FCS octets) and were otherwise well formed.
OversizePkts	The total number of packets received that were longer than 1518 octets (excluding framing bits but including FCS octets) and were otherwise well formed.
OversizePkts	The total number of packets received that were longer than 1518 octets (excluding framing bits but including FCS octets) and were otherwise well formed.
Fragments	The total number of packets received that were less than 64 octets in length (excluding framing bits but including FCS octets) and had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). It is entirely normal for etherStatsFragments to increment 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 were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a 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) that were transmitted and received on this port between 1 and 64 octets in length (excluding framing bits but including FCS octets).

Field	Descriptions
65..127	The total number of packets (including bad packets) that were transmitted and received on this port between 65 and 127 octets in length (excluding framing bits but including FCS octets).
128..255	The total number of packets (including bad packets) that were transmitted and received on this port between 128 and 255 octets in length (excluding framing bits but including FCS octets).
256..511	The total number of packets (including bad packets) that were transmitted and received on this port between 256 and 511 octets in length (excluding framing bits but including FCS octets).
512..1023	The total number of packets (including bad packets) that were transmitted and received on this port between 512 and 1023 octets in length (excluding framing bits but including FCS octets).
1024..1518	The total number of packets (including bad packets) that were transmitted and received on this port between 1024 and 1518 octets in length (excluding framing bits but including FCS octets).

Types of statistics

Statistic	Description
Poll Interval	Statistics are updated based on the poll interval.
	Default: 10s
	Range: None, 2s, 5s, 10s, 30s, 1m, 5m, 30m 1h
Absolute	The total count since the last time counters were reset. A system reboot resets all counters.
Cumulative	The total count since the statistics tab was first opened. The elapsed time for the cumulative counter is shown at the bottom of the graph window.
Average/sec	The cumulative count divided by the cumulative elapsed time.
Min/sec	The minimum average for the counter for a given polling interval over the cumulative elapsed time.
Max/sec	The maximum average for the counter for a given polling interval over the cumulative elapsed time.
Last/sec	The average for the counter over the last polling interval.

See also

Viewing history

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:

- Buckets are gathered at 30-minute intervals.

- Number of buckets gathered is 50.

Both the time interval and the number of buckets is configurable. 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.

To view RMON history:
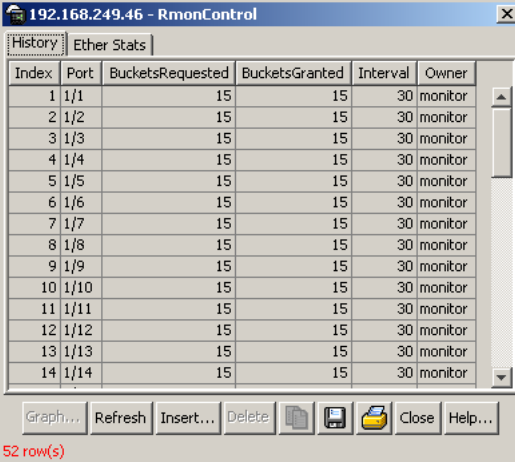
Step	Action

1   Open the **RmonControl** screen by selecting **Serviceability > RMON > Control** from the menu. This screen is illustrated below.

RmonControl screen - History tab

"History tab fields" (page 102) describes the fields on the **History** tab.

—End—

Creating a history

RMON can be used to collect statistics at intervals. For example, if switch performance will be monitored over a weekend, enough buckets to cover two days must be set aside. To do this, set the history to gather one bucket each hour, thus covering a 48-hour period. After history characteristics are set, they cannot be modified; the history must be deleted and another created.

To establish a history for a port and set the bucket interval:

Step	Action
1	Open the **RmonControl** screen by selecting **Serviceability > RMON > Control** from the menu.
2	Click **Insert**.

The **Insert History** screen opens. This screen is illustrated below.

Insert History screen

Step	Action
3	In the fields provided, enter the information for the new RMON history. The fields on this screen are described in the table below.
4	Click **Insert**.

—End—

The following table describes the **History** tab of the RmonControl dialog box.

History tab fields

Field	Description
Index	A unique value assigned to each interface. An index identifies an entry in a table.
Port	Any Ethernet interface on the device.
BucketsRequested	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.

Field	Description
BucketsGranted	The number of discrete sampling intervals over which data is saved in the part of the media-specific table associated with this entry. There are instances when the actual number of buckets associated with this entry is less than the value of this object. In this case, at the end of each sampling interval, a new bucket is added to the media-specific table.
Interval	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 may overflow at their maximum value with no indication, note the possibility of overflow in any of the associated counters. It is important to 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 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 one hour at the Ethernet's maximum utilization.
Owner	The network management system that created this entry.

