CONFIGURATION GUIDE

Extreme NetIron Monitoring Configuration Guide, 06.0.00g

Supporting NetIron OS 06.0.00g

53-1004197-09
29 June 2018

Contents

Preface

Document conventions

The document conventions describe text formatting conventions, command syntax conventions, and important notice formats used in Extreme technical documentation.

Notes, cautions, and warnings

Notes, cautions, and warning statements may be used in this document. They are listed in the order of increasing severity of potential hazards.

NOTE
A Note provides a tip, guidance, or advice, emphasizes important information, or provides a reference to related information.

ATTENTION
An Attention statement indicates a stronger note, for example, to alert you when traffic might be interrupted or the device might reboot.

CAUTION
A Caution statement alerts you to situations that can be potentially hazardous to you or cause damage to hardware, firmware, software, or data.

DANGER
A Danger statement indicates conditions or situations that can be potentially lethal or extremely hazardous to you. Safety labels are also attached directly to products to warn of these conditions or situations.

Text formatting conventions

Text formatting conventions such as boldface, italic, or Courier font may be used to highlight specific words or phrases.

Format	Description
bold text	Identifies command names.
	Identifies keywords and operands.
	Identifies the names of GUI elements.
	Identifies text to enter in the GUI.
italic text	Identifies emphasis.
	Identifies variables.
	Identifies document titles.
`Courier font`	Identifies CLI output.

Format	Description
	Identifies command syntax examples.

Command syntax conventions

Bold and italic text identify command syntax components. Delimiters and operators define groupings of parameters and their logical relationships.

Convention	Description
bold text	Identifies command names, keywords, and command options.
italic text	Identifies a variable.
[ ]	Syntax components displayed within square brackets are optional. Default responses to system prompts are enclosed in square brackets.
{ **x** | **y** | **z** }	A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options.
x | **y**	A vertical bar separates mutually exclusive elements.
< >	Nonprinting characters, for example, passwords, are enclosed in angle brackets.
...	Repeat the previous element, for example, *member*[*member*...].
\	Indicates a "soft" line break in command examples. If a backslash separates two lines of a command input, enter the entire command at the prompt without the backslash.

Extreme resources

Visit the Extreme website to locate related documentation for your product and additional Extreme resources.

White papers, data sheets, and the most recent versions of Extreme software and hardware manuals are available at www.extremenetworks.com. Product documentation for all supported releases is available to registered users at www.extremenetworks.com/support/documentation.

Document feedback

Quality is our first concern at Extreme, and we have made every effort to ensure the accuracy and completeness of this document. However, if you find an error or an omission, or you think that a topic needs further development, we want to hear from you.

You can provide feedback in two ways:

- Use our short online feedback form at http://www.extremenetworks.com/documentation-feedback-pdf/
- Email us at internalinfodev@extremenetworks.com

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

Contacting Extreme Technical Support

As an Extreme customer, you can contact Extreme Technical Support using one of the following methods: 24x7 online or by telephone. OEM customers should contact their OEM/solution provider.

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

- GTAC (Global Technical Assistance Center) for immediate support
  - Phone: 1-800-998-2408 (toll-free in U.S. and Canada) or +1 408-579-2826. For the support phone number in your country, visit: www.extremenetworks.com/support/contact.
  - Email: support@extremenetworks.com. To expedite your message, enter the product name or model number in the subject line.
- GTAC Knowledge - Get on-demand and tested resolutions from the GTAC Knowledgebase, or create a help case if you need more guidance.
- The Hub - A forum for Extreme customers to connect with one another, get questions answered, share ideas and feedback, and get problems solved. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.
- Support Portal - Manage cases, downloads, service contracts, product licensing, and training and certifications.

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

- Your Extreme Networks service contract number and/or serial numbers for all involved Extreme Networks products
- A description of the failure
- A description of any action(s) already taken to resolve the problem
- A description of your network environment (such as layout, cable type, other relevant environmental information)
- Network load at the time of trouble (if known)
- The device history (for example, if you have returned the device before, or if this is a recurring problem)
- Any related RMA (Return Material Authorization) numbers

About This Document

What's new in this document

On October 30, 2017, Extreme Networks, Inc. acquired the data center networking business from Brocade Communications Systems, Inc. This document has been updated to remove or replace references to Brocade Communications, Inc. with Extreme Networks., Inc., as appropriate.

Supported hardware and software

The hardware platforms in the following table are supported by this release of this guide.

TABLE 1 Supported devices

ExtremeRouting XMR Series	ExtremeRouting MLX Series	ExtremeSwitching CES 2000 Series	ExtremeRouting CER 2000 Series
XMR 4000	MLX-4	CES 2024C	CER 2024C
XMR 8000	MLX-8	CES 2024F	CER-RT 2024C
XMR 16000	MLX-16	CES 2048C	CER 2024F
XMR 32000	MLX-32	CES 2048CX	CER-RT 2024F
	MLXe-4	CES 2048F	CER 2048C
	MLXe-8	CES 2048FX	CER-RT 2048C
	MLXe-16		CER 2048CX
	MLXe-32		CER-RT 2048CX
			CER 2048F
			CER-RT 2048F
			CER 2048FX
			CER-RT 2048FX

Supported software

For the complete list of supported features and the summary of enhancements and configuration notes for this release, refer to the *Extreme NetIron Release Notes*.

How command information is presented in this guide

Starting with Extreme NetIron 5.6.00, command syntax and parameter descriptions are removed from commands that are referenced in configuration tasks. To find the full description of a specific command, including all required and optional keywords and variables, refer to the *Extreme NetIron Command Reference* for your software release.

Hardware Monitoring

Configuring optical monitoring

You can configure your Extreme device to monitor XFPs or SFPs in the system either globally or by specified port. If monitoring is enabled, console messages, syslog messages, and SNMP traps are sent when XFP or SFP operating conditions warrant it and a port is enabled.

Configure all XFP and SFP ports for optical monitoring, using the following command.

```
device(config)# optical-monitor
```

Configure a specific XFP or SFP port for optical monitoring, using the following command.

```
device(config)# interface ethernet 1/1
device(config-if-e10000-1/1)# optical-monitor
```

Configure a range of XFP or SFP ports for optical monitoring, using the following command.

```
device(config)# interface ethernet 1/1 to 1/2
device(config-mif-e10000-1/1-1/2)# optical-monitor
```

Syntax: [no] optical-monitor *alarm-interval*

The optional *alarm-interval* variable sets the interval in minutes between which alarms or messages are sent. The default interval is 3 minutes.

You can view the XFP optical monitoring information using the show optic command as displayed in the following.

```
device#show optic 4
 Port  Temperature    Tx Power      Rx Power      Tx Bias Current
+----+-----------+----------+------------+------------------+
4/1    30.8242 C   -001.8822 dBm -002.5908 dBm    41.790 mA
       Normal        Normal          Normal          Normal
4/2    31.7070 C   -001.4116 dBm -006.4092 dBm    41.976 mA
       Normal        Normal          Normal          Normal
4/3    30.1835 C                 -000.5794 dBm     0.000 mA
       Normal       Low-Alarm       Normal         Low-Alarm
4/4     0.0000 C                                   0.000 mA
       Normal        Normal          Normal          Normal
```

For Temperature, Tx Power, Rx Power, and Tx Bias Current, values are displayed along with one of the following status values: Low-Alarm, Low-Warn, Normal, High-Warn or High-Alarm. The thresholds that determine these status values are set by the manufacturer of the XFPs. Table 2 describes each of these status values.

TABLE 2 Status value description

Status value	Description
Low-alarm	The monitored level has dropped below the "low-alarm" threshold set by the XFP or SFP manufacturer.
Low-warn	The monitored level has dropped below the "low-warn" threshold set by the XFP or SFP manufacturer.
Normal	The monitored level is within the "normal" range set by the XFP or SFP manufacturer.
High-warn	The monitored level has climbed above the "high-warn" threshold set by the XFP or SFP manufacturer.
High-alarm	The monitored level has climbed above the "high-alarm" threshold set by the XFP or SFP manufacturer.

When the **show optic** command is issued on a BR-MLX-100GX interface card, the following conditions apply.

- The temperature is averaged over all lanes.
- TX bias, RX power and RX power are aggregate values.

NOTE

This function takes advantage of information stored and supplied by the SFP or XFP device. This information is an optional feature of the Multi-Source Agreement standard defining the SFP or XFP interface. Not all component suppliers have implemented this feature set. In such cases where the SFP or XFP device does not supply the information, a "Not Available" message will be displayed for the specific port that the device is installed

Displaying optical monitoring thresholds

To display information about the optical monitoring thresholds, enter the following command.

```
device#show optic threshold 3
Port 3/1
Transceiver Temperature High alarm     4600       70.0000 C
Transceiver Temperature High warning   4400       68.0000 C
Transceiver Temperature Low warning    0200        2.0000 C
Transceiver Temperature Low alarm      0000        0.0000 C
VCC Voltage High alarm                 875a        3.4650  mV
VCC Voltage High warning               8610        3.4320  mV
VCC Voltage Low warning                7bc0        3.1680  mV
VCC Voltage Low alarm                  7a76        3.1350  mV
SOA Bias Current High alarm            0000        0.000  mA
SOA Bias Current High warning          0000        0.000  mA
SOA Bias Current Low warning           0000        0.000  mA
SOA Bias Current Low alarm             0000        0.000  mA
Auxiliary 1 Monitor High alarm         0000
Auxiliary 1 Monitor High warning       0000
Auxiliary 1 Monitor Low warning        0000
Auxiliary 1 Monitor Low alarm          0000
Auxiliary 2 Monitor High alarm         0000
Auxiliary 2 Monitor High warning       0000
Auxiliary 2 Monitor Low warning        0000
Auxiliary 2 Monitor Low alarm          0000
Laser Bias Current High alarm          ea60      120.000  mA
Laser Bias Current High warning        e09c      115.000  mA
Laser Bias Current Low warning         445c       35.000  mA
Laser Bias Current Low alarm           3a98       30.000  mA
Laser TX Power High alarm              6e18      004.5000 dBm
Laser TX Power High warning            621f      004.0000 dBm
Laser TX Power Low warning             1049     -003.7996 dBm
Laser TX Power Low alarm               0e83     -004.3004 dBm
Laser Temperature High alarm           3700       55.0000 C
Laser Temperature High warning         3500       53.0000 C
Laser Temperature Low warning          1b00       27.0000 C
Laser Temperature Low alarm            1900       25.0000 C
Laser RX Power High alarm              6e18      004.5000 dBm
Laser RX Power High warning            621f      004.0000 dBm
Laser RX Power Low warning             01f5     -013.0016 dBm
Laser RX Power Low alarm               00fb     -016.0032 dBm
```

```
Port 3/2
Transceiver Temperature High alarm     4600       70.0000 C
Transceiver Temperature High warning   4400       68.0000 C
Transceiver Temperature Low warning    0200        2.0000 C
Transceiver Temperature Low alarm      0000        0.0000 C
VCC Voltage High alarm                 875a        3.4650  mV
VCC Voltage High warning               8610        3.4320  mV
VCC Voltage Low warning                7bc0        3.1680  mV
VCC Voltage Low alarm                  7a76        3.1350  mV
SOA Bias Current High alarm            0000        0.000  mA
SOA Bias Current High warning          0000        0.000  mA
SOA Bias Current Low warning           0000        0.000  mA
SOA Bias Current Low alarm             0000        0.000  mA
Auxiliary 1 Monitor High alarm         0000
Auxiliary 1 Monitor High warning       0000
Auxiliary 1 Monitor Low warning        0000
Auxiliary 1 Monitor Low alarm          0000
Auxiliary 2 Monitor High alarm         0000
Auxiliary 2 Monitor High warning       0000
Auxiliary 2 Monitor Low warning        0000
Auxiliary 2 Monitor Low alarm          0000
Laser Bias Current High alarm          ea60      120.000  mA
Laser Bias Current High warning        e09c      115.000  mA
Laser Bias Current Low warning         445c       35.000  mA
Laser Bias Current Low alarm           3a98       30.000  mA
Laser TX Power High alarm              6e18      004.5000 dBm
Laser TX Power High warning            621f      004.0000 dBm
Laser TX Power Low warning             1049     -003.7996 dBm
Laser TX Power Low alarm               0e83     -004.3004 dBm
Laser Temperature High alarm           3700       55.0000 C
Laser Temperature High warning         3500       53.0000 C
Laser Temperature Low warning          1b00       27.0000 C
Laser Temperature Low alarm            1900       25.0000 C
Laser RX Power High alarm              6e18      004.5000 dBm
Laser RX Power High warning            621f      004.0000 dBm
Laser RX Power Low warning             01f5     -013.0016 dBm
Laser RX Power Low alarm               00fb     -016.0032 dBm
Show optics thresholds done
device#
```

The example above displays information about the optical monitoring thresholds.

Syntax: show optic thresholds slot-number

Displaying media information

To display media information for SFP and XFP devices installed in a specific slot, enter the following command at any CLI level.

```
device#show media slot 3
Port 3/1:
    Type  : 10GBASE-ER/EW 1547.50nm (XFP)
    Vendor:     BOOKHAM         , Version:                  01
    Part# :     IGF-32511J      , Serial#:     BTH0622357
Port 3/2:
    Type  : 10GBASE-LR/LW 1310.00nm (XFP)
    Vendor:     foundry networks, Version:                  00
    Part# :     FTRX-1411E3     , Serial#:     K68034S
Port 3/3:
    Type  : 10GBASE-ER/EW 1547.50nm (XFP)
    Vendor:     BOOKHAM         , Version:                  01
    Part# :     IGF-32511J      , Serial#:     BTH0622410
Port 3/4:
    Type  : 10GBASE-SR/SW 854.00nm (XFP)
    Vendor:     Foundry Networks, Version:                  02
    Part# :     JXPR01SW05306   , Serial#:     F74340380372
```

The example above displays all optical devices on slot 3.

Syntax: show media slot slot-number

To display media information for SFP and XFP devices installed in an ethernet port, enter the following command at any CLI level.

```
device#show media ethernet 3/4
Port 3/4:
    Type  : 10GBASE-SR/SW 854.00nm (XFP)
    Vendor:     Foundry Networks, Version:                   02
    Part# :     JXPR01SW05306   , Serial#:     F74340380372
```

Syntax: show media [**ethernet slot-port** [**to slot-port**]]

You can display media information for all ports in an Extreme device by using the **show media** command without options.

The **ethernet** *slot-port* parameter limits the display to a single port.

The **to** *slot-port* parameter displays information for a range of ports.

This results displayed from this command provide the Type, Vendor, Part number, Version and Serial number of the SFP or XFP optical device installed in the port.

If no SFP or XFP device is installed in a port, the "Type" field will display "N/A", the "Vendor" field will be empty and the other fields will display "Unknown".

Multi-rate optical transceivers are supported. In this case, if a multi-rate optical transceiver is inserted in an Interface module, the "Type" parameter will display the transmission code for the correct value for the port as determined by either the Interface module type or the configuration of the port. There is one exception to this rule however. If a port is in the disabled state only one type will be displayed. Once the port is enabled, the correct "Type" will be displayed in accordance with the configuration.

Optics compatibility checking

This feature checks the installation of the following optical transceivers into Interface module ports and shuts down the port if the transceiver is incompatible with the port:

- 10 GbE XFP - This interface is brought up if the XFP is compliant with Ethernet transmission compliance.
- 1 Gb (100/1000) Ethernet interface will be enabled if the SFP is Ethernet capable

If the interface is incompatible with the optical transceiver installed, the port will not come up and the syslog message "**Incompatible optical trans-receiver detected on port n** " is displayed. An SNMP trap is also generated and the port is described as "down" because of "(incompatible transceiver)" in the output from the **show interface** command.

Multi-rate optical transceivers (XFP and SFP) are supported as described in the following:

- In Multi-rate SFPs and XFPs, the EEPROM is programmed for multi-rate - for example both Ethernet 1 Gb and SONET compliance codes can be programmed in the internal EEPROM of a multi-rate optical transceiver.
- Multi-rate SFPs and XFPs are supported. The system software checks for transmission compatibility against the interface configuration. Therefore an OC-12 interface will be brought up if the SFP is compatible for both OC-12 SONET and 1 Gb Ethernet transmission. The same SFP can also be used in a 1 Gb Ethernet interface.
- The **show media** command described in Configuring optical monitoring on page 17 continues to show only one transmission rate even for multi-rate SFPs and XFPs. If the interface is enabled and the SFP or XFP is compatible, the **show media** command only displays the compatible transmission code in the "Type" field. If the interface is disabled, the **show media** display depends on the module type. For Ethernet interface modules, the Ethernet compliance code is shown. If the Ethernet compliance code is not set then the SONET compliance code is displayed.

Disabling transceiver type checking

When transceiver type checking is disabled, the syslog message "**Incompatible optical trans-receiver detected on port n** " is still displayed but the port is not shut down. You can disable transceiver type checking with the **no transceiver-type-check** command as shown in the following.

```
device(config)# no transceiver-type-check
```

Syntax: [no] transceiver-type-check

Transceiver type checking is on by default and the command is not included in the configuration.

The **no** option of the **transceiver-type-check** command, disables transceiver type checking as described, sends a syslog message and places the command in the configuration.

Using the **transceiver-type-check** command without the no option, enables transceiver type checking, sends a syslog message and removes the command from the configuration.

Monitoring dynamic memory allocation

After a configured system-max value is accepted, it is possible that the dynamic memory allocation may fail in a running system. To monitor the amount of available memory on the Management Module and the Interface Module, a timer will check the memory every 10 seconds. If the available memory falls below 5 percent of the total installed memory, the timer will log the following warning message.

```
device# show log
...
Jan 17 22:55:55:N: WARN: Current Total Free Memory on MP is below 5 percent of Installed Memory.
...
Jan 17 23:53:55:N: WARN: Current Total Free Memory on LP 8 is below 5 percent of Installed Memory.
```

The warning message is displayed at a frequency of 1 log per 5 minutes.

> NOTE
> Notifications and traps are
> sent.

When the memory allocation fails, an alert message is logged immediately. The alert message is displayed at a frequency of 1 log per 5 minutes. The following example below displays an alert message on the Management Module and the Interface Module.

```
device# show log
...
Jan 17 22:55:55:A: ALERT: Failed to allocate memory on MP
...
Jan 17 23:52:55:A: ALERT: Failed to allocate memory on LP 8
...
```

The NULL value is returned to the calling routine. The calling routine will decide how to proceed after the memory allocation fails.

> NOTE
> Notifications and traps are
> sent.

At any time, you can display the status of all recorded memory that is available on the Management Module by entering the **show memory** command. The amount of available memory is displayed in percentage values.The following example displays a show memory output on a Management Module.

```
device#show memory
==================================================================
```

```
NetIron XMR active MP slot 33:
Total SDRAM         :  2147483648 bytes
Available Memory    :  1774059520 bytes
Available Memory (%):  82 percent
Free Physical Pages :  428503 pages
<...>
================================================================
NetIron XMR LP SL 2:
Total SDRAM         :   536870912 bytes
Available Memory    :   45821952 bytes
Available Memory (%):   8 percent
```

Switch fabric fault monitoring

With this feature, you can display information about the current status of links between the switch fabric modules (SFM) and interface modules in a XMR Series or MLX Serieschassis. This feature also provides log messages to the console when there is a change in the "UP" or "DOWN" status of links to the SFM and when an individual fabric element (FE) cannot be accessed by the management module. The device can also be configured to automatically shut down an SFM when failure is detected. The following sections describe the capabilities of this feature.