See also

- "Working with RMON information" (page 97)

- "Viewing statistics" (page 97)

- "Disabling history" (page 103)

- "Disabling Ethernet statistics gathering" (page 107)

Disabling history
To disable RMON history on a port:

Step	Action
1	Open the **RmonControl** screen by selecting **Serviceability > RMON > Control** from the menu.
2	Highlight the row that contains the record to delete.
3	Click **Delete**.

—End—

See also

- "Working with RMON information" (page 97)
- "Viewing statistics" (page 97)
- "Creating a history" (page 102)
- "Enabling ethernet statistics gathering" (page 105)
- "Disabling Ethernet statistics gathering" (page 107)

Viewing RMON history statistics

To display Rmon History statistics:

Step	Action
1	Open the **RmonControl** screen by selecting **Serviceability > RMON > Control** from the menu.
2	Select a port in the **RMON History** tab.
3	Click **Graph**.
4	The **RMON History** screen opens for the selected port. This screen is illustrated below.

Rmon History statistics

—End—

The following table describes the **RMON History** screen fields.

RMON History screen fields

Field	Description
SampleIndex	Indicates the sample number. As history samples are taken, they are assigned greater sample numbers.
Utilization	Estimates the percentage of a link's capacity that was 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.
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 were dropped because of system resource constraints.
CRCAlignErrors	The number of packets received during a sampling interval that were between 64 and 1518 octets long. This length included 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 non-integral number of octets (Alignment Error).
UndersizePkts	The number of packets received during the sampling interval were less than 64 octets long (including FCS octets, but not framing bits.
OversizePkts	The number of packets received during the sampling interval were longer than 1518 octets (including FCS octets, but not framing bits, and were otherwise well formed.
Fragments	The number of packets received during the sampling interval were 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 non-integral number of octets (Alignment Error).
Collisions	The best estimate of the number of collisions on an Ethernet segment during a sampling interval.

Enabling ethernet statistics gathering
To gather ethernet statistics:

Step	Action
1	Open the **RmonControl** screen by selecting **Serviceability > RMON > Control** from the menu.

2 Select the **Ether Stats** tab. This tab is illustrated below.

RmonControl screen - Ether Stats tab

3 Select an **Index** and click **Insert**.

The **Insert Ether Stats** screen opens.

Insert Ether Stats screen

4 Enter the ports to be used. Port numbers can be manually entered into the **Port** field or selected by clicking the ellipsis (...) and using the **Port List** screen to make the selections.

Insert Ether Stats Port List screen

5 Enter the owner of this RMON entry in the **Owner** field.

6 Click **Insert**.

—End—

The following table describes the **Ether Stats** tab fields.

Ether Stats tab fields

Field	Description
Index	A unique value assigned to each interface. An index identifies an entry in a table.
Port	Any port on the device.
Owner	The network management system that created this entry.

Disabling Ethernet statistics gathering

To disable Ethernet statistics, follow this procedure:

Step	Action
1	Open the **RmonControl** screen by selecting **Serviceability > RMON > Control** from the menu.
2	Select the **Ether Stats** tab.
3	Highlight the row that contains the record to delete.
4	Click **Delete**.

—End—

Alarms

Alarms are useful when you need to know when the values of a variable go out of range. Define an RMON alarm for any MIB variable that resolves to an integer value. String variables cannot be used. All alarms share the following characteristics:

- An upper and lower threshold value is defined.

- A corresponding rising and falling event occurs.

- An alarm interval or polling period is reached.

When alarms are activated, view the activity in a log or a trap log, or a script can be created to provide notification by beeping a console, sending e-mail messages, or calling a pager.

How RMON alarms work

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, then the alarm fires 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 fires as a rising event. During the first interval that the data drops below the falling value, the alarm fires as a falling event ("How alarms fire" (page 108)).

How alarms fire

It is important to note that the alarm fires 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, then 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 should provide the notification the system administrator needs 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 fires. When outbound traffic other than spanning tree ceases, the falling alarm fires. This process provides the system administrator with time intervals of any non-baseline outbound traffic.

If the alarm is defined with a falling threshold less than 260 (assuming the alarm polling interval is 10 seconds), say 250, then the rising alarm can fire only once ("Alarm example - threshold less than 260" (page 109)). The reason is that 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 would cause 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.

Alarm example - threshold less than 260

See also

- "Alarms" (page 107)
- "Creating alarms" (page 109)
- "Alarm Manager" (page 110)

Creating alarms

When creating an alarm, select a variable from the variable list and a 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 for absolute value. Therefore, an alarm could be created with a rising value of 2 and a falling value of 1 to alert a user to whether the card is up or down.

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 per 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, therefore, the result is twice the actual value. (This result is not an error in the software.)

Alarm Manager

> *Note:* The example alarm described in the following procedure
> generates at least one alarm every five minutes. The example is
> intended only to demonstrate how alarms fire; it is not a useful alarm.

Creating an Alarm

To create an alarm to receive statistics and history using default values:

Step	Action

---

1 Open the **Alarm Manager** screen by selecting **Serviceability >
RMON > Alarm Manager** from the menu. This screen is illustrated
below.

Alarm Manager screen

2 In the **Variable** field, select a variable and a port (or other ID) from
the list to set the alarm on below.

Alarm Manager Variable List

Alarm variables are in three formats, depending on the type:

- A chassis alarm ends in .x where the x index is hard-coded. No further information is required.

- A card, spanning tree group (STG) or EtherStat alarm ends with a dot (.). A card number, STG ID, IP address, or EtherStat information must be entered.

- A port alarm ends with no dot or index and requires using the port shortcut menu. An example of a port alarm is ifInOctets (interface incoming octet count).

3   In the remaining fields, enter the information for the alarm. The fields for this screen are described in table below.

4   Click **Insert**.

—End—

The following table describes the **RMON Insert Alarm** dialog box fields.

RMON Insert Alarm dialog box fields

Field	Description
Variable	Name and type of alarm--indicated by the format:
	alarmname.x where x=0 indicates a chassis alarm.
	alarmname. where the user must specify the index. This will be 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 results in display of the port selection tool.
Sample Type	Can be either absolute or delta.
	For more information about sample types, refer to "Creating alarms" (page 109).
Sample Interval	Time period (in seconds) over which the data is sampled and compared with the rising and falling thresholds.

Field	Description	
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.	
Threshold Type	Rising Value	Falling Value
Value	When the current sampled value is greater than or equal to this threshold, and the value at the last sampling interval was less than this threshold, generates a single event.	When the current sampled value is less than or equal to this threshold, and the value at the last sampling interval was greater than this threshold, generates a single event.
Event Index	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.)	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.)

Deleting an alarm

To delete an alarm:

Step	Action

1   Open the **Alarms** screen by selecting **Serviceability > RMON > Alarms** from the menu. This screen is illustrated below.

RmonAlarms screen - Alarms tab

2   Select the alarm to be deleted.

3   Click **Delete**.

—End—

The following table describes the fields on the **Alarms** tab.

Alarms tab fields

Field	Description
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.
Interval	The interval in seconds over which data is sampled and compared with the rising and falling thresholds. When setting this variable, note that in the case of deltaValue sampling, you should set the interval short enough that the sampled variable is very unlikely to increase or decrease by a delta of more than 2^31 - 1 during a single sampling interval.
Variable	The object identifier of the particular variable to be sampled. Only variables that resolve to an ASN.1 primitive type of INTEGER (INTEGER, Counter, Gauge, or TimeTicks) may be sampled.
Sample Type	The method of sampling the selected variable and calculating the value to be compared against the thresholds. If the value of this object is absoluteValue(1), the value of the selected variable will be compared directly with the thresholds at the end of the sampling interval. If the value of this object is deltaValue(2), the value of the selected variable at the last sample will be subtracted from the current value, and the difference compared with the thresholds.
Value	The value of the statistic during the last sampling period. For example, if the sample type is deltaValue, this value is the difference between the samples at the beginning and end of the period. If the sample type is absoluteValue, this value is the sampled value at the end of the period. This is the value that is compared with the rising and falling thresholds. The value during the current sampling period is not made available until the period is completed and remains available until the next period is completed.
StartupAlarm	The alarm that may be sent when this entry is first set to Valid. If the first sample after this entry becomes valid is greater than or equal to the risingThreshold and alarmStartupAlarm is equal to risingAlarm(1) or risingOrFallingAlarm(3), then a single rising alarm is generated. If the first sample after this entry becomes valid is less than or equal to the fallingThreshold and alarmStartupAlarm is equal to fallingAlarm(2) or risingOrFallingAlarm(3), then a single falling alarm is generated.