Displaying switch fabric information

You can display information about the current status of links between the SFMs and interface modules in a XMR Series or MLX Series chassis using the following command. Each line represents a link between an SFM and an interface module (LP).

```
device#show sfm-links all
 SFM#/FE# | FE link#   | LP#/TM# | TM link#   | link state
----------+-----------+---------+-----------+---------
  2 / 1   |     32     | 3 / 1   |     13     |      UP
  2 / 1   |     31     | 3 / 2   |     01     |      UP
  2 / 1   |     11     | 3 / 1   |     01     |      UP
  2 / 1   |     12     | 3 / 2   |     13     |      UP
  2 / 3   |     32     | 3 / 1   |     19     |      UP
  2 / 3   |     31     | 3 / 2   |     07     |      UP
  2 / 3   |     11     | 3 / 1   |     07     |      UP
  2 / 3   |     12     | 3 / 2   |     19     |      UP
  3 / 1   |     32     | 3 / 1   |     16     |      UP
  3 / 1   |     31     | 3 / 2   |     04     |      UP
  3 / 1   |     11     | 3 / 1   |     04     |      UP
  3 / 1   |     12     | 3 / 2   |     16     |      UP
  3 / 3   |     32     | 3 / 1   |     22     |      UP
  3 / 3   |     31     | 3 / 2   |     10     |      UP
  3 / 3   |     11     | 3 / 1   |     10     |      UP
  3 / 3   |     12     | 3 / 2   |     22     |      UP
WARN: LP 3 has 8 links up, less than minimum to guarantee line rate traffic   forwarding
```

Syntax: show sfm-links sfm-number | **all** [**errors**]

The *sfm-number* variable specifies an SFM that you want to display link information for.

The **all** option displays link information for all SFMs in the chassis.

The *errors* option only displays information for SFM links that are in the DOWN state.

The output of this command can also be filtered using an output modifier. To use an output modifier, type a vertical bar (|) followed by a space and one of the following parameters:

- *begin* - begin output with the first matching line
- *exclude* - exclude matching lines from the output
- *include* - include only matching lines in the output

A warning statement is sent if the number of operational links falls below the minimum threshold. This warning is displayed to warn users that the line rate traffic will not be maintained.

The **show sfm-links** command displays the following information.

TABLE 3 CLI display of SFM link information

This field...	Displays...
SFM#	The switch fabric module number.
FE#	The FE number.
FE link#	The number of the interconnect between the SFM and the FE.
LP#	The slot number where the Interface module (LP) is installed.
TM#	The number of the traffic manager used in the link.
TM link#	The link number on the traffic manager.
link state	The link state is either: UP - In an operating condition DOWN - In a non-operational condition

Displaying switch fabric module information

To display the state of all switch fabric modules in the chassis, enter the following command at any level of the CLI.

```
device> show module
M1 (upper): NI-MLX-MR Management Module Active
M2 (lower): NI-MLX-MR Management Module Standby (Ready State)
F1: NI-X-SF Switch Fabric Module Powered off (By Health Monitoring)
F2:
F3:
F4: NI-X-HSF Switch Fabric Module Active
...
```

Syntax: show module

The **show module** command displays the modules currently connected to the chassis and their state. For switch fabric modules, the command shows "Active" if the module is operational or "Powered off" and the reason for the shutdown.

Powering a switch fabric link on or off manually

To manually power on a switch fabric link, use a command such as the following.

```
device# power-on snm-link 3 3 37
```

To manually power off a switch fabric link, use a command such as the following.

```
device# power-off snm-link 3 3 37
```

Syntax: [no] power-on snm-link sfm-number fe-number link-number

Syntax: [no] power-off snm-link sfm-number fe-number link-number

Powering a switch fabric module off automatically on failure

To configure the device to automatically power off a switch fabric module (SFM) or high speed switch fabric module (hSFM) on which an access error has been detected, enter the following command at the CONFIG level of the CLI.

```
device(config)# system-init fabric-failure-detection
```

Syntax: [no] system-init fabric-failure-detection

> NOTE
> You must restart the device for automatic SFM shutdown to take effect.

Once you have configured automatic SFM shutdown on the device and restarted it, the management module will automatically detect access failure (see Access failure messages on page 25) and shut down the unresponsive SFM. You can restart the SFM at any time (manually, by removing and re-inserting the module, or by initiating a system restart), but if another access error is detected, the management module will shut the SFM down again. If an SFM is automatically powered down, SFM power-off status (and the associated reason) are synced to the standby management module, and in the event of failover the standby module will keep the faulty SFM powered off.

Fabric link balancing

On Extreme NetIron MLXe-16 and MLXe-32, when one of the fabric link connected between an egress LP and an SFM/FE goes down, and if the ingress traffic is entering to the same SFM/FE through a different fabric link, there might be congestion on SFM/FE during the line rate traffic because of this link imbalance.

Supported Line Cards

The fabric link balancing feature is supported for Line cards BR-MLX-10GX20-X2, BRMLX-100GX2-CFP2-M, and BR-MLX-100GX2-CFP2-X2 on NetIron MLXe-16 chassis and BR-MLX-10GX24-DM line card on both MLXe-16 and MLXe-32 chassis.

Switch fabric log messages

Information about the state of each switch fabric module and whether it can be accessed by the Management Module is also provided in the form of syslog messages.

Link up/down messages

The Switch Fabric modules (SFM) in a Extreme chassis send a log message when they first become operational or when they change state between "UP" and "DOWN". The following is an example of the message sent when a link first becomes operational (UP) or when it changes state from non-operational (DOWN) to operational (UP).

```
Apr  6 10:57:20:E: Fabric Monitoring Link Up : SFM 3/FE 3/Link 37, LP 5/TM 1
```

The following is an example of the message sent when a link is detected going from operational (UP) to non-operational (DOWN).

```
Apr  6 10:56:00:E: Fabric Monitoring Link Down : SFM 3/FE 3/Link 37, LP 5/TM 1
```

Once a link has been detected as going down and "auto-tune" is disabled or "auto-tune" is enabled but the link has already been tuned (see Auto-tune enhancement on page 39), it is automatically shut down by the Multi-Service IronWare software . The following is an example of the message sent when a link is either brought down automatically or manually using the command described in Powering a switch fabric link on or off manually on page 23.

```
Apr  6 10:56:00:E: Fabric Monitoring Link Admin Shut Down : SFM 3/FE 3/Link 37, LP 5/TM 1
```

This contents of the message are defined as described in the following.

`Apr 6 10:57:20:` - The time that the link changed state.

`Fabric Monitoring Link Up` - the link went "UP"`Fabric Monitoring Link Down` - the link went "DOWN"

`SFM 3` - The switch fabric module (SFM) number

`FE 3` - The Fabric Element number

`Link 37` - The number of the interconnect between the SFM and the FE

`LP 5` - The slot number where the Interface Module (LP) is installed.

`TM 1` – The number of the traffic manager (TM) used in the link.

Access failure messages

The management module attempts to access each fabric element for every poll period (1 second by default). If the number of access failures in a poll window (default 10 seconds) exceeds the threshold (3 by default), the management module sends a log message similar to the following:

```
Apr 6 20:33:57:A:System: Health Monitoring: FE access failure detected on SFM 2/FE 1
```

The contents of the message are defined as described in the following.

Apr 6 20:33:57: - the time at which the error threshold was exceeded

`FE access failure detected` - the management module failed to access the specified FE

`SFM 2` - the switch fabric module (SFM) number

`FE 1` - the Fabric Element (FE) number

If the device has been configured to shut down a switch fabric module when failure is detected (see Powering a switch fabric module off automatically on failure on page 24), the management module will shut down the failed switch fabric module, then send a log message similar to the following:

```
Oct 4 20:33:57:A:System: Health Monitoring: Switch fabric 2 powered off due to failure detection
```

The message above indicates that a failure was detected in attempting to access switch fabric module 2, and the module was powered off on October 4th at 20:33:57.

Fabric error interrupt

Fabric error interrupts are monitored and logged for both the switch fabric module, and the Extreme MLX 24-port 10 GbE module and Extreme MLX 2-port 100 GbE module.

The following example shows a fabric error interrupt for a switch fabric module on an MP:

```
Dec 4 20:33:57: SFM 1 / FE 1 Reg offset 0x00000800 value 0x0000000c Overflow ( DCQ) Interrupt
```

The following example shows a fabric error interrupt for a slot on an LP:

```
Mar 4 20:33:57: Slot 17 FE1 Reg offset 0x00000800 value 0x0000000c Overflow ( DCQ) Interrupt
```

Link pair down messages

The Switch Fabric modules (SFM) in a chassis send a log message when a link pair is brought down by software. The following is an example of the message sent when a link pair is brought down.

```
Fabric Link Pair shut down for link balancing: LP 14/TM 1/Link 86 -> SFM 1/FE 2/Link 17
```

The following is an example of the message sent when a link pair going from operational (UP) to non-operational (DOWN).

```
Jun 13 05:35:18: Fabric Monitoring Link Pair shut down for link balancing SFM 2/FE 2/Link 73.
Jun 13 05:32:36:I: Fabric Monitoring Link Admin Shut Down : SFM 2/FE 2/Link 75
```

he following syslog message appears in **show log** command output when the link pair is brought down by the software.

```
Nov 14 06:45:46:I:System: Health Monitoring: Fabric Link Pair shut down for link balancing: LP 10/FE 1/link
52 -> SFM 1/FE 1/ Link 92
Nov 14 06:45:46:I:System: Health Monitoring: Fabric Link Pair shut down for link balancing: LP 10/FE 1/link
51 -> SFM 1/FE 1/ Link 91
Nov 14 06:45:46:I:System: Health Monitoring: Fabric Link Pair shut down for link balancing: LP 10/FE 1/link
88 -> SFM 1/FE 1/ Link 88
Nov 14 06:45:46:I:System: Health Monitoring: Fabric Link shut down due to CRC errors: LP 10/FE 1/link 87 ->
SFM 1/FE 1/ Link 87
```

The following SFM log message appears in the **show sfm logging** command output when the link pair is brought down by the software.

```
Jun 13 05:35:18: Fabric Monitoring Link Pair shut down for link balancing SFM 2/FE 2/Link 73.
Jun 13 05:35:18: Fabric Monitoring Link Admin Shut Down SFM 2/FE 2/Link 74.
```

The following TM log message appears in the **show tm logging** command output when the link pair is brought down by the software.

```
May 15 06:52:01: TM Link Pair Shutdown for balancing: SFM 1/FE 2/Link 62 -> LP 6/TM 1/Link 22
May 15 06:52:01: TM Link Shutdown due to CRC Errors: SFM 3/FE 1/Link 64 -> LP 6/TM 1/Link 16
```

This contents of the message are defined as described in the following.

`Jun 13 05:35:18:` - The time that the link changed state.

`Fabric Link Pair shut down for link balancing` - - the link pair went "DOWN"

`SFM 1` - The switch fabric module (SFM) number

`FE 2` - The Fabric Element number

`Link 17` - The number of the interconnect between the SFM and the FE

`LP 10` - The slot number where the Interface Module (LP) is installed.

`TM 1` – The number of the traffic manager (TM) used in the link.

Switch fabric utilization monitoring

With this feature, you can monitor the percentage of the total bandwidth used on the SFM for the timing intervals of 1 sec, 5 sec, 1 min, and 5 min. For example, to display bandwidth usage on all SFMs on the device, enter the following command.

```
device#show sfm-utilization all
SFM#2
----------+-----------+---------+-----------+---------
last 1 second utilization = 0.4%
last 5 seconds utilization = 0.3%
last 1 minute utilization = 0.1%
last 5 minutes utilization = 0.0%
SFM#3
----------+-----------+---------+-----------+---------
last 1 second utilization = 0.4%
last 5 seconds utilization = 0.4%
```

```
last 1 minute utilization = 0.1%
last 5 minutes utilization = 0.0%
```

To display bandwidth usage on one SFM, enter the following command.

```
device#show sfm-utilization 2
SFM#2
----------+----------+---------+----------+---------
last 1 second  utilization = 0.4%
last 5 seconds utilization = 0.3%
last 1 minute  utilization = 0.1%
last 5 minutes utilization = 0.0%
```

Syntax: show sfm-utilization [**all** | *sfm-number*]

The *sfm-number* variable specifies an SFM that you want to utilization information for.

The **all** option displays utilization information for all SFMs in the chassis.

Link fault signaling

You can enable link fault signaling on 10 or 100 gigabit interfaces. Link fault signaling (LFS) is a physical layer protocol that enables communication on a link between two 10 or 100 Gigabit Ethernet devices. When configured on the Extreme 10 or 100 Gigabit Ethernet port, the port can detect and report fault conditions on transmit and receive ports.

If LFS is configured on an interface, the following Syslog messages are generated when that interface goes up or down or when the TX or RX fiber is removed from one or both sides of the link that has LFS configured:

- SYSTEM: port 2/1 is down (remote fault)
- SYSTEM: Interface ethernet 2/1, state down - remote fault
- SYSTEM: Interface ethernet 2/1, state up

Traditionally, in MLX Series and XMR Series devices, LFS was disabled in both TX and RX directions. The **link-fault-signaling** command was used to enable LFS in both TX and RX directions. When RX LFS is enabled, a port will be brought up only when the PHY-MAC link is up, and there is no link fault received by the MAC. When RX LFS is disabled, a port will be brought up as long as the PHY-MAC link is up, regardless of any RX fault indication to MAC.

The RX LFS is always enabled by default and cannot be disabled. The **link-fault-signaling** command only applies to enabling or disabling the TX LFS. While RX LFS is recommended to be enabled at all times, for some applications it is requested to have the means to disable RX LFS.

There are two independent link-fault signaling commands **link-fault-signaling** and **link-fault-signaling ignore-rx.** These commands are applicable at both the global (system-level) and per-port level. Both global and per-port configurations are considered jointly to determine the resulting per-port configuration. When a global configuration is applied, it will override the corresponding per-port configuration already present. It is recommended to configure the global configuration prior to applying per-port configurations.

To configure LFS, enter the following commands.

```
device(config)# interface ethernet 1/4
device(config-if-e1000-1/4)# link-fault-signaling
```

Syntax: [no] link-fault-signaling

LFS is disabled by default.

> NOTE
> Ensure both sides are LFS ON when using LFS with RX (always on) and another router (which can be configured ON or OFF). Do not assume all boxes have LFS ON or OFF by default. Be sure and check.

To to disable RX LFS on a specified port, enter the **link-fault-signaling ignore-rx** command.

```
device(config)# interface ethernet 1/4
device(config-if-e1000-1/4)# link-fault-signaling ignore-rx
```

Syntax: [no] link-fault-signaling ignore-rx

RX LFS is ignored on the specified port.

Configuration examples

The following configuration examples show global and port configurations.

```
device(config)# link-fault-signaling
Extreme(config)#show run
Current configuration:
!
ver V5.4.0iT163
module 1 ni-mlx-8-port-10g-m
module 3 ni-mlx-8-port-10g-m
!
link-fault-signaling
!3  3ffff(R) 0.0.0.0/0          N/A            Dis N/A  Drop   00094
```

TX LFS and RX LFS are enabled on all ports.

```
device(config)# interface e 3/1
Extreme(config-if-e100000-3/1)#link-fault-signaling ignore-rx
Extreme(config-if-e100000-3/1)#show run
Current configuration:
!
ver V5.4.0iT163
module 1 ni-mlx-8-port-10g-m
module 3 ni-mlx-8-port-10g-m
!
link-fault-signaling
!
interface ethernet 3/1
link-fault-signaling ignore-rx
!
```

TX LFS is enabled on all ports. RX LFS is enabled on all ports except 3/1.

Port configuration overwritten by global configuration.

```
device(config)# show run
Current configuration:
!
ver V5.3.0pT183
!
no spanning-tree
!
vlan 1 name DEFAULT-VLAN
!
hostname Extreme
!
end
device(config)# show link-fault-signaling
Global Link Fault : RX ON   TX OFF
PORT   #:  LINK FAULT:
PORT 2/1:  RX ON   TX OFF
PORT 2/2:  RX ON   TX OFF
```

TX LFS is disabled on all ports and RX LFS is enabled on all ports.

```
device(config)# link-fault-signaling
device(config)# show run
```

```
Current configuration:
!
ver V5.3.0pT183
!
no spanning-tree
!
vlan 1 name DEFAULT-VLAN
!
hostname Extreme
link-fault-signaling
!
end
device(config)# show link-fault-signaling
Global Link Fault : RX ON   TX ON
PORT   #:  LINK FAULT:
PORT 2/1:  RX ON   TX ON
PORT 2/2:  RX ON   TX ON
```

TX LFS is enabled on all ports and RX LFS is enabled on all ports.

```
device(config)# int e 2/1
device(config-if-e10000-2/1)# no link-fault-signaling
device(config-if-e10000-2/1)# show run
Current configuration:
!
ver V5.3.0pT183
!
no spanning-tree
!
vlan 1 name DEFAULT-VLAN
!
hostname Extreme
link-fault-signaling
!
interface ethernet 2/1
 no link-fault-signaling
!
end
device(config-if-e10000-2/1)# show link-fault-signaling
Global Link Fault : RX ON   TX ON
PORT   #:  LINK FAULT:
PORT 2/1:  RX ON   TX OFF
PORT 2/2:  RX ON   TX ON
```

TX LFS is enabled on all ports except 2/1 and RX LFS is enabled on all ports.

```
device(config-if-e10000-2/1)# exit
device(config)# link-fault-signaling
device(config)# show run
Current configuration:
!
ver V5.3.0pT183
!
no spanning-tree
!
vlan 1 name DEFAULT-VLAN
!
hostname Extreme
link-fault-signaling
!
end
device(config)# show link-fault-signaling
Global Link Fault : RX ON   TX ON
PORT   #:  LINK FAULT:
PORT 2/1:  RX ON   TX ON
PORT 2/2:  RX ON   TX ON
```

TX LFS is enabled on all ports and RX LFS is enabled on all ports. The previously configured no link-fault-signaling on port 2/1 is overwritten by the global TX LFS enable.

Configuring RX LFS on all ports and enabling TX LFS on one port.

```
device(config)# show run
Current configuration:
!
ver V5.3.0pT183
!
no spanning-tree
!
vlan 1 name DEFAULT-VLAN
!
hostname Extreme
!
end

device(config)# show link-fault-signaling
Global Link Fault : RX ON   TX OFF
PORT   #:  LINK FAULT:

PORT 2/1:  RX ON   TX OFF
PORT 2/2:  RX ON   TX OFF
```

TX LFS is disabled on all ports and RX LFS is enabled on all ports.

```
device(config)# int e 2/1
device(config-if-e10000-2/1)# link-fault-signaling
device(config-if-e10000-2/1)# exit
device(config)# show run
Current configuration:
!
ver V5.3.0pT183
!
no spanning-tree
!
vlan 1 name DEFAULT-VLAN
!
hostname Extreme
!
interface ethernet 2/1
 link-fault-signaling
!
end

device(config)# show link-fault-signaling
Global Link Fault : RX ON   TX OFF
PORT   #:  LINK FAULT:

PORT 2/1:  RX ON   TX ON
PORT 2/2:  RX ON   TX OFF
```

TX LFS is enabled only on port 2/1 and RX LFS is enabled on all ports.

Configuring TX LFS on all ports and enabling RX LFS on all ports except one port.

```
device(config)# link-fault-signaling
device(config)# no link-fault-signaling ignore-rx
device(config)# interface e 2/1
device(config-if-e10000-2/1)# link-fault-signaling ignore-rx
device(config-if-e10000-2/1)# exit
device(config)# show run
Current configuration:
!
ver V5.3.0pT183
!
no spanning-tree
!
vlan 1 name DEFAULT-VLAN
```

```
!
hostname Extreme
link-fault-signaling
!
interface ethernet 2/1
 link-fault-signaling ignore-rx
!
end

device(config)# show link-fault-signaling
Global Link Fault : RX ON   TX ON
PORT   #:  LINK FAULT:

PORT 2/1:  RX OFF  TX ON
PORT 2/2:  RX ON   TX ON
```

TX LFS is enabled on all ports and RX LFS is enabled on all ports except 2/1.

Displaying link-fault-signaling information

You can display information for link-fault-signaling in a Extreme device by using the **show link-fault-signaling** command.