Field	Description
RisingThreshold	A threshold for the sampled statistic. When the current sampled value is greater than or equal to this threshold, and the value at the last sampling interval was less than this threshold, a single event is generated. A single event is also generated if the first sample after this entry becomes valid is greater than or equal to this threshold and the associated alarmStartupAlarm is equal to risingAlarm(1) or risingOrFallingAlarm(3). After a rising event is generated, another such event is not generated until the sampled value falls below this threshold and reaches the alarmFallingThreshold.
RisingEventIndex	The index of the eventEntry that is used when a rising threshold is crossed. The eventEntry identified by a particular value of this index is the same as identified by the same value of the eventIndex object. If there is no corresponding entry in the eventTable, then no association exists. In particular, if this value is zero, no associated event is generated, because zero is not a valid event index.
FallingThreshold	A threshold for the sampled statistic. When the current sampled value is less than or equal to this threshold, and the value at the last sampling interval was greater than this threshold, a single event is generated. A single event is also generated if the first sample after this entry becomes valid is less than or equal to this threshold and the associated alarmStartupAlarm is equal to fallingAlarm(2) or risingOrFallingAlarm(3). After a falling event is generated, another such event is not generated until the sampled value rises above this threshold and reaches the alarmRisingThreshold.
FallingEventIndex	The index of the eventEntry that is used when a falling threshold is crossed. The eventEntry identified by a particular value of this index is the same as identified by the same value of the eventIndex object. If there is no corresponding entry in the eventTable, then no association exists. In particular, if this value is zero, no associated event is generated, because zero is not a valid event index.
Owner	The network management system that created this entry.
Status	The status of this alarm entry.

Events

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.

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, two default events are generated:

- RisingEvent

- FallingEvent

The default events specify that when an alarm goes out of range, the "firing" of the alarm will be tracked in both a trap and a log. For example, when an alarm fires 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 be sent to a trap and a log.

See also

- "Alarms" (page 107)

- "How RMON alarms work" (page 107)

- "Creating alarms" (page 109)

- "Viewing an event" (page 115)

- "Log information" (page 118)

Viewing an event

To view a table of events:

Step	Action
1	Open the **Alarms** screen by selecting **Serviceability > RMON > Alarms** from the menu.
2	Select the **Events** tab. This tab is illustrated below.

RmonAlarms screen - Events tab

The following table describes the **Events** tab fields.

Events tab fields

Field	Description
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 Device Manager 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.
LastTimeSent	The value of sysUpTime at the time this event entry last generated an event. If this entry has not generated any events, this value is zero.
Owner	If traps are specified to be sent to the owner, then this is the name of the machine that will receive alarm traps.
Status	Normally valid. A not-valid field indicates that an SNMP agent other than the Device Manager has tried to modify an RMON parameter or that network conditions have corrupted an SNMP packet sent by the Device Manager. The status would temporarily appear as "under creation" and then the status would become either "valid" or the field would be deleted.

—End—

See also

• "Alarms" (page 107)

• "How RMON alarms work" (page 107)

• "Creating alarms" (page 109)

• "Viewing an event" (page 115)

• "Log information" (page 118)

Creating an event

To create an event:

Step	Action
1	Open the **Alarms** screen by selecting **Serviceability > RMON > Alarms** from the menu. This screen is illustrated below.
2	Select the **Events** tab.
3	Click **Insert**. The **Insert Events** screen opens. This screen is illustrated below.

Insert Events dialog box

Step	Action
4	In the **Description** field, type a name for the event.
5	Select the type of event in the **Type** field.
6	Enter the community information in the **Community** field.
7	Enter the owner information in the **Owner** field.
8	Click **Insert**.

—End—

Deleting an event

To delete an event:

Step	Action
1	Open the **Alarms** screen by selecting **Serviceability > RMON > Alarms** from the menu.
2	Select the **Events** tab.
3	Select an event from the list.
4	Click **Delete**.

—End—

Log information

The **Log** tab chronicles and describes the alarm activity.

To view the **Log** tab follow this procedure:

Step	Action
1	Open the **Alarms** screen by selecting **Serviceability > RMON > Alarms** from the menu.
2	Select the **Log** tab. This tab is illustrated below.

Log tab

The following table describes the **Log** tab fields.

Log tab fields

Item	Description
Time	Specifies when an event occurred that activated the log entry.
Description	Specifies whether the event is a rising or falling event.

—End—

Chapter 4
IP Flow Information Export (IPFIX)

IP Flow Information Export (IPFIX) is a protocol used for the export of flow information from traffic observed on a switch. Since IPFIX is still in development with the IETF, the current implementation is based on Netflow V9.

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

* protocol type

* destination IP address

* source IP address.

* ingress port

* TOS

   Note: You can't use IPFIX on secondary interfaces.

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

* source port

* destination port

Software Release 5.0 and up supports IPFIX through the creation and display of sampled information as well as the ability to export this sampled information. IPFIX functionality can be accessed through the Java Device Manager or Web-based Management Interface.

   Note: The IPFIX feature shares resources with QoS. If the IPFIX feature is enabled, a QoS policy precedence will be used. For further information about QoS policies, refer to the *Nortel Ethernet Routing Switch 5500 Series Configuration - Quality of Service* (Part Number NN47200-504).

IPFIX configuration using the Java Device Manager

This section describes the configuration and management of IPFIX functionality using the Java Device Manager.

Global IPFIX configuration

IPFIX functionality can be globally enabled or disabled from the Java Device Manager. By default, IPFIX is disabled and must be enabled before it will start to collect flow information. This section contains the procedures for enabling and disabling IPFIX on a switch.

Global configuration using the JDM

To enable or disable IPFIX using the JDM, follow this procedure:

Step	Action
1	Select **Serviceability > IPFIX** from the Device Manager menu. The **IPFIX** dialog opens with the **Global** tab selected. This screen is illustrated below.

IPFIX dialog - Global tab

2	On the **Global** tab, select the operational state of IPFIX functionality from the **State** area.
3	Click **Apply**.

—End—

Configuring IPFIX flows

Once IPFIX has been enabled on a switch, the ports IPFIX will monitor must be configured. Configuration of flow information sources can be performed in the Java Device Manager.

Configuring flows using the JDM

Flow configuration using the JDM is performed on the **Exporters** and **Ports** tab of the **IPFIX** dialog.

To configure IPFIX flows in the JDM, perform the following procedure:

Step	Action

1   Select **Serviceability > IPFIX** from the Device Manager menu. The **IPFIX** dialog will open with the **Global** tab selected.

2   Select the **Exporters** tab. This tab is illustrated below.

IPFIX dialog - Exporters tab

3   The **Exporters** tab lists the IPFIX exporters that are currently available. If connected to a standalone unit, the export properties of that unit are listed. If connected to a stack, the export properties of all units in the stack are listed. Using the fields provided, set up the IPFIX export properties. These fields are explained in the table below.

Exporters tab fields

Field	Description
Slot(Unit)	The switch that is exporting IPFIX flows. This number corresponds to the unit number in a stack or is 1 for a standalone unit.
AgingIntv	The aging interval of the flow record in seconds. This is an integer value between 0 and 2147400.
ExportIntv	The frequency of data exports to the collector in seconds. This is an integer value between 10 and 3600.
ExportState	The current state of the exporter.
TempRefIntvSec	The template refresh time out in seconds. The template is sent out to the collector either at the interval specified in this value or after the number of packets specified in the **TempRefIntvPkts** value, whichever occurs first. This is an integer value between 300 and 3600.
TempRefIntvPkts	The template refresh time out in numbers of packets. The template is sent out to the collector either at the interval specified in this value or after the number of seconds specified in the **TempRefIntvSec** value, whichever occurs first. This is an integer value between 10000 and 100000.

4   Click **Apply**.

5   To continue with the export configuration process, continue with the
     port configuration procedure outlined below.

—End—

Configuring IPFIX collectors

IPFIX collectors are used to collect and analyze data exported from an
IPFIX-compliant switch. At this time, up to two collectors can be supported.

IPFIX data is exported from the switch in *Netflow version 9* format. Data
is exported using UDP port 9995.

> *Note:* IPFIX data is not load balanced when two collectors are in use.
> Identical information is sent to both collectors.

To configure an IPFIX collector, perform this procedure:

Step   Action

---

1   Select **Serviceability > IPFIX** from the Device Manager menu. The
     IPFIX dialog will open with the **Global** tab selected.

2   Select the **Collectors** tab. This tab is illustrated below.

IPFIX dialog - Collectors tab

3   To modify the configuration of a collector, use the fields provided on
     the tab. These fields are described in the table below.

Collectors tab fields

Field	Description
Slot(Unit)	The unit number of the collector. Currently up to two collectors are supported.
AddressType	The address type of the IP address of the collector. Currently only IPv4 addresses are supported.
Address	The IP address of the collector.

Protocol	The protocol used to transport the IPFIX data to the collector. Currently only the UDP protocol is supported for this task.
DestPort	The port on which the collector will be listening for IPFIX data. Currently only port 9995 is supported for this task.
ExporterIpType	The address type of the IP address of the IPFIX exporter. Currently only IPv4 addresses are supported.
ExporterIp	The IP address of the IPFIX exporter.
ProtoVer	The format in which IPFIX data is provided to the collector. Currently only Netflow version 9 formatting is supported for this task.
Enable	The operational state of this collector.