To display if LFS is configured on an interface, enter the following command.

```
device# show link-fault-signaling
Global Link Fault : RX ON   TX OFF
PORT   #:  LINK FAULT:
PORT 2/1:  RX ON   TX OFF
PORT 2/2:  RX ON   TX OFF
PORT 2/3:  RX ON   TX OFF
PORT 2/4:  RX ON   TX OFF
PORT 2/5:  RX ON   TX OFF
PORT 2/6:  RX ON   TX OFF
PORT 2/7:  RX ON   TX OFF
PORT 2/8:  RX ON   TX OFF
PORT 3/1:  RX ON   TX OFF
PORT 3/2:  RX ON   TX OFF
PORT 3/3:  RX ON   TX OFF
PORT 3/4:  RX ON   TX OFF
PORT 3/5:  RX ON   TX OFF
PORT 3/6:  RX ON   TX OFF
PORT 3/7:  RX ON   TX OFF
PORT 3/8:  RX ON   TX OFF
```

NOTE

The **show link-fault-signaling** command does not display RX and TX information for 1 Gb Ethernet ports.

Displaying BIP error information

The **show bip slot** command is used to display a table that contains the lane number for a Physical Coding Sublayer (PCS) lane and a count of Bit Interleaved Parity (BIP) errors for the specific PCS lane. The command output is provided for a lane where a counter is active. The output helps the user to identify the bit parity errors on each physical interface lane of the Extreme MLX 100 GbE modules.

The following example displays an output from the **show bip slot** command on the Extreme NetIron devices.

```
device# show bip slot 3
Port 3/1:
PCS Lane BIP Error Counters :
*****************************
Lane00 : 001 Lane01 : 001
Lane02 : 001 Lane03 : 001
Lane04 : 001 Lane05 : 001
Lane06 : 001 Lane07 : 001
Lane08 : 001 Lane09 : 001
```

```
Lane10 : 001 Lane11 : 001
Lane12 : 001 Lane13 : 001
Lane14 : 001 Lane15 : 001
Lane16 : 001 Lane17 : 001
Lane18 : 001 Lane19 : 001
Port 3/2:
PCS Lane BIP Error Counters :
*****************************
Lane00 : 000 Lane01 : 000
Lane02 : 000 Lane03 : 000
Lane04 : 000 Lane05 : 000
Lane06 : 000 Lane07 : 000
Lane08 : 000 Lane09 : 000
Lane10 : 000 Lane11 : 000
Lane12 : 000 Lane13 : 000
Lane14 : 000 Lane15 : 000
Lane16 : 000 Lane17 : 000
Lane18 : 000 Lane19 : 000
All show BIP done
```

NOTE

The BIP error counter is reset to zero when the command is run. When the counter reaches 255, it does not exceed 255. The counter is increased by a link going up or down and this is expected behavior.

Displaying Network Processor statistics

The Network Processor (NP) counters track the packets and bytes that enter the ingress NP and exit the egress NP. Counts displayed are since the last time the **clear np statistics** command was issued.

The **show np statistics command** displays the NP statistics for all interface modules within a device or for an interface in a specified slot or port. A routed packet drop counter is added to the **show np statistics** command. For more information on the routed packet drop counter, see Table 4. The following example displays an output from the **show np statistics** command on the XMR Series, CES 2000 Series and CER 2000 Series.

Output of the XMR Series is as follows.

```
device # show np statistics ethernet 10/4
NP STATs IPC reply from slot 10 length =1608
Port 10/4 RX
NP Rx Raw Good Packet          = (115458)
NP Rx Forward Packet            = (115458)
NP Rx Discard Packet           = (0)
NP Rx Unicast Packet           = (44571)
NP Rx Broadcast Packet         = (0)
NP Rx Multicast Packet         = (70887)
NP Rx Send to TM Packet        = (115458)
NP Rx Bad Packet               = (0)
NP Rx Lookup Unavailable       = (0)
NP Rx ACL Drop                 = (0)
NP Rx Priority 0/1 Drop        = (0)
NP Rx Priority 2/3 Drop        = (0)
NP Rx Priority 4/5 Drop        = (0)
NP Rx Priority 6/7 Drop        = (0)
NP Rx Suppress RPF Drop        = (0)
NP Rx RPF Drop                 = (0)
NP Rx IPv4 Packet              = (0)
NP Rx IPv6 Packet              = (0)
NP Rx Route-only Drop          = (0)
NP Rx IPv6 Suppress RPF Drop   = (0)
NP Rx IPv6 RPF Drop Count      = (0)
NP Rx IPv4 Byte                = (0)
NP Rx IPv6 Byte                = (0)
NP Rx Routed Packet Drop     = (0)
Port 10/4 TX
NP Tx Sent to MAC Packet       = (1365518)
```

```
NP Tx Raw Good Packet          = (1365518)
NP Tx Source Port Supress Drop = (0)
NP Tx Bad Packet Count         = (0)
NP Tx Unicast Packet           = (1324427)
NP Tx Broadcast Packet         = (1)
NP Tx Multicast Packet         = (41090)
NP Tx IPX HW Forwarded Packet  = (41090)
NP Tx Receive from TM          = (1365518)
NP Tx ACL Drop                 = (0)
NP Tx IPv4 Packet              = (0)
NP Tx IPv6 Packet              = (0)
NP Tx IPv4 Byte                = (0)
NP Tx IPv6 Byte                = (0)
```

Syntax: show np statistics [**ethernet** *slot/port*] [**slot** *slot-num*]

You can use the **ethernet** option and specify a *slot/port* variable to display NP statistics for an individual port.

You can use the slot option and specify a *slot-num* variable to display NP statistics for an individual interface module.

Output of the CES 2000 Series is as follows.

```
device#show np statistics
TD: Traffic Descriptor. Each TD has size of 512 Bytes
MODULE # 0 PPCR # 0 :
Ingress Counters :
Received packets                                    = 0
Received TDs on traffic class 0                     = 0
Received TDs on traffic class 0                     = 0
Received TDs on traffic class 1                     = 0
Received TDs on traffic class 2                     = 0
Received TDs on traffic class 3                     = 0
Received TDs on traffic class 4                     = 0
Received TDs on traffic class 5                     = 0
Received TDs on traffic class 6                     = 0
Received TDs on traffic class 7                     = 0
Egress Counters :
Transmitted unicast packets                         = 0
Transmitted multicast packets                       = 0
Transmitted broadcast packets                       = 0
Filtered packets due to VLAN spanning tree          = 0
Tail dropped packets                                = 0
Control packets                                     = 0
Packets filtered due to egress forward restrictions = 0)
```

Syntax: show np statistics [**slot** *slot-num*]

You can use the **slot** option and specify a *slot-num* variable to display NP statistics for an individual interface module.

For CES 2000 Series and CER 2000 Series, you can either use **show np statistics** command or **show np statistics slot** *slot-num* command to display the NP statistics for an interface in a specified slot.

The **Tx** and **Rx** counters displayed are described in the following tables.

TABLE 4 Rx counters

Rx counter (per port)	Explanation
Rx Raw Good Packet	Number of good packets received from MAC
Rx Forward Packet	Number of forwarded packets by packet evaluation engine
Rx Discard Packet	Number of packets flagged for discard by packet evaluation engine
Rx Unicast Packet	Number of unicast (indicated by MAC DA) packets received
Rx Broadcast Packet	Number of broadcast (indicated by MAC DA) packets received
Rx Multicast Packets	Number of multicast (indicated by MAC DA) packets received
Rx Send to TM Packets	Number of packets sent to TM (= Rx Forward Packet - RL drops)

TABLE 4 Rx counters (continued)

Rx counter (per port)	Explanation
Rx Bad Packets	Number of packets that have MAC to NP interface errors
Rx Loopup Unavailable	Number of packets that have been dropped due to unavailability of the CAM interface for packet lookups
Rx ACL Drop	Drop counter for ACL drop on the ingress path
Rx Priority 0/1 Drop	Drop counter for ingress priority 0,1 packets
Rx Priority 2/3 Drop	Drop counter for ingress priority 2,3 packets
Rx Priority 4/5 Drop	Drop counter for ingress priority 4,5 packets
Rx Priority 6/7 Drop	Drop counter for ingress priority 6,7 packets
Rx Supress RPF Drop	Counter for suppressed RPF drops on the ingress path due to ACL override
Rx RPF Drop	Counter for RPF drop on the ingress
Rx IPv4 Packet	Raw packet count that have IPv4 EType (0x0800) and IP version of 0x4
Rx IPv6 Packet	Raw packet count that have IPv6 EType (0x86DD) and IP version of 0x6
Rx IPv6 Supress RPF Drop	Counter for IPv6 suppressed RFP drops on the ingress path due to ACL override
Rx IPv6 RPF Drop Count	Counter for IPv6 drop on the ingress
NP Rx Route-only Drop	Counts packets that have been dropped due to Route-Only configuration during MAC-DA processing.
Rx IPv4 Byte	Raw packet Bytes (+FCS) that have IPv4 etype (0x0800) and IP version equals 0x4
Rx IPv6 Byte	Raw packet Bytes (+FCS) that have IPv6 etype (0x86DD) and IP version equals 0x6
Rx Routed Packet Drop	Number of received IPv4 or IPv6 routed packets that are dropped because the TTL is 0, or because routing is not enabled on the given virtual interface.

TABLE 5 Tx counters

TX counter (per port)	Explanation
Tx Sent to MAC Packet	Total number of packets sent to MAC for transmit
Tx Raw Good Packet	Total number of packets sent to egress processing logic that pass the initial length checks (min, max, offsets, bad packet etc.)
Tx Source Port Suppression Drop	Number of packets dropped because of transmit source port suppression
Tx Bad Packet Count	Total number of packets dropped in egress logic that fail the initial length checks (min, max, bad packet etc.)
Tx Unicast Packet	Number of unicast packets transmitted (from MAC DA)
Tx Broadcast Packet	Number of broadcast packets transmitted (from MAC DA)
Tx Multicast Packet	Number of multicast packets transmitted (from MAC DA)
Tx Receive From TM	Number of packets received from TM
Tx ACL Drop	Number of packets that have been dropped by the Outbound ACL Logic
Tx IPv4 Packet	Number of IPv4 packets transmitted out the port (Etype==0x0800 & IPver == 0x4)
Tx IPv6 Packet	Number of IPv6 packets transmitted out the port (Etype==0x86DD & IPver == 0x6)
Tx IPv4 Byte	Counts packet Bytes (+FCS) that have IPv4 etype (0x0800) and IP version equals 0x4
Tx IPv6 Byte	Counts packet Bytes (+FCS) that have IPv6 etype (0x86DD) and IP version equals 0x6

Relationships between some counters

Some of the values for counters displayed using the **show np statistics** command are the result of adding the contents of more than one counter. The following tables describe these relationships between NP counters displayed.

TABLE 6 Relationships between RX counters

Total RX Packets	=	Rx Bad Packets + Rx Lookup Unavailable Packets + Rx Raw Good Packets
Rx Raw Good Packets	===	Rx Unicast Packets + Rx Multicast Packets + Rx Broadcast Packets+Rx IPv4 Packets + Rx IPv6 Packets + Rx Other Packets+Rx Forward Packets + Rx Discard Packets
Rx Forward Packets	=	Rx Sent to TM Packets + Rx RL drop packets
Rx Discard Packets	=	ACL drop + TTL drop + route-only drop + RPF drop + tag mismatch drop+ VLAN blocking drop + segment filtering drop+ drop by packet evaluation decisions +miscellaneous
Rx Priority Drops	=	RL drop + Rx Discard Packets

TABLE 7 Relationships between TX counters

Tx Raw Good Packets	==	Tx Receive From TM Packets - Tx Bad Packets+Tx Unicast Packets + Tx Broadcast Packets + Tx Multicast Packets + Tx Source Port Suppression Drop
Tx Sent to MAC	==	Tx IPv4 Packets + Tx IPv6 Packets + Tx Others+Tx Raw Good Packets - Tx Source Port Suppression Drop - Tx ACL drop - Tx RL Drop - Tx Multicast TTL drop

Clearing the NP statistics counters

You can clear the NP statistics counters for an entire device or selectively by port or slot using the **clear np statistics** command as shown in the following.

```
device# clear np statistics
```

Syntax: clear np statistics [**ethernet** *slot/port*] [**slot** *slot-num*]

You can use the **ethernet** option and specify a *slot/port* variable to clear NP statistics for an individual port.

You can use the **slot** option and specify a *slot-num* variable to clear NP statistics for an individual interface module.

Capturing hardware errors from Tsec statistics and logging in syslog and console

All the IPC communication between LP and MP happens through backplane ethernet. Each LP has a Tsec chip through which RX and TX of the packets from MP are processed. Any drops in the backplane LP ethernet controller and errors on packets encountered during LP-MP communications are recorded in Tsec. Viewing these errors only through the Tsec commands does not help you to identify these errors at run time as soon as they are encountered.

This feature monitors some of the errors encountered in Tsec like FCS error, code error, and carrier sense error while receiving the packet. When any of these Rx error counters are set, appropriate error message are displayed on LP console. The log message are displayed in the following format.

```
Jun 27 10:32:21 2016: TSEC device: 1, Rx Carrier Sense error: 41
Jun 27 10:33:06 2016: TSEC device: 0, Rx FCS Error: 1
Jun 27 10:33:06 2016: TSEC device: 0, Rx Code error: 1
```

In addition to this, those messages which are displayed in the LP console while trying to process a free buffer will also be displayed in syslog. Please refer the *Syslog* section for additional information on syslog.

About ICMP request handler offload

With the NetIron 6.2.0 release, a new feature, ICMP request handler offload is being introduced. This feature is supported on Extreme NetIron XMR Series and Extreme MLX Series devices. Enabling this feature helps hosts to receive a faster ICMP response than that observed in pre-6.2.0 releases. This feature is designed for optimizing response time for tools that monitor or troubleshoot Extreme XMR and MLX devices for network issues periodically.

CPU populates a table with the IPs of hosts for which PBIF needs to send ICMP response, and the destination FID where the packet to be sent out. Without any performance degradation, the CPU processes the ICMP, resolves the destination identifiers for the response, and populates a table in PBIF so that from the next ping packet onwards, the response is sent out from PBIF. When ICMP request packet arrives in PBIF, and if no entry is found in the table, or the packets source IP does not match with entries in the table, the packet is sent to CPU. The PBIF creates an ICMP response packet with MAC and IP addresses swapped, sets ICMP type to Response, maintains the same ICMP payload recalculate all the checksums, and then sends it back to the destination port as configured in the table. The PBIF validates the packet for IP and ICMP checksums. In case of a failure, it sends the packet to CPU as is and error handling and statistics are updated by CPU.

Configuring ICMP request handler offload

Use the following steps to configure ICMP request handler offload.

Use the following steps to assign an alphanumeric name to multiple ports in an interface.

1. Enter global configuration mode.

```
switch# configure terminal
```

2. Enter interface configuration mode.

```
switch(config)# interface ethernet 12/11
```

3. Enable the interface

```
switch(config-if-e1000-12/11)# enable
```

4. Configure an IP address.

```
switch(config-if-e1000-12/11)#  ip address 4.4.4.2/24
```

5. Disable 802.3ae 10G link fault signaling

```
switch(config-if-e1000-12/11)# no link-fault-signaling
```

6. Set the global Gig port default option to non-autonegotiation.

```
switch(config-if-e1000-12/11)# gig-default neg-off
```

7. Configure ICMP fast-echo-reply.

```
switch(config)# ip icmp fast-echo-reply 4.4.4.2 4.4.4.0 255.255.255.0 10
switch(config)# ip icmp fast-echo-reply 4.4.4.2 100.100.100.0 255.255.255.0 360
switch(config)# ip icmp fast-echo-reply any any 5
```

8. Configure an IP static route.

```
switch(config)# ip route ip route 100.100.100.0/24 3.3.3.1
```

9. Configure a VRF instance.

```
switch(config)# vrf test
```

10. enter address family command mode and configure an IPv4 address.

```
switch(config)# address-family ipv4 max-route 200000
```

11. enter address family command mode and configure an IPv4 address.

```
switch(config)# address-family ipv4 max-route 200000
```

Reliability, Availability, and Serviceability

Auto-tune enhancement

The RAS feature set is extended with automatic monitoring and tuning of transmit and receive parameters on SERDES links between line modules and switch fabric modules that are down due to excessive CRC errors. Tuning is done on both the line module and switch fabric module i.e. both ends of the the link, at the same time.

If tuning fails on the switch fabric module then the **sysmon fe link action** configuration defines the action taken i.e. a syslog message is generated or the link is shut down and a syslog message is generated. If the link has already been tuned and goes down a second time, then the link is powered down and a syslog message generated.

Auto-tune is enabled by default on MLXe-16 and MLXe-32 chassis. This enhancement is not needed on MLXe-4 and MLXe8 chassis, because the CRC error condition does not occur on these chassis.

To disable auto-tuning on FE for slow or burst CRC errors, enter the following command:

```
device(config)# no sysmon fe link auto-tune
```

To enable auto-tuning again, use the following command:

```
device(config)# sysmon fe link auto-tune
```

Syntax: [no] sysmon fe link auto-tune

To disable auto-tuning on TM for slow or burst CRC errors, enter the following command:

```
device(config)# no sysmon tm link auto-tune
```

To enable auto-tuning again, use the following command:

```
device(config)# sysmon tmlink auto-tune
```

Syntax: [no] sysmon tm link auto-tune

Operations, Administration, and Maintenance

Operations, Administration, and Maintenance (OAM) implementation refers to the tools and utilities for installing, monitoring, and troubleshooting the network.

IEEE 802.1ag Connectivity Fault Management

IEEE 802.1ag Connectivity Fault Management (CFM) refers to the ability of a network to monitor the health of a service delivered to customers as opposed to just links or individual bridges.

The IEEE 802.1ag CFM standard specifies protocols, procedures, and managed objects to support transport fault management. This allows for the discovery and verification of the path, through bridges and LANs, taken by frames addressed to and from specified network users and the detection, and isolation of a connectivity fault to a specific bridge or LAN.

Ethernet CFM defines proactive and diagnostic fault localization procedures for point-to-point and multipoint Ethernet Virtual Connections that span one or more links. It operates end-to-end within an Ethernet network.

Ethernet OAM capabilities

Ethernet OAM is able to:

- Monitor the health of links (because providers and customers might not have access to the management layer)
- Check connectivity of ports
- Detect fabric failures
- Provide the building blocks for error localization tools
- Give appropriate scope to customers, providers and operators (hierarchical layering of OAM)
- Avoid security breaches

IEEE 802.1ag purpose

Bridges are increasingly used in networks operated by multiple independent organizations, each with restricted management access to each other's equipment. CFM provides capabilities for detecting, verifying and isolating connectivity failures in such networks.

There are multiple organizations involved in a Metro Ethernet Service: Customers, Service Providers and Operators.

Customers purchase Ethernet Service from Service Providers. Service Providers may utilize their own networks, or the networks of other Operators to provide connectivity for the requested service. Customers themselves may be Service Providers, for example a Customer may be an Internet Service Provider which sells Internet connectivity.

Operators will need minimal Ethernet OAM. Providers will need more comprehensive Ethernet OAM for themselves and to allow customers better monitoring functionality.

FIGURE 1 OAM Ethernet tools

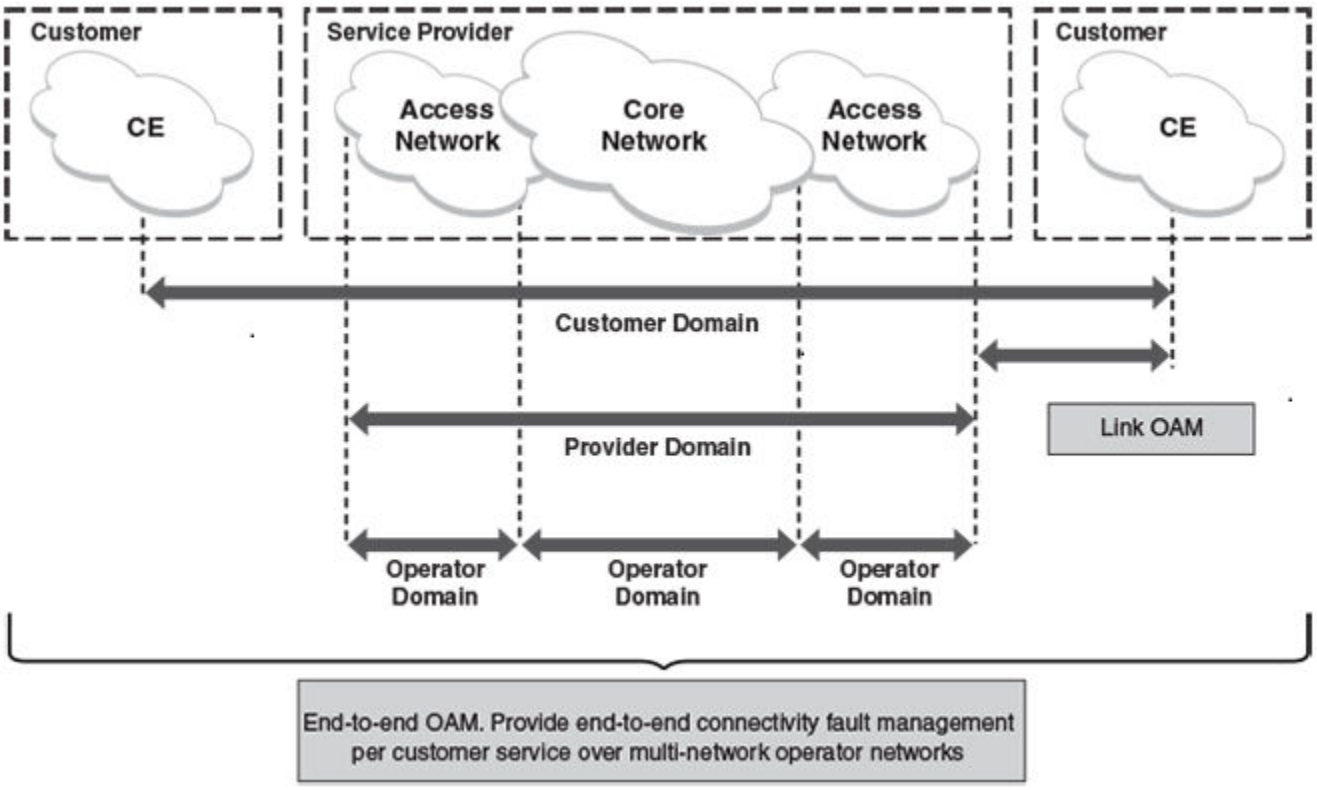

IEEE 802.1ag provides hierarchical network management

Maintenance Domain

A Maintenance Domain (MD) is part of a network controlled by a single operator. Figure 1 on page 42, shows the customer domain, provider domain and operator domain.

Maintenance Domain level

The Maintenance Domain levels (MD level) are carried on all CFM frames to identify different domains. For example, in Figure 1 on page 42, some bridges belong to multiple domains. Each domain associates a MD level.

- Customer Level: 5-7
- Provider Level: 3-4
- Operator Level: 0-2

Maintenance Association

Every MD can be further divided into smaller networks having multiple Maintenance End Points (MEP). Usually a Maintenance Association (MA) is associated with service instances (for example a VLAN or a VPLS).

Maintenance End Point (MEP)

Maintenance End Point (MEP) is located on the edge of a Maintenance Association (MA). It defines the endpoint of the MA. Each MEP has unique ID (MEPID) within MA. The connectivity in a MA is defined as connectivity between MEPs. MEP generates Continuity Check Message and multicasts to all other MEPs in same MA to verify the connectivity.

Maintenance Intermediate Point

Maintenance Intermediate Point (MIP) is located within a Maintenance Association (MA). It responds to Loopback and Linktrace messages for Fault isolation.

Mechanisms of Ethernet IEEE 802.1ag OAM

Mechanisms supported by IEEE 802.1ag include Connectivity Check (CC), Loopback, and Link trace. Connectivity Fault Management allows end-to-end fault management that is generally reactive (through Loopback and Link trace messages) and connectivity verification that is proactive (through Connectivity Check messages).

Fault detection (Continuity Check Message)

The Continuity Check Message (CCM) provides a means to detect hard and soft faults such as software failure, memory corruption, or misconfiguration. The failure detection is achieved by each Maintenance End Point (MEP) transmitting a CCM periodically within its associated Service Instance.

As a result, MEPs also receive CCMs periodically from other MEPs. If a MEP on local Bridge stops receiving the periodic CCMs from peer MEP on a remote Bridge, it can assume that either the remote Bridge has failed or failure in the continuity of the path has occurred. The Bridge can subsequently notify the network management application about the failure and initiate the fault verification and fault isolation steps either automatically or through operator command.

A CCM requires only N transmissions within its member group, where N is the number of members within the member group. In other words, if a Virtual Bridge LAN Service has N members, only N CCMs need to be transmitted periodically, one from each.

Continuity Check (CC) messages are periodic hello messages multicast by a MEP within the maintenance domain, at the rate of X; All Maintenance association Intermediate Points (MIPs) and MEPs in that domain will receive it but will not respond to it. The receiving MEPs will build a MEP database that has entities of the format. MEPs receiving this CC message will catalog it and know that the various maintenance associations (MAs) are functional, including all intermediate MIPs.

NOTE
The Extreme NetIron CES does not support sub-second values.

CCMs are not directed towards any specific; rather they are multicast across the entire point-to- point or multipoint service on a regular basis. Accordingly, one or more service flows, including the determination of MAC address reachability across a multipoint network, are monitored for connectivity status with IEEE 802.1ag.

Fault verification (Loopback messages)

A unicast Loopback Message is used for fault verification. To verify the connectivity between MEP and its peer MEP or a MIP, the Loopback Message is initiated by a MEP with a destination MAC address set to the MAC address of either a Maintenance association Intermediate Point (MIP) or the peer MEP. The receiving MIP or MEP responds to the Loopback Message with a Loopback Reply.

A Loopback message helps a MEP identify the precise fault location along a given MA. A Loopback message is issued by a MEP to a given MIP along an MA. The appropriate MIP in front of the fault will respond with a Loopback reply. The MIP behind the fault will not respond. For Loopback to work, the MEP must know the MAC address of the MIP to ping.

Fault isolation (Linktrace messages)

Linktrace mechanism is used to isolate faults at Ethernet MAC layer. Linktrace can be used to isolate a fault associated with a given Virtual Bridge LAN Service. It should be noted that fault isolation in a connectionless (multi-point) environment is more challenging than a connection oriented (point-to-point) environment. In case of Ethernet, fault isolation can be even more challenging since a MAC address can age out when a fault isolates the MAC address. Consequently a network-isolating fault results in erasure of information needed for locating the fault.

A Linktrace Message uses a set of reserved multicast MAC address. The Linktrace Message gets initiated by a MEP and traverses hop-by-hop and each Maintenance Point (a MEP or MIP) along the path intercepts this Linktrace Message and forwards it onto the next hop after processing it until it reaches the destination MEP. The processing includes looking at the destination MAC address contained in the Linktrace Message.

Each MP along the path returns a unicast Linktrace Reply back to the originating MEP. The MEP sends a single LTM to the next hop along the trace path; however, it can receive many Linktrace Responses from different MPs along the trace path and the destination MEP as the result of the message traversing hop by hop. As mentioned previously, the age-out of MAC addresses can lead to erasure of information at MIPs, where this information is used for the Linktrace mechanism. Possible ways to address this behavior include:

- Carrying out Linktrace following fault detection or verification such that it gets exercised within the window of age-out.
- Maintaining information about the destination MEP at the MIPs along the path using CCMs.
- Maintaining visibility of path at the source MEPs through periodic LTMs.

Linktrace may also be used when no faults are apparent in order to discover the routes normally taken by data through the network. In the rare instances during network malfunctions where Linktrace cannot provide the information needed to isolate a fault, issuing Loopback Messages to MPs along the normal data path may provide additional useful information.

The Linktrace message is used by one MEP to trace the path to another MEP or MIP in the same domain. It is needed for Loopback (Ping). All intermediate MIPs respond back with a Link trace reply to the originating MEP. After decreasing the TTL by one, intermediate MIPs forward the Link trace message until the destination MIP or MEP is reached. If the destination is a MEP, every MIP along a given MA responds to the originating MEP. The originating MEP can then determine the MAC address of all MIPs along the MA and their precise location with respect to the originating MEP.

Configuring IEEE 802.1ag CFM

Enabling or disabling CFM

To enable or disable the CFM protocol globally on the devices and enter into the CFM Protocol Configuration mode, enter a command such as the following.

```
device(config)#cfm-enable
device(config-cfm)#
```

Syntax: [no] cfm-enable

The **no** form of the command disables the CFM protocol.

Creating a Maintenance Domain

A Maintenance Domain is the network or the part of the network for which faults in connectivity are to be managed. A Maintenance Domain consists of a set of Domain Service Access Points.

A Maintenance Domain is, or is intended to be, fully connected internally. A Domain Service Access Point associated with a Maintenance Domain has connectivity to every other Domain Service Access Point in the Maintenance Domain, in the absence of faults.

Each Maintenance Domain can be separately administered.

The **domain-name** command in CFM protocol configuration mode creates a maintenance domain with a specified level and name and enters the Specific Maintenance Domain mode specified in the command argument.

```
device(config-cfm)#domain-name VPLS-SP level 4
device(config-cfm-md-VPLS-SP)#
```

Syntax: [no] domain-name *name* [**id** *md-id*] [**level** *level*]

The *name* variable specifies the domain name. The *name* is case-sensitive.

The **id** *md-id* parameter is the Maintenance Domain Index. It is an optional parameter. The range is 1 - 4090.

The **level** *level* parameter sets the domain level in the range 0 - 7. When the domain already exists, the **level** argument is optional. The levels are.

- Customer's Domain Levels: 5 - 7
- Provider Domain Levels: 3 - 4
- Operator Domain Levels: 0 - 2

The **no** form of the command removes the specified domain from the CFM Protocol Configuration mode.

Setting Maintenance Domain parameters

Creating Maintenance Associations

The Maintenance Association Identifier is unique over the domain. If the Maintenance Association Identifier is globally unique, then that domain is global. CFM can detect connectivity errors only for a list of MEPs with unique MAIDs.

The **ma-name** command, in Maintenance Domain mode, creates a maintenance association within a specified domain. The **ma-name** command changes the Maintenance Domain mode to a Specific Maintenance Association mode.

```
device(config-cfm-md-VPLS-SP)# ma-name ma_1 vlan-id 30 priority 4
device(config-cfm-md-VPLS-SP-ma-ma_1)#
```

Syntax: [no] ma-name *name* [**id** *ma-id*] [**esi** *esi-id*] [**vlan-id** *vlan-id*] [**vpls-id vpls-id**] [**priority** *priority*]

The **ma-name** *name* parameter specifies the maintenance association name. The NAME attribute is case-sensitive.

The **id** *ma-id* is the Maintenance Association Index. It is an optional parameter. The range is 1 - 4090.

The *esi-id* specifies a unique ESI identifier of the maintenance association. In case of creating a MA a ESI ID should be set. This option is available only on platforms that support the Ethernet Service Instance (ESI) framework.

The *vlan-id* specifies a unique VLAN identifier of the maintenance association in the range *1-4090* . In case of creating a MA a VLAN ID should be set.

The *vpls-id* specifies a unique VPLS identifier of the maintenance association. In case of creating a MA, a VPLS ID should be set.

The *priority* parameter specifies the priority of the CCM messages, sent by MEPs, in the range *0-7* . When the maintenance association is already created, the *priority* argument is optional.

The **no** form of the command removes the created MA.

Tag-type configuration

For the NetIron CES, the following two VLAN tag-types are allowed that can be configured globally:

- tag1 - applies to customer edge ports (CVLAN) by default.
- tag2 - applies to provider-network, backbone-edge, and backbone-network port types (SVLAN and BVLAN) by default.

  NOTE
  The tag1 and tage2 are independent of port-types, so the system can be configured to use tag1 for SVLAN, BVLAN and tag2 for CVLAN.

Configuring tag-types

You can set the ISID value using a separate command similar to NetIron XMR.

Syntax: [no] tag-value isid *num*

You can configure CVLAN, SVLAN, and BVLAN tag-types as shown below.

```
device(config)# tag-value tag1 8100
device(config)# tag-value tag2 9100
device(config)# tag-type cvlan tag1 svlan tag2 bvlan tag2
```

Syntax: [no] tag-value *num*

Syntax: tag-type *tag-n*

The *num* parameter specifies the value assigned to the tag. The default value for *tag1* is 0x8100 and for *tag2* is 0x88a8.

The *tag-n* parameter can be either tag1 or tag2.

Tag type can be changed from a default value to a specific port as shown in the following example.

```
device(config-if-e1000-1/1)# tag-type tag2 ethernet 1/1
device(config-if-e1000-1/1)# tag-type tag1 ethernet 1/2
```

Syntax: tag-type *tagid* **ethernet** *interface_id*

The *tagid* parameter can be either **tag1** or **tag2**.

The *interface_id* parameter specifies the Ethernet slot and port ID.

Restrictions

The tag-type has the following restrictions:

- CVLAN and SVLAN cannot have the same tag-type.
- SVLAN and BVLAN must have the same tag-type.
- Port-type must be set to the default to configure the port-level tag-type.

Configuring a CCM interval for a Maintenance Association

The **ccm-interval** command sets the time interval between two successive Continuity Check messages (CCMs) that are sent by MEPs in the specified Maintenance Domain. The default value is 10 seconds.

```
device(config-cfm)#domain name VPLS-SP level 4
device(config-cfm-md-VPLS-SP)#ma-name ma_1 vlan-id 30 priority 3
device(config-cfm-md-VPLS-SP-ma-ma_1)#ccm-interval 10-second
device(config-cfm-md-VPLS-SP-ma-ma_1)#
```

Syntax: [no] ccm-interval [**1-second** | **1-minute** | **10-second** | **10-minute** | **3.3-ms** | **10-ms** | **100-ms**]

The **1-second** parameter sets the time interval between two successive CCM packets to 1 second.

The **1-minute** parameter sets the time interval between two successive CCM packets to 1 minute.

The **10-second** parameter sets the time interval between two successive CCM packets to 10 seconds.

The **10-minute** parameter sets the time interval between two successive CCM packets to 10 minutes.

The **3.3-ms** parameter sets the time interval between two successive CCM packets to 3.3 milliseconds.

The **10-ms** parameter sets the time interval between two successive CCM packets to 10 milliseconds.

The **100-ms** parameter sets the time interval between two successive CCM packets to 100 milliseconds.

Configuring local ports

The **mep** command, in Maintenance Association mode, adds local ports as MEP to a specific maintenance association. If configuring a CFM packet to a "down" MEP, it will need to be sent out on the port on which it was configured. If configuring a CFM packet to an "up" MEP, it will need to be sent to the entire VLAN for multicast traffic, and unicast traffic will need to be sent to a particular port according to the MAC table.

Configuring a MED for each of the Domain Service Access Points of a service instance creates a MA to monitor the connectivity:

- The list of MEPs configured with identical values for MA ID defines an MA.
- Each Bridge has its own Maintenance Association managed object for an MA.
- Each individual MEP is configured with a ID that is unique within that MA.
- Each MEP is associated with a Service Access Point that provides access to a single service instance.

  NOTE
  When configuring 802.1ag over VPLS, if the VPLS endpoint is deleted from the configuration, the MEP configuration is deleted under CFM without warning.

To add local ports to an upstream MEP, enter commands such as the following.

```
device(config-cfm)# domain name VPLS-SP level 4
device(config-cfm-md-VPLS-SP)# ma-name ma_1 vlan-id 30 priority 3
device(config-cfm-md-VPLS-SP-ma_1)# mep 1 up port eth 2/1
device(config-cfm-md-VPLS-SP-ma_1)#
```

Syntax: [no] mep *mep-id* [**up** | **down**] [**vlan** *vlan-id* **port ethernet** *slot/port* | **port ethernet** *slot/port*]

The *mep-id* parameter specifies the maintenance end point ID (mandatory) in the range *1-8191* .

The **up** parameter sets the MEP direction away from the monitored VLAN.

The **down** parameter sets the MEP direction towards the monitored VLAN.

The *vlan-id* parameter specifies the VLAN end-points. It is configured only for MAs associated with VPLS and not configured for MAs with a VLAN.

The *port-id* parameter specifies the target interface on which it is used.

The **no** form of the command removes the specified MEPs.

Configuring remote MEPs

The **remote-mep** command is used to configure the remote MEP's you are expecting. If a remote MEP is not specified, the remote MEP database is built based on the CCM. If one remote MEP never sends CCM, the failure can not be detected.

```
device(config-cfm-md-VPLS-SP)# ma-name ma_1 vlan-id 30
device(config-cfm-md-VPLS-SP-ma_1)# remote-mep 1 to 120
device(config-cfm-md-VPLS-SP-ma_1)#
```

Syntax: [no] remote-mep *mep-id* [**to** *mep-id*]

The *mep-id* parameter specifies the maintenance end point ID (mandatory) in the range 1-8191.

The **no** form of the command removes the specified remote MEPs.

Setting the Remote Check Start-Delay

When configuring the remote MEPs range, you can set a wait time before the MEPs come up and the CCM check operation is started. The default is set to 30 seconds.

```
device(config)#cfm-enable
device(config-cfm)#rmep-check start-delay 120
device(config-cfm)#
```

Syntax: [no] rmep-check start-delay *seconds*

The *seconds* parameter is the wait time interval before the CCM check is started. The range is 10 - 600 seconds.

Specifying MIP creation policy

The **mip-policy** command, in Maintenance Association mode, specifies the conditions in which MIPs are automatically created on ports.

> NOTE
> MIP functionality of 802.1ag over VPLS with sub-second timer will have all the configuration restrictions of the VPLS CPU-protection.

A MIP can be created on a port and VLAN, only when explicit or default policy has been defined for them. For a specific port and VLAN a MIP will be created at the lowest of the levels. Additionally, the level created should be the next higher than the MEP level defined for these port and VLAN.

```
device(config-cfm)#domain name VPLS-SP level 4
device(config-cfm-md-VPLS-SP)#ma-name ma_1 vlan-id 30
device(config-cfm-md-VPLS-SP-ma_1)#mip-policy explicit
device(config-cfm-md-VPLS-SP-ma_1)#
```

Syntax: [no] mip-policy [**explicit** | **default**]

Use the **explicit** explicit parameter to specify that explicit MIPs are configured only if a MEP exists on a lower MD Level.

Use the **default** parameter to specify that MIPs will always be created.

The **no** form of the command restores the default policy.

Y.1731 performance management

The Y.1731 feature provides the following performance monitoring capability for point-to-point links as defined in ITU-T Rec Y.1731:

- Two-way Frame Delay Measurement (ETH-DM)
- Two-way Frame Delay Measurement Variation

NOTE
One-way ETH-DM is not supported in this release of the Multi-Service IronWare.

About Y.1731

Figure 2 shows an ETH-DM requester and responder.

FIGURE 2 ETH-DM requester and responder.

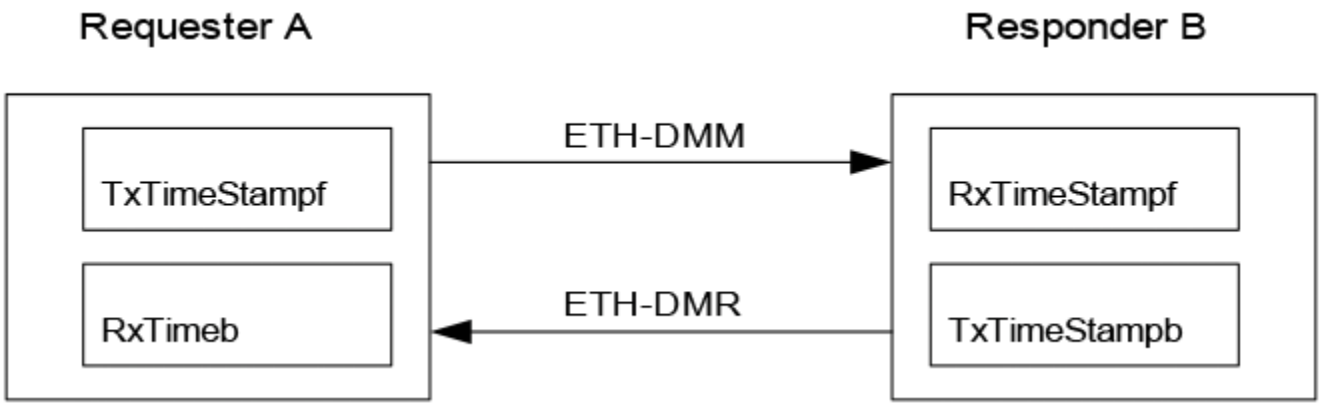

ETH-DM packets are transmitted, received and processed by LP CPU and are timestamped by the hardware on transmit and receive path.