4   To create a new collector, click **Insert**.

5   The **Insert Collectors** dialog opens. This dialog is illustrated below.

Insert Collectors dialog

6   Using the fields provided on the **Insert Collectors** dialog, configure the new collector. These fields are described in the table below.

Insert Collectors fields

Field	Description
Slot(Unit)	The unit number of the collector. Currently up to two collectors are supported so the values 1 or 2 are valid.
AddressType	The address type of the IP address of the collector. Currently only IPv4 addresses are supported.
Address	The IP address of the collector.

Protocol	The protocol used to transport the IPFIX data to the collector. Currently only the UDP protocol is supported for this task.
DestPort	The port on which the collector will be listening for IPFIX data. Currently only port 9995 is supported for this task.
ProtoVer	The format in which IPFIX data is provided to the collector. Currently only Netflow version 9 formatting is supported for this task.
Enable	The operational state of this collector.

—End—

Configuring IPFIX ports

Use the **Ports** tab to configure port settings for IPFIX data export. To configure IPFIX ports, use the following procedure:

Step	Action

1   Select **Serviceability > IPFIX** from the Device Manager menu. The **IPFIX** dialog will open with the **Global** tab selected.

2   Select the **Ports** tab. This tab is used to configure the individual ports on the exporting units. This tab is illustrated below.

IPFIX dialog - Ports tab
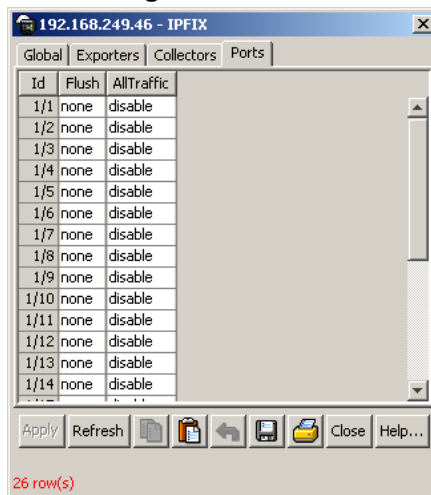
3   Using the fields provided, configure the IPFIX parameters for the individual ports. The fields on this tab are outlined in the table below.

Ports tab fields

Field	Description
Id	The individual port on which the IPFIX parameters are being configured. Ports are itemized in the format *Unit / Port*.
Flush	***Note:*** Although this field is displayed on a per port basis, flushing is only supported on a per unit basis in Software Release 5.0. Determines the flushing action to take on the port. Flushing the port of data involves deleting all previously gathered information on that port. This field provides three options: • none - The port data is not flushed. • flush - The port data is flushed; deleting it from switch memory. • exportAndFlush - The port data is exported to a configured collector and the data is then flushed.
AllTraffic	Determines whether IPFIX data is collected on this port. This field provides two options: • enable - IPFIX data is collected. • disable - IPFIX data is not collected.

If a single port is selected, packets are sampled every second. If multiple ports are selected, sampling is performed on every port that has a link in succession. Sampling rotates between the selected ports with each port having a sampling window of 1 second. For example, if 10 ports were selected on a switch, each port would be sampled every 10 seconds.

4   Click **Apply**.

—End—

Graphing Exporter Statistics

To view IPFIX exporter statistics, use the following procedure:

Step	Action

1   Select **Serviceability > IPFIX** from the Device Manager menu. The **IPFIX** dialog will open with the **Global** tab selected.

2   Select the **Collectors** tab.

3   On the **Collectors** tab, select an entry and click **Graph**. The **IPFIX
Exporter Stats** screen opens with the **Exporter** tab selected. This
tab is illustrated below.

IPFIX Exporter Stats screen - Exporter tab

	AbsoluteValue	Cumulative	Average/sec	Minimum/sec	Maximum/sec	LastVal/sec
OutPkts	0	0	0	0	0	0
OutOctets	0	0	0	0	0	0
PktsLoss	0	0	0	0	0	0

192.168.249.46 - IPFIX Exporter Stats: 1.ipv4.192.168.249.46

Exporter | Clear Time

Reset Counters | Clear Counters | Close | Help... | Poll Interval: 10s | 0 day, 00h:00m:32s

4   The following table outlines the fields on this tab.

Exporter tab fields

Field	Description
OutPkts	Indicates the total number of packets sent.
OutOctets	Indicates the total number of bytes sent.
PktsLoss	Indicates the total number of records lost.

—End—

Exporter Stats Clear Time

In conjunction with the **Exporters** tab, the **Clear Time** tab indicates the
system time when exporter statistics were last cleared (**none** if this has
never occurred). This tab is illustrated below.