An ETH-DM packet contains 4 timestamps for measuring the round-trip delay.

Requester A, transmits ETH-DMM packets with TxTimeStampf (timestamp at the transmission time of the packet).

Responder B, responds with an ETH-DMR packet using two timestamps to account for its processing time: RxTimeStampf (Timestamp at the time of receiving the DMM packet) and TxTimeStampb (timestamp at the time of transmitting the DMR packet).

Upon receiving an ETH-DMR packet, requester A stamps the packet with RxTimeb (timestamp at the time the DMR packet is received).

Frame Delay = (RxTimeb – TxTimeStampf) – (TxTimeStampb – RxTimeStampf)

Y.1731 support in NetIron 6.0.00

NetIron 6.0.00 provides Y.1731 support for the following:

- VLANs
- VPLS—Both VC-mode tagged and raw
- VLL—Both tagged and raw modes
- Up and Down MEPs for VLANs, VPLS, and VLL
- Over LAG ports—The active primary port of the trunk would be used to transmit ETH-DM frames in case of down MEP
- Through 802.1ag MIPs—MIP would behave as a transient node for ETH-DM frames

Y.1731 support for Extreme BR-CER-2024C/F-4X(RT)

Beginning with NetIron 6.2.0 release, Y.1731 is supported on Extreme BR-CER-2024C/F-4X(RT).

The Y.1731 support on Extreme BR-CER-2024C/F-4X(RT) includes the following.

- Y1731-based 1-DM—one-way delay measurement using Y1731 frames between pair of point-to-point MEP.
- Y1731-based SLM—synthetic loss measurement to check the packet loss using artificially generated Y1731 frames.
- Y1731-based LMM—loss measurement using actual data traffic.
- Features should be supported for VLAN/VPLS/VLL/LAG.

  NOTE
  Currently, CFM support over VLL is not available.

Configuration considerations:

When using Y.1731, consider the following:

- ETH-DM is reliable only if the transmitted DM frame (DMM) and received DM reply (DMR) are on the same line processor (LP). In the event that they are different, results will not be accurate.
- Maximum frame-delay that can be measured is 4 seconds. If a DMR packet is received with a delay greater than 4 seconds, the packet is discarded and ignored.
- ETH-DM does not gather path data. To determine which path the DM applies to, use the **cfm linktrace domain** command, since ETH-DM frames follow the same path.
- One-way ETH-DM is not supported in this release of the Multi-Service IronWare.

Configuring Y.1731 performance monitoring

Use the **cfm delay_measurement domain** command to issue the delay measurement. If the number of delay measurement frame is greater than 16, then the last 16 delay measurement replies are printed.

You can issue the **cfm delay measurement** command from different sessions if they are for different **src-meps** . However, if it is for same **src-mep** , it only completes one session at a time.

```
device# cfm delay_measurement domain md2 ma ma2 src-mep 3 target-mep 2
 Y1731: Sending 10 delay_measurement to 0000.00f7.3931, timeout 1000 msec
        Type Control-c to abort
        Reply from 0000.00f7.3931: time= 32.131 us
        Reply from 0000.00f7.3931: time= 31.637 us
        Reply from 0000.00f7.3931: time= 32.566 us
        Reply from 0000.00f7.3931: time= 34.052 us
        Reply from 0000.00f7.3931: time= 33.376 us
        Reply from 0000.00f7.3931: time= 31.501 us
```

```
 Reply from 0000.00f7.3931: time= 33.016 us
 Reply from 0000.00f7.3931: time= 32.537 us
    Reply from 0000.00f7.3931: time= 32.492 us
 Reply from 0000.00f7.3931: time= 32.552 us
 sent = 10 number = 10 A total of 10 delay measurement replies received.
 Success rate is 100 percent (10/10)
 ================================================================================
Round Trip Frame Delay Time : min = 31.501 us avg = 32.586 us max = 34.052 us
Round Trip Frame Delay Variation : min = 45 ns avg = 839 ns max = 1.875 us
================================================================================
```

Syntax: cfm delay_measurement domain *domain-name* **ma** *ma-name* **src-mep** *mep-id* **target-mep** *mep-id* [**timeout** *timeout*] [**number** *number*]

The **domain** *domain-name* parameter specifies the maintenance domain to be used for a delay measurement message. The *domain-name* attribute is case-sensitive.

The **ma** *ma-name* parameter specifies the maintenance association to be used for a delay measurement message. The *ma-name* attribute is case-sensitive.

The **src-mep** *mep-id* parameter specifies the source mep-id in the range 1-8191.

The **target-mep** *mep-id* parameter specifies the destination mep-id in the range 1-8191.

The **number** *number* parameter specifies the number of delay_measurement messages to be sent. The range is 1-1000. The default value is 10. This is an optional parameter.

The **timeout** *timeout* parameter specifies the timeout used to wait for previous delay_measurement reply before sending the next delay_measurement message. The range is 1-4 seconds. The default value is 1second. This is an optional parameter.

If a **delay_measurement** reply is received before the timeout, then the next delay measurement frame is sent immediately after processing the delay measurement reply. However, if the **delay measurement** reply is not received within the specified timeout, then the next **delay measurement** frame will be sent.

Y. 1731 show commands

Use the **show cfm statisticdelay_measurement domain** command to display delay measurement statistics. If the command is issued gain, the output is replaced with the new values.

```
device#show cfm statistics delay_measurement domain md2 ma ma2 rmep-id 2
Domain: md2 Level: 7
Maintenance association: ma2 VLAN ID: 2 Priority: 7
================================================================================
Round Trip Frame Delay Time : min =  31.501 us  avg =  32.586 us  max =  34.052 us
Round Trip Frame Delay Variation : min = 45 ns     avg =     839 ns     max =   1.875 us
================================================================================
```

Port Used to transmit delay_measurement: 2/2

Number of delay_measurement frames Used to calculate Statistics: 10

Syntax: show cfm statistics delay_measurement domain *domain-name* **ma** *ma-name* **rmep** *rmep-id*

The **domain** *domain-name* parameter specifies the maintenance domain to be used for a delay measurement message. The *domain-name* attribute is case-sensitive.

The **ma** *ma-name* parameter specifies the maintenance association to be used for a delay measurement message. The *ma-name* attribute is case-sensitive.

The **rmep** *rmep-id* parameter specifies the remote mep id to be used for a delay measurement message.

Sample configuration

1. MEP configuration (prerequisite for ETH-DM to work).

   Requester-A :

```
device(config)#cfm-enable
device(config-cfm)# domain-name md2 level 7
device(config-cfm-md-md2)#  ma-name ma2 vlan-id 2 priority 7
device(config-cfm-md-md2-ma-ma2)# mep 3 down port ethe 2/2
device(config-cfm-md-md2-ma-ma2)#
Responder-B:
device(config)#cfm-enable
device(config-cfm)# domain-name md2 level 7
device(config-cfm-md-md2)#  ma-name ma2 vlan-id 2 priority 7
device(config-cfm-md-md2-ma-ma2)# mep 2 down port ethe 2/2
device(config-cfm-md-md2-ma-ma2)#
```

2. Issue the **cfm delay_measurement** command.

```
device# cfm delay_measurement domain md2 ma ma2 src-mep 3 target-mep 2
Y1731: Sending 10  delay_measurement to 0000.00f7.3931, timeout 1000 msec
Type Control-c to abort
Reply from 0000.00f7.3931: time= 32.131 us
Reply from 0000.00f7.3931: time= 31.637 us
Reply from 0000.00f7.3931: time= 32.566 us
Reply from 0000.00f7.3931: time= 34.052 us
Reply from 0000.00f7.3931: time= 33.376 us
Reply from 0000.00f7.3931: time= 31.501 us
Reply from 0000.00f7.3931: time= 33.016 us
Reply from 0000.00f7.3931: time= 32.537 us
Reply from 0000.00f7.3931: time= 32.492 us
Reply from 0000.00f7.3931: time= 32.552 us
sent = 10 number = 10 A total of 10 delay measurement replies received.
Success rate is 100 percent (10/10)
================================================================================
Round Trip Frame Delay Time      : min =  31.501 us  avg =  32.586 us  max =  34.052 us
Round Trip Frame Delay Variation : min =      45 ns  avg =     839 ns  max =   1.875 us
================================================================================
```

3. Issue the **show cfm statistic delay_measurement domain** command.

```
device#show cfm statistics delay_measurement domain md2 ma ma2 rmep-id 2
Domain: md2 Level: 7
Maintenance association: ma2 VLAN ID: 2 Priority: 7
================================================================================================
Round Trip Frame Delay Time : min =  31.501 us  avg =  32.586 us  max =  34.052 us
Round Trip Frame Delay Variation : min = 45 ns     avg =     839 ns     max =   1.875 us
================================================================================================
Port Used to transmit delay_measurement: 2/2
Number of delay_measurement frames Used to calculate Statistics: 10
```

CFM monitoring and show commands

Sending linktrace messages

The **cfm linktrace domain** command sends a linktrace message to a specified MEP in the domain. Enter a command such as the following to send a linktrace message to a specified MEP in the domain.

```
device# cfm linktrace domain VPLS-SP ma ma_1 src-mep 21 target-mep 1 timeout 10 ttl 4
Linktrace to 0000.00fb.5378 on Domain VPLS-SP, level 4: timeout 10ms, 4 hops
```

```
-----------------------------------------------------------------
Hops           MAC             Ingress     Ingress Action   Relay Action
            Forwarded          Egress      Egress Action    Nexthop
-----------------------------------------------------------------
   1    0000.00e2.6ea0                                        RLY_FDB
            Forwarded               5/4           EgrOK
   2    0000.00fb.5378              7/2           IgrOK       RLY_HIT
        Not Forwarded
Destination 0000.00fb.5378 reached
```

Syntax: [no] cfm linktrace domain *name* **ma** *ma-name* **src- mep** *mep-id* **target-mip** *HH:HH:HH:HH:HH:HH* | **target-mep** *mep-id* } [**timeout** *timeout*] [**ttl** *TTL*]

The **domain** *name* parameter specifies the maintenance domain to be used for a linktrace message. The *name* attribute is case-sensitive.

The **ma** *ma-name* parameter specifies the maintenance association to be used for a linktrace message. The *ma-name* attribute is case-sensitive.

The **src-mep** *mep-id* parameter specifies the Source ID in the range 1 - 8191.

The **target-mip** *HH:HH:HH:HH:HH:HH* parameter specifies the MAC-address of the MIP linktrace destination.

The **target-mep** *mep-id* parameter specifies the Destination ID of the linktrace destination.

The **timeout** *timeout* parameter specifies the time to wait for a linktrace reply. The range is 1 - 30 seconds.

The **ttl** *TTL* parameter specifies the initial TTL field value in the range 1 - 64.The default is 8 seconds.

Sending loopback messages

The **cfm loopback domain** command, sends a loopback message to a specific MIP in a specified domain.

```
device#cfm loopback domain VPLS-SP ma ma_1 src-mep 2 target-mep 1 timeout 10 number 10
cfm: Sending 10  Loopback to 0000.00fb.5378, timeout 10 msec
Type Control-c to abort
Reply from 0000.00fb.5378: time=1ms
Reply from 0000.00fb.5378: time<1ms
Reply from 0000.00fb.5378: time<1ms
Reply from 0000.00fb.5378: time<1ms
Reply from 0000.00fb.5378: time<1ms
Reply from 0000.00fb.5378: time<1ms
Reply from 0000.00fb.5378: time<1ms
Reply from 0000.00fb.5378: time<1ms
Reply from 0000.00fb.5378: time<1ms
Reply from 0000.00fb.5378: time<1ms
A total of 10 loopback replies received.
Success rate is 100 percent (10/10), round-trip min/avg/max=0/0/1 ms.
```

Syntax: [no] cfm loopback domain *name* **ma** *ma-name* **scr-mep** *mep-id* { **target-mip** *HH:HH:HH:HH:HH:HH* | **target-mep** *mep-id* } [**number** *number*] [**timeout** *timeout*]

The **domain** *name* parameter specifies the maintenance domain to be used for a linktrace message. The*name* attribute is case-sensitive.

The **ma** *ma-name* parameter specifies the maintenance association to be used for a linktrace message. The *ma-name* attribute is case-sensitive.

The **src-mep** *mep-id* parameter specifies the Source ID in the range *1-8191* .

The **dst- mip** *HH:HH:HH:HH:HH:HH* parameter specifies the MAC address of the MIP linktrace destination.

The **target-mep** *mep-id* parameter specifies the Destination ID in the range *1-8191* .

The **number** *number* parameter specifies the number of loopback messages to be sent.

The **timeout** *timeout* parameter specifies the timeout used to wait for linktrace reply.

Displaying CFM configurations

The **show cfm** command, displays the current configuration and status of CFM. For the **show cfm** command to take effect, CFM should first be enabled in Protocol Configuration mode.

```
device#show cfm
Domain: md2
Index: 1
Level: 6
  Maintenance association: ma2
  Ma Index: 1
  CCM interval: 10000 ms
  VLAN ID: 2
  Priority: 6
  MEP   Direction  MAC             PORT
  ====  =========  =========       ====
     3  DOWN       0000.00f7.3831  ethe 2/2
```

Syntax: show cfm [**domain** *name*] [**ma** *ma-name*]

The **domain** *name* parameter specifies a domain for display. By default, all defined domains are shown.

The **ma** *ma-name* parameter specifies the maintenance association name. By default, all defined domains are shown.

TABLE 8 Show CFM output descriptions

Description	Field
Domain	The Domain is the network or the part of the network for which faults in connectivity are displayed.
Index	The Domain Index.
Level	The level is the domain level in the range *0-7*. The levels can be: • Customer's MD levels: 5 - 7 • Provider's MD levels: 3 - 4 • Operator's MD levels: 0 - 2
Maintenance Association	The maintenance association name.
Ma Index	The Maintenance Association Index.
CCM interval	The time interval between two successive Continuity Check messages (CCMs) that are sent by MEPs in the specified Maintenance Domain.
VLAN ID	The VLAN identifier of the maintenance association.
VPLS ID	The VPLS identifier of the maintenance association.
Priority	The priority of the CCM messages, sent by MEPs, in the range *0-7* .
MEP	The maintenance end point ID
Direction	DIsplays the direction the MEP was sent: Up - The MEP direction away from the monitored VLAN. Down - The MEP direction is towards the monitored VLAN.
MAC	DIsplays the associated MAC Address.
PORT	DIsplays the associated port.
MIP	DIsplays the associated MIP
VLAN	DIsplays the associated VLAN.

The **show cfm brief show** command displays a summary of the configured MEPs and RMEPs.

```
device#show cfm brief
Domain: md2
Index: 1
Level: 6  Num of MA: 1
  Maintenance association: ma2
  MA Index: 1
  CCM interval: 10000 ms
  VLAN ID: 2
  Priority: 6
  Num of MEP: 1               Num of RMEP: 1
  rmepstart: 0 rmepfail: 0 rmepok 1
```

Syntax: show cfm [**domain** *name*] [**ma** *ma-name*] **brief**

TABLE 9 Show cfm brief output description

Description	Field
Domain	The Domain is the network or the part of the network for which faults in connectivity are displayed.
Index	The Domain Index.
Level	The level is the domain level in the range *0-7* . The levels can be: • Customer's MD levels: 5 - 7 • Provider's MD levels: 3 - 4 • Operator's MD levels: 0 - 2
Maintenance Association	The maintenance association name.
Ma Index	The Maintenance Association Index.
CCM interval	The time interval between two successive Continuity Check messages (CCMs) that are sent by MEPs in the specified Maintenance Domain.
VLAN ID	The VLAN identifier of the maintenance association.
VPLS ID	The VPLS identifier of the maintenance association.
Priority	The priority of the CCM messages, sent by MEPs, in the range *0-7* .
Numof MEP	The number of MEPs configured.
Num of RMEP	The number of remote MEPs configured
remepstart	The number of RMEPs in the start state.
rmepfail	The number of RMEPs that have failed.
rmepok	The number of RMEPs in an OK state.

Displaying connectivity statistics

The **show cfm connectivity** command displays connectivity statistics for the remote database. For the **show cfm connectivity** command to take effect, CFM should first be enabled in the Protocol Configuration mode.

```
device#show cfm connectivity
Domain: md2 Index: 1
Level: 6
  Maintenance association: ma2
  MA Index: 1
  CCM interval: 10000 ms
  VLAN ID: 2
  Priority: 6
  RMEP  MAC             VLAN/PEER         AGE     PORT     SLOTS
  ====  =====           ==========  =======  ======== =========
     2  0000.00f7.3931        2              20       2/2 2
```

Syntax: show cfm connectivity

TABLE 10 Show CFM connectivity output descriptions

Description	Field
Domain	The Domain is the network or the part of the network for which faults in connectivity are displayed.
Index	The Domain Index.
Level	The level is the domain level in the range *0-7* . The levels can be: • Customer's MD levels: 5 - 7 • Provider's MD levels: 3 - 4 • Operator's MD levels: 0 - 2
Maintenance association	The maintenance association name.
Ma Index	The Maintenance Association Index.
CCM interval	The time interval between two successive Continuity Check messages (CCMs) that are sent by MEPs in the specified Maintenance Domain.
VPLS ID	The VPLS identifier of the maintenance association.
Priority	The priority of the CCM messages, sent by MEPs, in the range *0-7* .
RMEP	The remote maintenance end point ID
MAC	DIsplays the associated MAC Address.
VLAN or VC	VLAN ID or VC label learned from the CCM packet. VC label is in hexadecimal format.
Age	Uptime since RMEP discovery or from last age out
PORT	DIsplays the associated port.
SLOTMASK	Mask of slots that are receiving CCM packets which are used for multi-slot trunks. For example a value of 0005 indicates Slots 1 and 3.