IPFIX Exporter Stats screen - Clear Time tab

192.168.249.46 - IPFIX Exporter S...

Exporter | Clear Time

ClearTime: none

Refresh | Close | Help...

IPFIX Configuration using the Command Line Interface

This section describes the commands used in the configuration and
management of IPFIX using the CLI.

ip ipfix collector command

The **ip ipfix collector** command is used to configure IPFIX collectors. IPFIX collectors are used to collect and analyze data exported from an IPFIX compliant switch. In Software Release 5.0, the only external collector supported is **NetQOS**. At this time, up to two collectors can be supported.

IPFIX data is exported from the switch in *Netflow version 9* format. Data is exported using UDP port 9995.

> *Note:* IPFIX data is not load balanced when two collectors are in use. Identical information is sent to both collectors.

The syntax of the **ip ipfix collector** command is:

```
ip ipfix collector <unit_number> <collector_ip_address>
[dest-port <destination_port>] [enable] [exporter-ip
<exporter_ip_address>]
```

The following table describes the parameters for this command.

ip ipfix collector Parameters

Parameter	Description
<unit_number>	The unit number of the collector. Currently up to two collectors are supported so the values 1 or 2 are valid.
<collector_ip_address>	The IP address of the collector.
<destination_port>	The port on which the collector will be listening for IPFIX data. Currently only port 9995 is supported for this task.
<exporter_ip_address>	The IP address of the IPFIX exporter.

The **ip ipfix collector** command is executed in the Global Configuration mode.

ip ipfix enable command (Global Configuration)

The **ip ipfix enable** command is used to globally enable IPFIX on the switch.

The syntax of the **ip ipfix enable** command is:

```
ip ipfix enable
```

The **ip ipfix enable** command is executed in the Global Configuration mode.

ip ipfix slot command

The **ip ipfix slot** command is used to configure unit specific IPFIX parameters.

The syntax of the **ip ipfix slot** command is:

```
ip ipfix slot <unit_number> [aging-interval <aging_interval>]
[export-interval <export_interval>] [exporter-enable]
[template-refresh-interval <template_refresh_interval>]
[template-refresh-packets <template_refresh_packets>]
```

The parameters of this command are described in the following table.

ip ipfix slot Parameters

Parameter	Description
<unit_number>	The unit number of the collector. Currently up to two collectors are supported so the values 1 or 2 are valid.
<aging_interval>	The IPFIX aging interval. This is a value in seconds from 0 to 2147400.
<export_interval>	The IPFIX export interval. This is the interval at which IPFIX data is exported. This is a value in seconds from 10 to 3600.
<template_refresh_interval>	The IPFIX template refresh interval. This is a value in seconds from 300 to 3600.
<template_refresh_packets>	The IPFIX template refresh packet setting. This is a value in number of packets from 10000 - 100000.

The **ip ipfix slot** command is executed in the Global Configuration mode.

ip ipfix enable command (Interface Configuration)

The **ip ipfix enable** command is used to enable IPFIX on the interface.

The syntax of the **ip ipfix enable** command is:

```
ip ipfix enable
```

The **ip ipfix enable** command is executed in the Interface Configuration mode.

ip ipfix port command

The **ip ipfix port** command is used to enable the ports exporting data through IPFIX.

The syntax of the **ip ipfix port** command is:

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

The `<port_list>` parameter represents a single or comma-separated list of ports.

The `ip ipfix port` command is executed in the Interface Configuration mode.

ip ipfix flush command

The `ip ipfix flush` command is used to delete the collected IPFIX information for a port.

The syntax of the `ip ipfix flush` command is:

```
ip ipfix flush port <port_list> [export-and-flush]
```

The `<port_list>` parameter represents a single or comma-separated list of ports. The `export-and-flush` parameter is optional and is used to export data to a collector before it is deleted.

The `ip ipfix flush` command is executed in the Privileged EXEC mode.

show ip ipfix table command

The `show ip ipfix table` command is used to display IPFIX data collected from the switch.

The syntax of the `show ip ipfix table` command is:

```
show ip ipfix table <unit_number> sort-by <sort_by>
sort-order <sort_order> display <num_entries>
```

The following table outlines the parameters of this command:

show ip ipfix command parameters

Parameter	Description
<unit_number>	The unit number of the collector. Currently up to two collectors are supported so the values 1 or 2 are valid.

sort-by <sort_by>	The value on which the data is sorted. Valid options are: • byte-count • dest-addr • first-pkt-time • last-pkt-time • pkt-count • port • protocol • source-addr • TCP-UDP-dest-port • TCP-UDP-src-port • TOS
sort-order <sort_order>	The order in which the data is sorted. Valid options are ascending and descending.
display <num_entries>	The number of data rows to display. Valid options are: • all • top-10 • top-25 • top-50 • top-100 • top-200

The `show ip ipfix table` command is executed in the Privileged EXEC mode.