Sample configuration for a customer's domain

FIGURE 3 Sample configuration of a customer's domain

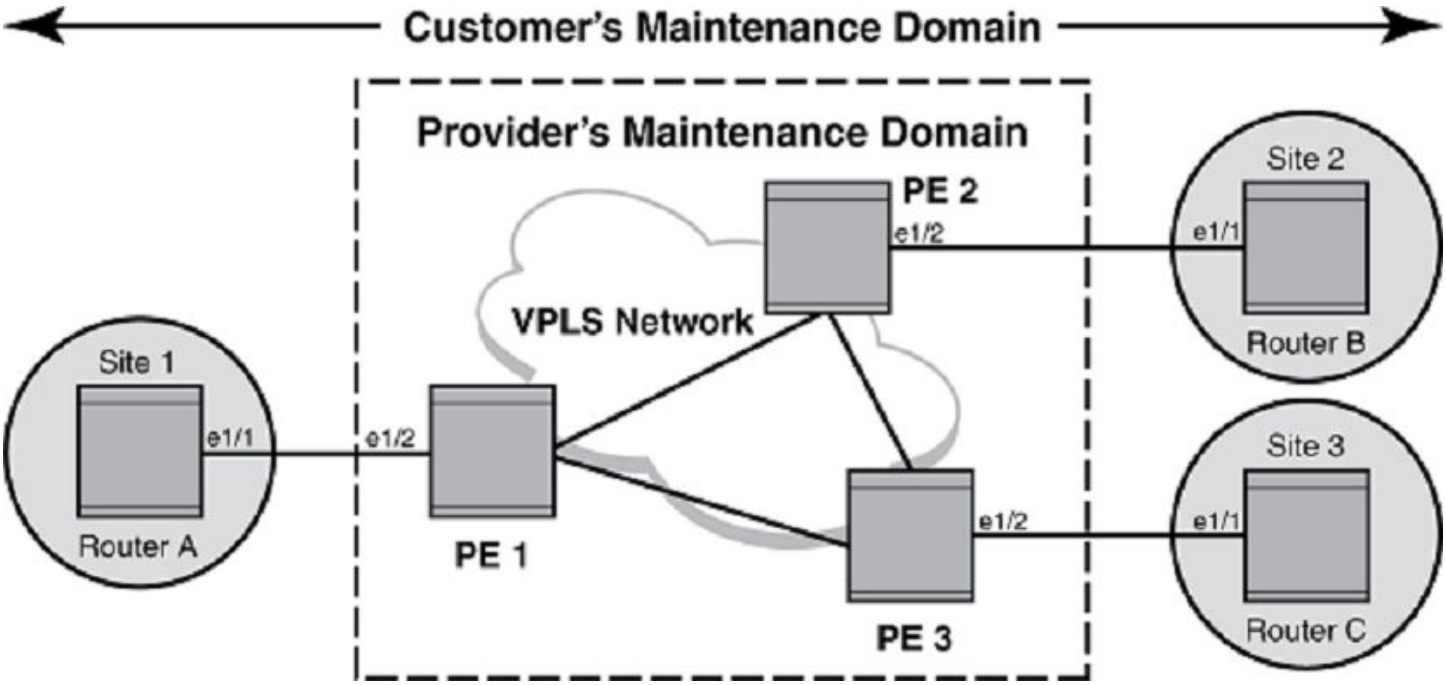

Configuring Router A

CFM configuration steps for Router A are listed below.

1. To enable CFM, enter the following command.

```
RouterA(config)#cfm-enable
```

2. Create a maintenance domain with a specified name **CUST_1** and level **7**.

```
RouterA(config-cfm)#domain-name CUST_1 level 7
```

3. Create a maintenance association within a specified domain of **vlan-id 30** with a priority 3.

```
RouterA(config-cfm-md-CUST_1)#ma-name ma_5 vlan-id 30 priority 3
```

4. Set the time interval between successive Continuity Check Messages to **10-seconds** .

```
RouterA(config-cfm-md-CUST_1-ma-ma_5)#ccm-interval 10-second
```

5. Configuring a MED for each of the Domain Service Access Points of a service instance creates a MA to monitor the connectivity. Add ethernet port **1/1** to a specified maintenance association.

```
RouterA(config-cfm-md-CUST_1-ma-ma_5)#mep 2 down vlan 30 port ethe 1/1
```

Configuring Router B

CFM configuration steps for Router B are listed below.

1. To enable CFM for VPLS, enter the following command.

```
RouterB(config)#cfm-enable
```

2. Create a maintenance domain with a specified name **CUST _1** and level **7** .

```
RouterB(config-cfm)#domain-name CUST_1 level 7
```

3. Create a maintenance association within a specified domain of **vlan-id 30** with a priority **5** .

```
RouterB(config-cfm-CUST_1)#ma-name ma_5 vlan-id 30 priority 5
```

4. Set the time interval between successive Continuity Check Messages to **10-seconds** .

```
RouterB(config-cfm-md-CUST_1-ma-ma_5)#ccm-interval 10-second
```

5. Configuring a MED for each of the Domain Service Access Points of a service instance creates a MA to monitor the connectivity. Add ethernet port **1/1** as MEP to a specified maintenance association.

```
RouterB(config-cfm-md-CUST_1-ma-ma_5)#mep 2 down vlan 30 port ethe 1/1
```

Configuring Router C

CFM configuration steps for Router C are listed below.

1. To enable CFM for VPLS, enter the following command.

```
RouterC(config)#cfm-enable
```

2. Create a maintenance domain with a specified name **CUST_1** and level **7**.

```
RouterC(config-cfm)#domain-name CUST_1 level 7
```

3. Create a maintenance association within a specified domain of **vlan-id 30** with a priority **4**.

```
RouterC(config-cfm-CUST_1)#ma-name ma_5 vlan-id 30 priority 4
```

4. Set the time interval between successive Continuity Check Messages to **10-seconds**.

```
RouterC(config-cfm-md-CUST_1-ma-ma_5)#ccm-interval 10-second
```

5. Configuring a MED for each of the Domain Service Access Points of a service instance creates a MA to monitor the connectivity. Add ethernet port **1/1** as MEP to a specified maintenance association.

```
Routerc(config-cfm-md-CUST_1-ma-ma_5)#mep 1 down vlan 30 port ethe 1/1
```

Configuring CFM using Provider Bridges

Below is an example for configuring CFM when using Provider Bridges configurations as in the figure on Sample configuration for a customer's domain on page 56.

Configuring Router A

CFM configuration steps for Router A are listed below.

1. Create the port-based VLAN that contains the tagged interface that you want to use by entering the following commands.

```
device(config)# vlan30
device(config-vlan-30)# tagged ethe 1/1
```

2. To enable CFM, enter the following command.

```
device(config)# cfm-enable
```

3. Create a maintenance domain with a specified name **CUST_1** and level **7**.

```
device(config-cfm)# domain-name CUST_1 level 7
```

4. Create a maintenance association within a specified ESI **Site1vlan30**, and a **vlan-id 30** with a priority **3** .

```
device(config-cfm-md-CUST_1)# ma-name ma_5 esi Site1vlan30 vlan-id 30 priority 3
```

5. Set the time interval between successive Continuity Check Messages to **10-seconds** .

```
device(config-cfm-md-CUST_1-ma-ma_5)# ccm-interval 10-second
```

6. Configuring a MED for each of the Domain Service Access Points of a service instance creates a MA to monitor the connectivity. Add ethernet port 1/1 to a specified maintenance association.

```
device(config-cfm-md-CUST_1-ma-ma_5)# mep 2 down port ethe 1/1
```

7. To configure the hostnameas **RouterA**, enter a command such as the following.

```
device(config)#hostname RouterA
```

8. Configure interface ethernet 1/1 as the custom-edge by entering the following commands.

```
device(config)# interface ethernet 1/1
device(config-if-e10000-1/1)# port-type customer-edge
device(config-if-e10000-1/1)# enable
device(config-if-e10000-1/1)# end
```

Configuring Router B

CFM configuration steps for Router B are listed below.

1. Create the port-based VLAN that contains the tagged interface that you want to use by entering the following commands.

```
device(config)# vlan30
device(
config-vlan-30)# tagged ethe 1/1
```

2. To enable CFM, enter the following command.

```
device(config)# cfm-enable
```

3. Create a maintenance domain with a specified name **CUST_1** and level **7**.

```
device(config-cfm)# domain-name CUST_1 level 7
```

4. Create a maintenance association within a specified ESI **Site2vlan30**, and a **vlan-id 30** with a priority **3**.

```
device(config-cfm-md-CUST_1)# ma-name ma_5 esi Site2vlan30 vlan-id 30 priority 3
```

5. Set the time interval between successive Continuity Check Messages to **10-seconds**.

```
device(config-cfm-md-CUST_1-ma-ma_5)# ccm-interval 10-second
```

6. Configuring a MED for each of the Domain Service Access Points of a service instance creates a MA to monitor the connectivity. Add ethernet port 1/2 to a specified maintenance association.

```
device(config-cfm-md-CUST_1-ma-ma_5)# mep 2 down port ethe 1/2
```

7. To configure the hostnameas **RouterB** , enter a command such as the following.

```
device(config)# hostname RouterB
```

8. Configure interface ethernet 1/1 as the custome-edge by entering the following commands.

```
device(config)#interface ethernet 1/1
device(config-if-e10000-1/1)# port-type customer-edge
device(config-if-e10000-1/1)# enable
device(config-if-e10000-1/1)# end
```

Configuring Router C

CFM configuration steps for Router C are listed below.

1. Create the port-based VLAN that contains the tagged interface that you want to use by entering the following commands.

```
device(config)# vlan30
device(
config-vlan-30)# tagged ethe 1/1
```

2. To enable CFM, enter the following command.

```
device(config)# cfm-enable
```

3. Create a maintenance domain with a specified name**CUST_1** and level **7**.

```
device(config-cfm)# domain-name CUST_1 level 7
```

4. Create a maintenance association within a specified ESI **Site3vlan30** , and a vlan-id 30 with a priority **3** .

```
device(config-cfm-md-CUST_1)# ma-name ma_5 esi Site3vlan30 vlan-id 30 priority 3
```

5. Set the time interval between successive Continuity Check Messages to **10-seconds** .

```
device(config-cfm-md-CUST_1-ma-ma_5)# ccm-interval 10-second
```

6. Configuring a MED for each of the Domain Service Access Points of a service instance creates a MA to monitor the connectivity. Add ethernet port 1/2 to a specified maintenance association.

```
device(config-cfm-md-CUST_1-ma-ma_5)# mep 2 down port ethe 1/2
```

7. To configure the hostnameas **RouterC**, enter a command such as the following.

```
device(config)# hostname RouterC
```

8. Configure interface ethernet 1/1 as the **custome-edge** by entering the following commands.

```
device(config)# interface ethernet 1/1
device(config-if-e10000-1/1)# port-type customer-edge
device(config-if-e10000-1/1)# enable
device(config-if-e10000-1/1 # end
```

Provider Bridge Extreme 1

1. Create the port-based VLAN that contains the tagged interface that you want to use by entering the following commands.

```
device(config)# vlan30
device(config-vlan-30)# tagged ethe 1/1
```

2. Create the ESI Extreme **vlan300** as an encapsulated SVLAN with the ESI client **Site1vlan30** by entering the following commands.

```
device(config)# esi Extremevlan300 encapsulation svlan#
device(config)# esi-client Site1vlan30
```

3. Add the port-based **VLAN300** that contains the tagged interfaces that you want to use by entering the following commands.

```
device(config)# vlan300
device(config-vlan-300)# tagged ethe 1/1 ethe 1/3
```

4. To enable CFM, enter the following command.

```
device(config)# cfm-enable
```

5. Create a maintenance domain with a specified name **CUST_1** and level **5** .

```
device(config-cfm)# domain-name CUST_1 level 5
```

6. Create a maintenance association within a specified ESI **Site1vlan30** , and a **vlan-id 30** with a priority **3** .

```
device(config-cfm-md-CUST_1)# ma-name ma_5 esi Site1vlan30 vlan-id 30 priority 3
```

7. Set the time interval between successive Continuity Check Messages to **10-seconds** .

```
device(config-cfm-md-CUST_1-ma-ma_5)# ccm-interval 10-second
```

8. Configuring a MED for each of the Domain Service Access Points of a service instance creates a MA to monitor the connectivity. Add ethernet port **1/2** as MEP to a specified maintenance association.

```
device(config-cfm-md-CUST_1-ma-ma_5)# mep 4 up port ethe 1/2
```

9. To configure the hostname as device, enter a command such as the following.

```
device(config)# hostname Extreme
```

10. Configure interface ethernet **1/1** as the **provider network** by entering the following commands.

```
device(config)# interface ethernet 1/1
device(config-if-e10000-1/1)# port-type provider-network
device(config-if-e10000-1/1)enable
device(config-if-e10000-1/1)end
```

11. Configure interface ethernet **1/2** as the **customer-edge** by entering the following commands.

```
device(config)# interface ethernet 1/2
device(config-if-e10000-1/2)# port-type custonmer-edge
device(config-if-e10000-1/2)# enable
device(config-if-e10000-1/2)# end
```

12. Configure interface ethernet **1/3** as the **provider network** by entering the following commands.

```
device(config)# interface ethernet 1/3
device(config-if-e10000-1/3)# port-type provider-network
device(config-if-e10000-1/3)# enable
device(config-if-e10000-1/3)# end
```

Provider Bridge Extreme2

1. Create the port-based **VLAN300** that contains the tagged interfaces that you want to use by entering the following commands.

```
device(config)# vlan300
device(config-vlan-300)# tagged ethe 1/1 ethe 1/3
```

2. To enable CFM, enter the following command.

```
device(config)# cfm-enable
```

3. Create a maintenance domain with a specified name **CUST_1** and level **5** .

```
device(config-cfm)# domain-name CUST_1 level 5
```

4. Create a maintenance association within a specified ESI **Site2vlan30** , and a **vlan-id 30** with a priority **3** .

```
device(config-cfm-md-CUST_1)# ma-name ma_5 esi Site2vlan30 vlan-id 30 priority 3
```

5. Set the time interval between successive Continuity Check Messages to **10-seconds** .

```
device(config-cfm-md-CUST_1-ma-ma_5)# ccm-interval 10-second
```

6. Configuring a MED for each of the Domain Service Access Points of a service instance creates a MA to monitor the connectivity. Add ethernet port **1/2** as MEP to a specified maintenance association.

```
device(config-cfm-md-CUST_1-ma-ma_5)# mep 5 up port ethe 1/2
```

7. To configure the hostname as device**1** , enter a command such as the following.

```
device(config)# hostname device1
```

8. Configure interface ethernet **1/1** as the **provider network** by entering the following commands.

```
device(config)# interface ethernet 1/1
device(config-if-e10000-1/1)# port-type provider-network
device(config-if-e10000-1/1)enable
device(config-if-e10000-1/1)end
```

9. Configure interface ethernet **1/2** as the **customer-edge** by entering the following commands.

```
device(config)# interface ethernet 1/2
device(config-if-e10000-1/2)# port-type custonmer-edge
device(config-if-e10000-1/2)enable
device(config-if-e10000-1/2)end
```

10. Configure interface ethernet **1/3** as the **provider network** by entering the following commands.

```
device(config)# interface ethernet 1/3
device(config-if-e10000-1/3)# port-type provider-network
device(config-if-e10000-1/3)enable
device(config-if-e10000-1/3)end
```

Provider Bridge Extreme3

1. Create the port-based VLAN that contains the tagged interface that you want to use by entering the following commands.

```
device(config)# vlan30
device(config-vlan-30)# tagged ethe 1/2
```

2. Create the ESI device**vlan300** as an encapsulated SVLAN with the ESI client **Site3vlan30** by entering the following commands.

```
device(config)# esi Extreme3vlan300 encapsulation svlan
device(config)# esi-client Site3vlan30
```

3. Add the port-based **VLAN300** that contains the tagged interfaces that you want to use by entering the following commands.

```
device(config)# vlan300
device(config-vlan-300)# tagged ethe 1/1 ethe 1/3
```

4. To enable CFM, enter the following command.

```
device(config)# cfm-enable
```

5. Create a maintenance domain with a specified name **CUST_1** and level **5** .

```
device(config-cfm)# domain-name CUST_1 level 5
```

6. Create a maintenance association within a specified ESI **Site3vlan30** , and a **vlan-id 30** with a priority **3** .

```
device(config-cfm-md-CUST_1)# ma-name ma_5 esi Site3vlan30 vlan-id 30 priority 3
```

7. Set the time interval between successive Continuity Check Messages to **10-seconds** .

```
device(config-cfm-md-CUST_1-ma-ma_5)# ccm-interval 10-second
```

8. Configuring a MED for each of the Domain Service Access Points of a service instance creates a MA to monitor the connectivity. Add ethernet port **1/2** as MEP to a specified maintenance association.

```
device(config-cfm-md-CUST_1-ma-ma_5)# mep 6 up port ethe 1/2
```

9. To configure the hostname as Extreme**3** , enter a command such as the following.

```
device(config)# hostname Extreme3
Configure interface ethernet 1/1
as the provider network
by entering the following commands:
device(config)# interface ethernet 1/1
device(config-if-e10000-1/1)# port-type provider-network
device(config-if-e10000-1/1)# enable
device(config-if-e10000-1/1)# end
```

10. Configure interface ethernet **1/2** as the **customer-edge** by entering the following commands.

```
device(config)# interface ethernet 1/2
device(config-if-e10000-1/2)# port-type custonmer-edge
device(config-if-e10000-1/2)# enable
device(config-if-e10000-1/2)# end
```

11. Configure interface ethernet **1/3** as the **provider network** by entering the following commands.

```
device(config)# interface ethernet 1/3
device(config-if-e10000-1/3)# port-type provider-network
device(config-if-e10000-1/3)# enable
device(config-if-e10000-1/3)# end
```

Displaying the connectivity status in a customer's domain

The following output are for 3 VPLS CEs. The 3 CEs are monitoring Ethernet LAN service in VLAN 30. The Ethernet SP is providing transport service for the customer's VLAN 30 through VPLS which is transparent to customer. The customer is only concerned about RMEPs from remote sites.

```
device# show cfm connectivity
Domain: CUST_1 Level: 7
  Maintenance association: ma_5
  CCM interval: 10000 ms
  VLAN ID: 30
  Priority: 3
  RMEP  MAC             VLAN/PEER       AGE   PORT     SLOTS
  ====  =====           ========== ====  ======== ====
   400  0000.00e2.8a00      30          879      1/2 1,
   200  0000.00f5.e500      30         1550      1/2 1,
device# show cfm connectivity
Domain: CUST_1 Level: 7
  Maintenance association: ma_5
  CCM interval: 10000 ms
  VLAN ID: 30
  Priority: 5
  RMEP  MAC              VLAN/PEER       AGE   PORT     SLOTS
  ====  =====           ========== ====  ======== ====
   400  0000.00e2.8a00       30          898       1/3 1,
   100  0000.00e2.b400       30         1569       1/3 1,
device# show cfm connectivity
Domain: CUST_1 Level: 7
  Maintenance association: ma_5
  CCM interval: 10000 ms
```

```
VLAN ID: 30
Priority: 4
RMEP  MAC              VLAN/PEER       AGE   PORT     SLOTS
====  =====            ================ ====  ======== ====
 200  0000.00f5.e500       30           907      1/4 1,
 100  0000.00e2.b400       30           904      1/4 1,
```

Sample configuration for a customer domain using MPLS VLL

The topology inside an MPLS networks can be managed by using LSP ping and LSP trace route to detect and diagnose LSP failures. CFM packets are Ethernet packets with well know CFM etype and are not shown in the MPLS cloud. Therefore, the topology inside MPLS cannot be managed by the CFM protocol. However, you can use CFM to monitor the health of a VPLS or VLL instances.

FIGURE 4 Sample configuration of a customer domain using MPLS VLL

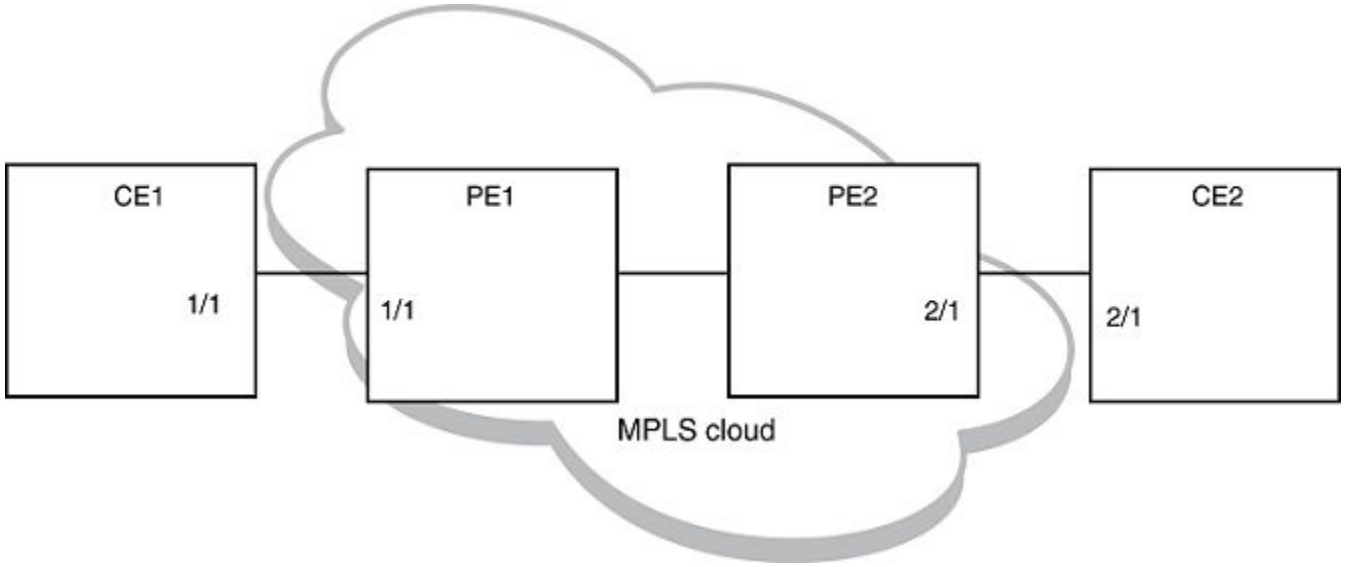

Figure 4 shows a deployment scenario where CE1 and CE2 are customer devices and PE1 and PE2 are provider routers. Port 1/1 on PE1 and port 2/1 on PE2 are VLL-end points. Port 1/1 on PE1 is connected to port 1/1 on CE1 and port 2/1 on PE2 is connected to port 2/1 on CE2.