IPFIX configuration using the Web-based Management Interface

This section outlines the configuration and management of IPFIX functionality in the Web-based Management Interface.

Global configuration using the Web-based Management Interface

IPFIX functionality can be globally enabled or disabled from the Web-based Management Interface. By default, IPFIX is disabled and must be enabled before it will start to collect flow information. This section contains the procedures for enabling and disabling IPFIX on a switch.

To enable or disable IPFIX using the Web-based Management Interface, follow this procedure:

Step	Action

1   Select **Applications > IP Fix > IP Fix Configuration** from the
Web-based Management Interface navigation pane. The **IP Fix
Configuration** page opens. This page is illustrated below.

IP Fix Configuration Page

2   Select the operational state of the IPFIX functionality from the **IP Fix**
drop down list located in the **IP Fix Global Setting** area.

3   Click **Submit**.

—End—

Configuring flows using the Web-based Management Interface

Once IPFIX has been enabled on a switch, the ports IPFIX will monitor must
be configured. Configuration of flow information sources can be performed
in the Web-based Management Interface.

Flow configuration in the Web-based Management Interface is performed
on the **IP Fix Configuration** page.

To configure IPFIX flows using the Web-based Management Interface,
perform the following procedure:

Step	Action

1   Select **Applications > IP Fix Configuration** from the Web-based
Management Interface navigation pane. The **IP Fix Configuration**
page opens. This page is illustrated below.

IP Fix Configuration Page

Application > IP Fix Configuration

IP Fix Global Setting
IP Fix [Disabled ▼]

[Submit]

IP Fix Port Setting																													
	Aging Time		**Observation Ports**																										
Port		All	1 25	2 26	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24		
	86400 seconds ☐		☐ ☐	☐ ☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐		

[Submit]

2 Using the fields provided in the **IP Fix Port Setting** area, configure the IPFIX flow for individual ports. The fields in this area are described in the table below.

IP Fix Port Setting fields

Field	Description
Aging Time	The aging interval of the flow record in seconds.
Observation Ports	Each port is represented by a check box. Select or de-select the appropriate check boxes to enable or disable IPFIX data collection on that port. Select or de-select all ports using the **All** check box.

3 Click **Submit**.

—End—

Viewing IPFIX data

IPFIX data can be viewed using the Web-based Management Interface. This viewing mechanism is provided for administrators who do not, or do not wish to have, IPFIX collectors configured on the network. Using this interface, data can be sorted, filtered, and cleared entirely.

To view IPFIX data, perform the following procedure:

Step Action

1 Choose **Applications > IP Fix > IP Fix Information** from the Web-based Management Interface navigation pane. The **IP Fix Information** page opens. This page is illustrated below.

IP Fix Information Page

Application > IP Fix Information

IP Fix Information (View By)

Sort On	Source Address
Sort Order	Ascending
Entries To Display	Top 10
Clear Statistics	No

Submit

IP Fix Information Table

Source Address	Destination Address	TOS	Protocol	Module ID	Port	TCP/UDP Src Port	TCP/UDP Dst Port	Packet Count	Byte Count	First Packet Time	Last Packet Time

2   Using the fields provided in the **IP Fix Information (View By)** area, configure the viewing or clearing of the IPFIX data. These fields are described in the table below.

IP Fix Information (View By) fields

Field	Description
Sort On	The item of data to sort the IPFIX data on. IPFIX data can be sorted on any item that is gathered.
Sort Order	The order to apply to the sorted data.
Entries To Display	The number of entries to display.
Clear Statistics	Whether or not to clear the current statistics from memory.

3   Click **Submit**.

4   The IPFIX data with be filtered and sorted based on the selections.

—End—

Index

Symbols/Numerics

A

B

C

T

test
   stack loopback 73
Threshold Type field 112
TimeExcds field 39, 40
TimestampReps field 38, 40
Timestamps field 38, 40
Total Octets field 78
Type field 116

U

Undersized Packets field 78
UndersizePkts field 54, 99, 105
unit stats 31
Utilization field 105

V

Value field 112, 113
Variable field 111, 113

Nortel Ethernet Routing Switch 5500 Series

Configuration - System Monitoring

To provide feedback, or report a problem in this document, go to http://www.nortel.com/documentfeedback.

Sourced in Canada and the United States of America.