Achieving end-to-end connectivity between CE1 and CE2

To achieve end-to-end connectivity between CE1 and CE2, configure DOWN MEP on 1/1 and 2/1. PE1 and PE2 acts as MIP.

Configuring CE1

1. To enable CFM, enter the following command.

```
device(config)#cfm-enable
```

2. Create a maintenance domain with a specified name CUST_1 and level 7.

```
device(config-cfm)#domain-name CUST_1 level 7
```

3. Create a maintenance association within a specified domain of vlan-id 30 with a priority 3.

```
device(config-cfm-md-CUST_1)#ma-name ma_5 vlan-id 30 priority 3
```

4. Set the time interval between successive Continuity Check Messages.The default is 10-seconds.

```
device(config-cfm-md-CUST_1-ma-ma_5)#ccm-interval 10-second
```

5. Configure a MEP on port 1/1 and vlan 30.

```
device(config-cfm-md-CUST_1-ma-ma_5)#mep 1 down vlan 30 port ethe 1/1
```

6. 6. Configure a remote-mep.

```
device(config-cfm-md-CUST_1-ma-ma_5)#remote-mep 2 to 2
```

Configuring CE2

1. To enable CFM, enter the following command.

```
device(config)#cfm-enable
```

2. Create a maintenance domain with a specified name CUST_1 and level 7.

```
device(config-cfm)#domain-name CUST_1 level 7
```

3. Create a maintenance association within a specified domain of vlan-id 30 with a priority 3.

```
device(config-cfm-md-CUST_1)#ma-name ma_5 vlan-id 30 priority 3
```

4. Set the time interval between successive Continuity Check Messages.The default is 10-seconds.

```
device(config-cfm-md-CUST_1-ma-ma_5)#ccm-interval 10-second
```

5. Configure a MEP on port 2/1 and vlan 30.

```
device(config-cfm-md-CUST_1-ma-ma_5)#mep 2 down vlan 30 port ethe 2/1
```

6. Configure a remote-mep.

```
device(config-cfm-md-CUST_1-ma-ma_5)#remote-mep 1 to 1
```

MPLS Configurations on PE1

Before configuring CFM on PE1, the MPLS Configuration on PE1 must be done.

Enter the following commands to configure VLL peers from PE1 to PE 2.

1. To create a VLL instance, enter commands such as the following.

```
device(config)#router mpls
device(config-mpls)vll pe1-to-pe2 30
```

2. To specify a VLL peer, enter a command such as the following.

```
device(config-mpls-vll)vll-peer 10.1.1.2
```

3. To specify an un-tagged endpoint for a VLL instance, enter the following commands.

```
device(config-mpls-vll)untagged ethe 1/1
```

   Tagged ports are configured under a VLAN ID.

4. To specify a tagged endpoint for a VLL instance, enter the following commands.

```
device(config-mpls-vll)vlan 30
device(config-mpls-vll-vlan)tagged ethe 1/1
```

IEEE 802.1ag Configuration on PE1

If the VLL configuration is not done prior to configuring the maintenance association, then the maintenance association is not allowed.

1. To enable CFM, enter the following command.

```
device(config)#cfm-enable
```

2. Create a maintenance domain with a specified name CUST_1 and level 7.

```
device(config-cfm)#domain-name CUST_1 level 7
```

3. Create a maintenance association within a specified domain of vll-id 30 with a priority 3.

```
device(config-cfm-md-CUST_1)#ma-name ma_5 vll-id 30 priority 3
```

4. Set the time interval between successive Continuity Check Messages.The default is 10-seconds.

```
device(config-cfm-md-CUST_1-ma-ma_5)#ccm-interval 10-second
```

   In the configuration, a MIP (Maintenance Intermediate Point) is created by default on the VLL port. You can also configure an explicit MIP on PE1. In that case, MIP is created on the VLL-port if there is a MEP (Maintenance End Point) created on the port at some lower Maintenance Domain Level.

5. To configure an explicit MIP on PE1, enter the following command.

```
device(config-cfm-md-CUST_1-ma-ma_5)# mip-creation explicit
```

6. To change back to default use the following command.

```
device(config-cfm-md-CUST_1-ma-ma_5)# mip-creation default
```

MPLS Configurations on PE2

Before configuring CFM on PE2, MPLS is configured on PE2.

Use the following steps to configure VLL peers from PE2 to PE 1.

1. To create a VLL instance, enter commands such as the following.

```
device(config)#router mpls
device(config-mpls)vll pe2-to-pe1 30
```

2. To specify a VPLS peer enter a command such as the following.

```
device(config-mpls-vll)vpls-peer 10.1.1.1
```

3. To specify an un-tagged endpoint for a VLL instance, enter the following commands.

```
device(config-mpls-vll)untagged ethe 2/1
```

   Tagged ports are configured under a VLAN ID.

4. To specify a tagged endpoint for a VLL instance, enter the following commands.

```
device(config-mpls-vll)vlan 30
device(config-mpls-vll-vlan)tagged ethe 2/1
```

IEEE 802.1ag Configurations on PE2

If the VLL configuration is not done prior to configuring the maintenance association, then the maintenance association is not allowed.

1. To enable CFM, enter the following command.

```
device(config)#cfm-enable
```

2. Create a maintenance domain with a specified name CUST_1 and level 7.

```
device(config-cfm)#domain-name CUST_1 level 7
```

3. Create a maintenance association within a specified domain of vll-id 30 with a priority 3.

```
device(config-cfm-md-CUST_1)#ma-name ma_5 vll-id 30 priority 3
```

4. Set the time interval between successive Continuity Check Messages.The default is 10-seconds.

```
device(config-cfm-md-CUST_1-ma-ma_5)#ccm-interval 10-second
```

   In the above configuration, MIP is created by default on the VLL-endpoint. You can also configure an explicit-mip on PE2. In that case, MIP is created on the VLL-port if there is a MEP is created on the endpoint at some lower MD Level.

5. To configure an explicit-mip on PE2, enter the following command.

```
device(config-cfm-md-CUST_1-ma-ma_5)# mip-creation explicit
```

6. To change back to default use the following command.

```
device(config-cfm-md-CUST_1-ma-ma_5)# mip-creation default
```

Verifying connectivity using IEEE 802.1ag

Once CE1,CE2,PE1 and PE2 are configured, you can determine the end-to-end connectivity by looking at the remote-mep status by using the following show commands.

```
device#show cfm connectivity
Domain: CUST_1 Level: 7
Maintenance association: ma_5
CCM interval: 10000 ms
VLAN ID: 30
Priority: 3
RMEP MAC                VLAN/PEER       AGE     PORT    SLOTS
==== =============      =======         ====    ====    ====
2    0000.00e2.8a00     30              879     1/2     1,

device#show cfm connectivity domain CUST_1 ma ma_5 rmep-id 2
Domain: CUST_1 Level: 7
Maintenance association: ma_5 VLAN ID: 30 Priority: 3
CCM interval: 10
RMEP MAC               PORT   Oper     Age    CCM    RDI   Port  Intf   Intvl
```

```
Seq
                State             Val                    Cnt           Status Status  Error  Error
==== ===============   ===    ===      =====  ===    ===   ====  ======= ======  ====
2    0000.00e2.8a00    1/1    OK       26000  2600   N     0     0       N       N=
```

Syntax: show cfm connectivity [**domain** *name*] [**ma** *ma-name*]

The **domain** *name* parameter displays the specific domain information. By default, all defined domains are shown.

The **ma** *ma-name* parameter specifies the maintenance association name. By default, all defined domains are shown.

TABLE 11 Show CFM connectivity output descriptions

This field...	Displays...
Domain	The Domain is the network or the part of the network for which faults in connectivity are displayed.
Level	The level is the domain level in the range *0-7*. The levels can be: • Customer's MD levels: 5 - 7 • Provider's MD levels: 3 - 4 • Operator's MD levels: 0 - 2
Maintenance association	The maintenance association name.
CCM interval	The time interval between two successive Continuity Check messages (CCMs) that are sent by MEPs in the specified Maintenance Domain.
VPLS ID	The VPLS identifier of the maintenance association.
Priority	The priority of the CCM messages, sent by MEPs, in the range *0-7* .
RMEP	The remote maintenance end point ID
MAC	DIsplays the associated MAC Address.
VLAN/VC	VLAN ID or VC label learned from the CCM packet. VC label is in hexadecimal format.
Age	Uptime since RMEP discovery or from last age out
PORT	DIsplays the associated port.
SLOTMASK	Mask of slots that are receiving CCM packets which are used for multi-slot trunks. For example a value of 0005 indicates Slots 1 and 3.

Verifying connectivity in a VLL network using IEEE 802.1ag Loopback

You can manually monitor the status of a VLL peer using IEEE 802.1ag CFM Loopback (MAC ping) as shown below.

```
device#cfm loopback domain CUST_1 ma ma_5 src-mep 1 target-mep 2
DOT1AG: Sending 10 Loopback to 0000.00e2.8a00, timeout 10000 msec
Type Control-c to abort
Reply from 0000.00e2.8a00: time=3ms
Reply from 0000.00e2.8a00: time<1ms
Reply from 0000.00e2.8a00: time<1ms
Reply from 0000.00e2.8a00: time<1ms
Reply from 0000.00e2.8a00: time=38ms
Reply from 0000.00e2.8a00: time<1ms
Reply from 0000.00e2.8a00: time<1ms
Reply from 0000.00e2.8a00: time<1ms
Reply from 0000.00e2.8a00: time<1ms
Reply from 0000.00e2.8a00: time<1ms
A total of 10 loopback replies received.
Success rate is 100 percent (10/10), round-trip min/avg/max=0/4/38 ms.
```

Syntax: [no] cfm linktrace domain *name* **ma** *ma-mame* **src- mep** *mep-id* **target-mip** *HH:HH:HH:HH:HH:HH* | **target-mep** *mep-id*} [**timeout** *timeout*] [**ttl** *TTL*]

The **domain** *name* parameter specifies the maintenance domain to be used for a linktrace message. The *name* attribute is case-sensitive.

The **ma** *ma-name* parameter specifies the maintenance association to be used for a linktrace message. The *ma-name* attribute is case-sensitive.

The **src-mep** *mep-id* parameter specifies the Source ID in the range 1 - 8191.

The **target-mip** *HH:HH:HH:HH:HH:HH* parameter specifies the MAC-address of the MIP linktrace destination.

The **target-mep** *mep-id* parameter specifies the ID of the linktrace destination.

The **timeout** *timeout* parameter specifies the time to wait for a linktrace reply. The range is 1 - 30 seconds.

The **ttl** *TTL* parameter specifies the initial TTL field value in the range 1 - 64. The default is 8 seconds.

Verifying Connectivity in a VLL Network Using IEEE 802.1ag Linktrace

You can manually monitor the status of a VLL peer using IEEE 802.1ag CFM Loopback (MAC Ping) as shown below.

Syntax: [no] cfm loopback domain *name* **ma** *ma-name* **scr-mep** *mep-id* { **target-mip** *HH:HH:HH:HH:HH:HH* | **target-mep** *mep-id* } [**number** *number*] [**timeout** *timeout*]

The **domain** *name* parameter specifies the maintenance domain to be used for a linktrace message. The *name* attribute is case-sensitive.

The **ma** *ma-name* parameter specifies the maintenance association to be used for a linktrace message. The *ma-name* attribute is case-sensitive.

The **src-mep** *mep-id* parameter specifies the Source ID in the range *1-8191* .

The **target- mip** *HH:HH:HH:HH:HH:HH* parameter specifies the MAC address of the MIP linktrace destination.

The **target-mep** *mep-id* parameter specifies the Destination ID in the range *1-8191* .

The **number** *number* parameter specifies the number of loopback messages to be sent.

The **timeout** *timeout* parameter specifies the timeout used to wait for linktrace reply.

If the linktrace and loopback to target-mep 2 fails, then the linktrace can be done on the MIPs on PE1 and PE2 to know the exact failure.

Deployment scenario with PEs functioning as DOWN MEP

It is also possible to configure DOWN MEP on VLL end-points. For example, in Sample configuration for a customer domain using MPLS VLL on page 64, the DOWN MEP can be configured to monitor the connectivity between CE1 and PE1 or to monitor the connectivity between CE2 and PE2.

Configuring CE1

1. To enable CFM, enter the following command.

```
device(config)#cfm-enable
```

2. Create a maintenance domain with a specified name CUST_2 and level 6.

```
device(config-cfm)#domain-name CUST_2 level 6
```

3. Create a maintenance association within a specified domain of vlan-id 30 with a priority 3.

```
device(config-cfm-md-CUST_2)#ma-name ma_6 vlan-id 30 priority 3
```

4. Set the time interval between successive Continuity Check Messages.The default is 10-seconds.

```
device(config-cfm-md-CUST_2-ma-ma_6)#ccm-interval 10-second
```

5. Configure a MEP on port 1/1 and vlan 30.

```
device(config-cfm-md-CUST_2-ma-ma_6)#mep 3 down vlan 30 port ethe 1/1
```

6. Configure a remote-mep.

```
device(config-cfm-md-CUST_1-ma-ma_5)#remote-mep 4 to 4
```

Configuring PE1

The MPLS-VLL configuration is the same as shown in the previous deployment scenario. If the VLL configuration is not done prior to configuring maintenance association, the MA configuration will not be allowed. Also the port and vlan in the MEP configuration should exist in the VLL configuration prior to MEP configuration, otherwise it is not allowed. The port in the MEP configuration can be either a tagged or untagged port already present in the VLL configuration.

1. To enable CFM, enter the following command.

```
device(config)#cfm-enable
```

2. Create a maintenance domain with a specified name CUST_2 and level 6.

```
device(config-cfm)#domain-name CUST_2 level 6
```

3. Create a maintenance association within a specified domain of vll-id 30 with a priority 3.

```
device(config-cfm-md-CUST_2)#ma-name ma_6 vll-id 30 priority 3
```

4. Set the time interval between successive Continuity Check Messages. The default is 10-seconds.

```
device(config-cfm-md-CUST_2-ma-ma_6)#ccm-interval 10-second
```

5. Configure a MEP on port 1/1 and vlan 30.

```
device(config-cfm-md-CUST_2-ma-ma_6)#mep 4 down vlan 30 port ethe 1/1
```

   To monitor the connectivity between CE-1 and PE-1, you can use the **show cfm connectivity** commands as mentioned in the previous scenario. You can also use the **loopback** or **linktrace** commands on CE-1 or PE-1.

Deployment scenario with PEs functioning as UP MEP

UP MEPs can also be configured on PEs. This monitors connectivity of VLL end points.

Configuring PE1

The MPLS-VLL configuration is the same as shown in the previous deployment scenario. If the VLL configuration is not done prior to configuring maintenance association, the MA configuration would not be allowed. Also the port and vlan in the MEP configuration should exist in VLL configuration prior to MEP configuration, otherwise it will not be allowed. The port in the MEP configuration can be either a tagged or untagged port already present in VLL configuration.

1. To enable CFM, enter the following command.

```
device(config)#cfm-enable
```

2. Create a maintenance domain with a specified name PROVIDER_1 and level 4.

```
device(config-cfm)#domain-name PROVIDER_1 level 4
```

3. Create a maintenance association within a specified domain of vll-id 30 with a priority 3.

```
device(config-cfm-md-PROVIDER_1)#ma-name ma_8 vll-id 30 priority 3
```

4. Set the time interval between successive Continuity Check Messages. The default is 10-seconds.

```
device(config-cfm-md-PROVIDER_1-ma-ma_8)#ccm-interval 10-second
device(config-cfm-md-PROVIDER_1-ma-ma_8)#mep 6 up vlan 30 port ethe 1/1
```

Configuring PE2

The configuration on PE1 is similar to the PE1 configuration.

1. To enable CFM, enter the following command.

```
device(config)#cfm-enable
```

2. Create a maintenance domain with a specified name PROVIDER_1 and level 4.

```
device(config-cfm)#domain-name PROVIDER_1 level 4
```

3. Create a maintenance association within a specified domain of vll-id 30 with a priority 3.

```
device(config-cfm-md-PROVIDER_1)#ma-name ma_8 vll-id 30 priority 3
```

4. Set the time interval between successive Continuity Check Messages. The default is 10-seconds.

```
device(config-cfm-md-PROVIDER_1-ma-ma_8)#ccm-interval 10-second
device  (config-cfm-md-PROVIDER_1-ma-ma_8)#mep 7 up vlan 30 port ethe 2/1
```

To monitor the connectivity between PE1 and PE2, you could use the **show cfm connectivity** commands as mentioned in the previous scenario. Also you could use either loopback or linktrace on PE1 or PE2.

Configuring PE2

1. To enable CFM, enter the following command.

```
device(config)#cfm-enable
```

2. Create a maintenance domain with a specified name PROVIDER_1 and level 4.

```
device(config-cfm)#domain-name PROVIDER_1 level 4
```

3. Create a maintenance association within a specified domain of vll-id 30 with a priority 3.

```
device(config-cfm-md-PROVIDER_1)#ma-name ma_8 vll-id 30 priority 3
```

4. Set the time interval between successive Continuity Check Messages.The default is 10-seconds.

```
device(config-cfm-md-PROVIDER_1-ma-ma_8)#ccm-interval 10-second
```

5. Configure MEP 4 down on port 1/1 and vlan 30

```
device(config-cfm-md-CUST_2-ma-ma_6)#mep 4 down vlan 30 port ethe 1/1
```

To monitor the connectivity between PE-1 and PE-2, we could use "**show cfm connectivity** " commands as mentioned in the previous scenario. Also we could use either loopback or linktrace on PE-1 or PE-2.

IEEE 802.1ag with VLL-LOCAL

In the case of IEEE 802.1ag over VLL-LOCAL, the PE acts as a MIP and VLL does VLAN translation. As shown in the figure below, MEP is configured on vlan-10 on CE1 and vlan-11 on CE2. On PE1, MIP is configured on VLL-LOCAL and which has vlan-10, port 1/1 and vlan-11,port 2/1 configured as end points.

UP MEP is not be allowed for VLL-Local.

FIGURE 5 IEEE 802.1ag over VLL-LOCAL

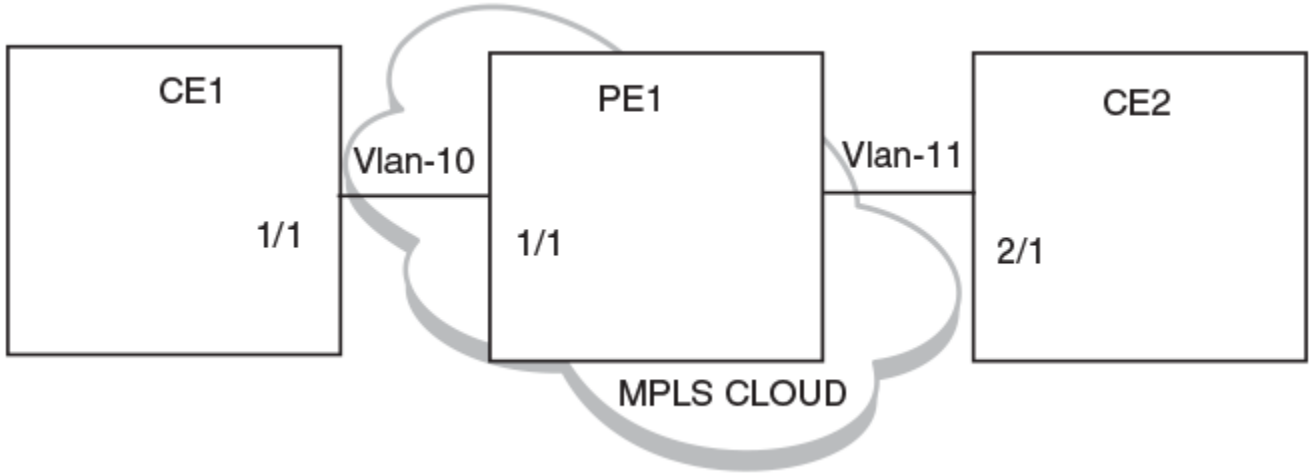

MPLS configurations on PE1

Before configuring CFM on PE1 we need to do MPLS Configuration on PE1.

Enter the following commands to configure VLL peers from PE1 to PE 2.

1. To create a VLL instance, enter commands such as the following.

```
device(config)#router mpls
device(config-mpls)vll-local test1
```

2. To specify an un-tagged endpoint for a VLL instance, enter the following commands.

```
device(config-mpls-vll-test1)untagged ethe 1/1
```

Tagged ports are configured under a VLAN ID.

3. To specify a tagged endpoint for a VLL instance, enter the following commands.

```
device(config-mpls-vll-test1)vlan 30
device(config-mpls-vll-vlan)tagged ethe 1/1
```

As in the previous case, to monitor the connectivity between CE1 and CE2, you can use the **show cfm connectivity** command as mentioned in the previous scenario. You can also use either loopback or linktrace on CE1 or CE2.

LAG-support for IEEE 802.1ag-over-vll

FIGURE 6 IEEE 802.1ag over VLL

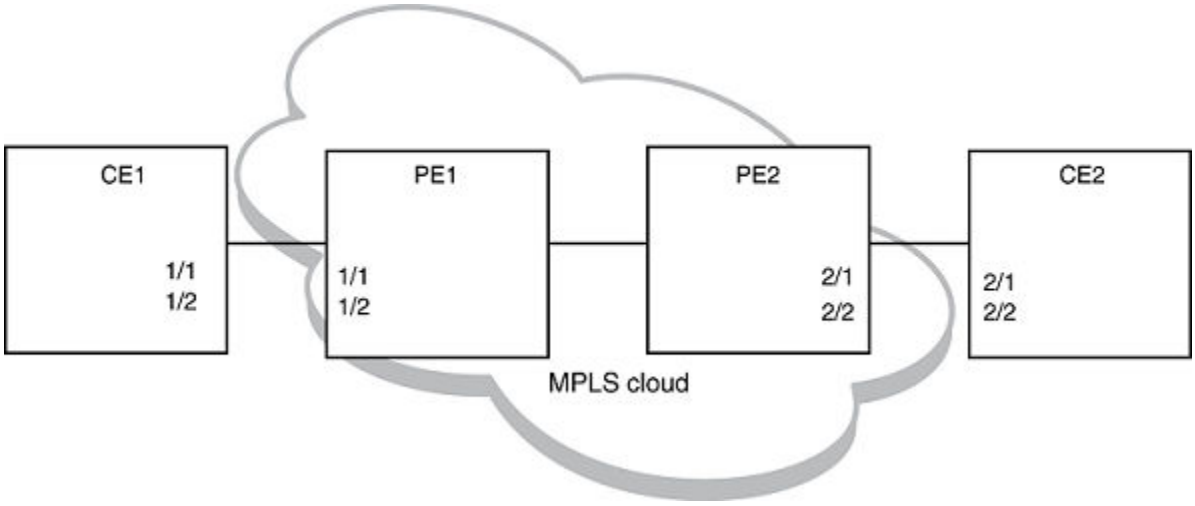

As shown in the previous figure, you can have MEP configuration over a LAG port. On CE1 and CE2 DOWN MEP is configured on VLAN and on PE1 and PE2 DOWN or UP MEP would be configured, depending on what to monitor.

The configuration and monitoring of MEPs is similar as mentioned in the previous examples.

Deletion of VLL

NOTE
Deletion of VLL would cause the deletion of Maintenance Association and corresponding MEPs on that MA.

Sub-second timer support

The **ccm-interval** command sets the time interval between two successive Continuity Check messages (CCMs) that are sent by MEPs in the specified Maintenance Domain. The default value is 10 seconds. There is support for sub-second timers 3.3 ms, 10 ms and 100 ms. As in the case of VLAN and VPLS, for sub-second timers pbif hardware assist is used to transmit and process the CCM packets.

NOTE
The sub-second timer functionality is not supported on VLL-Local.The sub-second timer functionality is not supported on NetIron CES devices

Hitless upgrade support

Hitless upgrade support for IEEE 802.1ag over VLL is similar to IEEE 802.1ag hitless upgrade support for VLAN or VPLS.

Monitoring the status of devices in a VPLS network in a Provider's Maintenance Domain

CFM provides capabilities to detect, verify, and isolate connectivity failures.

NOTE
When configuring 802.1ag over VPLS, if the VPLS endpoint is deleted from the configuration, the MEP configuration is deleted under CFM without warning.

In the Figure 7, CFM is applied over a VPLS network; ports 1/2 and 1/3 are customer facing networks; and port 1/1 is an uplink to a VPLS cloud.

FIGURE 7 VPLS cloud with CFM enabled

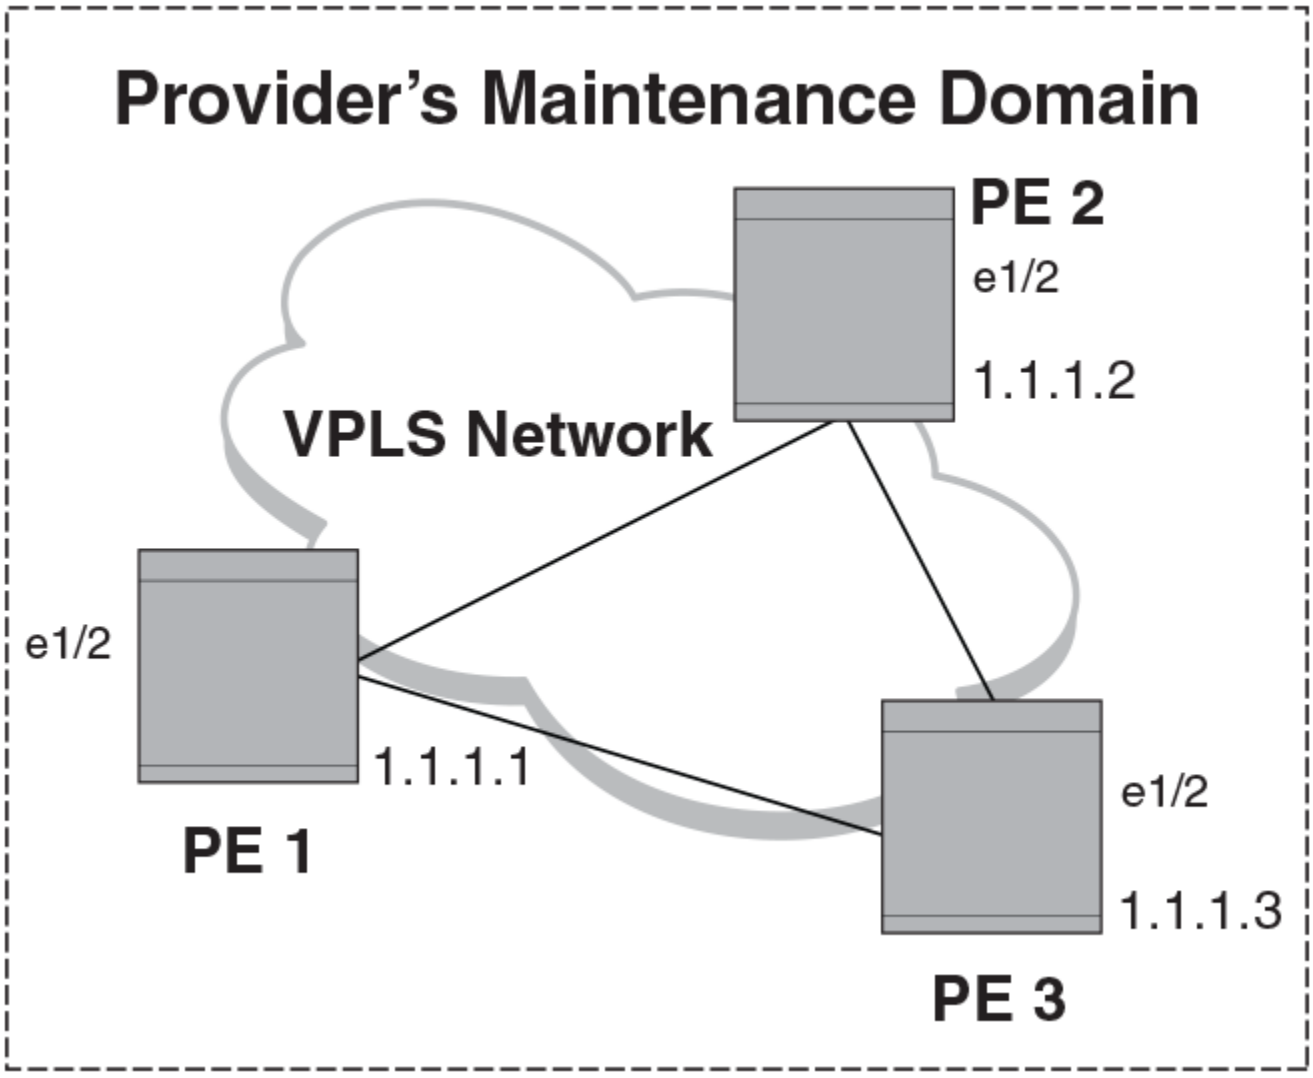

Configuring PE 1

1. To enable CFM for VPLS, enter the following command.

```
device(config)#cfm-enable
```

2. Create a maintenance domain with a specified name **VPLS-SP** and level **4** .

```
device(config-cfm)#domain-name VPLS-SP level 4
```

3. Create a maintenance association within a specified domain of **vpls-id 1** with a priority **3** .

```
device(config-cfm-md-VPLS-SP)#ma-name ma_1 vpls-id 1 priority 3
```

4. Configuring a MED for each of the Domain Service Access Points of a service instance creates a MA to monitor the connectivity. Add ethernet port **1/2** as MEP to a specified maintenance association.

```
device(config-cfm-md-VPLS-SP-ma-ma_1)#mep 1 up vlan 30 port ethe 1/2
```

MPLS configurations

Enter the following commands to configure VPLS peers from PE 2 to PE3.

1. To create a VPLS instance, enter commands such as the following.

```
device(config)#router mpls
device(config-mpls)#vpls 1 1
```

2. To specify two remote VPLS peers within a VPLS instance, enter a commands such as the following.

```
device(config-mpls-vpls-1)#vpls-peer 10.1.1.2
device(config-mpls-vpls-1)#vpls-peer 10.1.1.3
```

3. Tagged ports are configured under a VLAN ID. To specify a tagged endpoint for a VPLS instance, enter the following commands.

```
device(config-mpls-vpls-1)#vlan 30
device(config-mpls-vpls-1-vlan-30)#tagged ethe 1/2 to 1/3
```

Configuring PE 2

CFM configuration steps for Router 2 are listed below.

1. To enable CFM, enter the following command.

```
device(config)#cfm-enable
```

2. Create a maintenance domain with a specified name **VPLS-SP** and level **4** .

```
device(config-cfm)#domain-name VPLS-SP level 4
```

3. Create a maintenance association within a specified domain of **vpls-id 1** with a priority **3** .

```
device(config-cfm-VPLS-SP)#ma-name ma_1 vpls-id 1 priority 3
```

4. Configuring a MED for each of the Domain Service Access Points of a service instance creates a MA to monitor the connectivity. Add ethernet port **1/2** as MEP to a specified maintenance association.

```
device(config-cfm-md-VPLS-SP-ma-ma_1)#mep 2 up vlan 30 port ethe 1/2
```

MPLS configurations

Enter the following commands to configure VPLS peers from PE1 toPE 3.

1. To create a VPLS instance, enter commands such as the following.

```
device(config)#router mpls
device(config-mpls)vpls 1 1
```

2. To specify two remote VPLS peers within a VPLS instance, enter a command such as the following.

```
device(config-mpls-vpls-1)vpls-peer 10.1.1.1
device(config-mpls-vpls-1)vpls-peer 10.1.1.3
```

   Tagged ports are configured under a VLAN ID. To specify a tagged endpoint for a VPLS instance, enter the following commands.

```
device(config-mpls-vpls-1)vlan 30
device(config-mpls-vpls-1-vlan-30)tagged ethe 1/2
```

Configuring PE 3

CFM configuration steps for PE 3 are listed below.

1. To enable CFM for VPLS, enter the following command.

```
device(config)#cfm-enable
```

2. Create a maintenance domain with a specified name **VPLS-SP** and level **4** .

```
device(config-cfm)#domain-name VPLS-SP level 4
```

3. Create a maintenance association within a specified domain of **vpls-id 1** with a priority **3** .

```
device(config-cfm-md-VPLS-SP)#ma-name ma_1 vpls-id 1 priority 3
```

4. Configuring a MED for each of the Domain Service Access Points of a service instance creates a MA to monitor the connectivity. Add ethernet port 1/2 as MEP to a specified maintenance association.

```
device(config-cfm-md-VPLS-SP-ma-ma_1)#mep 3 up vlan 30 port ethe 1/2
```

MPLS configurations

Enter the following commands to configure VPLS peers from Router 1 to Router 2.

1. To create a VPLS instance, enter commands such as the following.

```
device(config)router mpls
device(cconfig-mpls)vpls 1 1
```

2. To specify two remote VPLS peers within a VPLS instance, enter a command such as the following.

```
device(config-mpls-vpls-1)vpls-peer 10.1.1.1
device(config-mpls-vpls-1)vpls-peer 10.1.1.2
```

3. Tagged ports are configured under a VLAN ID. To specify a tagged endpoint for a VPLS instance, enter the following commands.

```
device(config-mpls-vpls-1)vlan 30
device(config-mpls-vpls-1-vlan-30)tagged ethe 1/2
```

Verifying connectivity in a VPLS network using IEEE 802.1ag

To display VPLS IEEE 802.1ag connectivity, enter the following commands.

```
device#show cfm domain VPLS-SP
Domain: VPLS-SP
Level: 4
  Maintenance association: ma_1
  CCM interval: 10
  VPLS ID: 1
  Priority: 3
  MEP   Direction  MAC             PORT
  ====  =========  =========       ====
     1    UP       0000.00e3.8210  ethe 1/3
```

Syntax: show cfm [**domain** *name*] [**ma** *ma-name*]

The **domain** *name* parameter displays the specific domain information. By default, all defined domains are shown.

The **ma** *ma-name* parameter specifies the maintenance association name. By default, all defined domains are shown.

TABLE 12 Output for show CFM domain command

This field...	Displays...
Domain	The Domain is the network or the part of the network for which faults in connectivity are displayed.
Level (Maintenance Domain)	The level is the domain level in the range *0-7* . The levels can be: • Operator's MD levels: 0 - 2 • Provider's MD levels: 3 - 4 • Customer's MD levels: 5 - 7
Maintenance association	The maintenance association name.
CCM interval	The time interval between two successive Continuity Check messages (CCMs) that are sent by MEPs in the specified Maintenance Domain.
VLAN ID	The VLAN identifier of the maintenance association.
VPLS ID	The VPLS identifier of the maintenance association.
Priority	The priority of the CCM messages, sent by MEPs, in the range *0-7* .
MEP	The maintenance end point ID
Direction	DIsplays the direction the MEP was sent: Up - The MEP direction away from the monitored VLAN. Down - The MEP direction is towards the monitored VLAN.
MAC	DIsplays the associated MAC Address.
PORT	DIsplays the associated port.

The **show cfm connectivity** command, displays connectivity statistics for the remote database. For the **show cfm connectivity** command to take effect, CFM should first be enabled in the Protocol Configuration mode.

```
device#show cfm connectivity
  Domain: VPLS-SP Level: 4
  Maintenance association: ma_1
  CCM interval: 10
```

```
  VPLS ID: 1
  Priority: 3
  RMEP     MAC           VLAN/VC AGE   PORT SLOTMASK
  ====    =====          ======= ====  ==== ====
     4  0000.00e2.d80a  00f00a1 2157       0008
     2  0000.00e2.b560  00f00a0 2597       0008
device#show cfm connectivity domain VPLS-SP ma ma_1 rmep-id 2
Domain: VPLS-SP Level: 4
Maintenance association: ma_1 VPLS ID: 1 Priority: 3
CCM interval: 10
  RMEP   MAC             PORT    Oper  Age  CCM  RDI Port   Intf   Intvl Seq
                                 State Val  Cnt      Status Status Error Error
  ====  =============== ======= ===== ===  === ==== ====== ====== ===== =====
     2  0000.00e2.b560  00f00a0   OK  26000 2600   N     0      0     N     N
```

Syntax: show cfm connectivity [**domain** *NAME*] [**ma** *MA NAME*]

The **domain** *NAME* parameter displays information for a specific domain. By default, all defined domains are shown.

The **ma** *MA NAME* parameter specifies the maintenance association name. By default, all defined domains are shown.

TABLE 13 Output for show CFM connectivity command

This field...	Displays...
Domain	The Domain is the network or the part of the network for which faults in connectivity are displayed.
Level	The level is the domain level in the range *0-7* . The levels can be: • Customer's MD levels: 0 - 2 • Provider's MD levels: 3 - 4 • Operator's MD levels: 5 - 7
Maintenance association	The maintenance association name.
CCM interval	The time interval between two successive Continuity Check messages (CCMs) that are sent by MEPs in the specified Maintenance Domain.
VLAN ID	The VLAN identifier of the maintenance association.
VPLS ID	The VPLS identifier of the maintenance association.
Priority	The priority of the CCM messages, sent by MEPs, in the range *0-7* .
RMEP	The remote maintenance end point ID
MAC	DIsplays the associated MAC Address.
PORT	DIsplays the associated port.
Oper State	Defines the state of the port attached to the MEP. Possible values OK and Fail
Age Val	Age of the operational state of the port.
CCM Count	Displays the total number of Continuity Check messages (CCMs) that are sent.
RDI	Remote Defect Indicator
Port Status	The status of a port
Intf Status	The status of the interface
Intvl Error	Displays Y if there has been an interval error and N if no interval errors have been detected.
Seq Error	Displays Y if there has been a sequence error and N if no sequence errors have been detected.

Verifying connectivity in a VPLS network using IEEE 802.1ag Loopback

You can manually monitor the status of VPLS peers using IEEE 802.1ag CFM Linktrace (MAC traceroute) and CFM Loopback (MAC Ping) as shown below.

Syntax: [no] cfm linktrace domain *name* **ma** *ma-name* **src-mep** *mep-id* **target-mip** *HH:HH:HH:HH:HH:HH* | **target-mep***mep-id* } [**timeout** *timeout*] [**ttl** *TTL*]

The **domain** *name* parameter specifies the maintenance domain to be used for a linktrace message. The *name* attribute is case-sensitive.

The **ma** *ma-name* parameter specifies the maintenance association to be used for a linktrace message. The *name* attribute is case-sensitive.

The **src-mep** *mep-id* parameter specifies the Source ID in the range 1-8191.

The **target-mip** *HH:HH:HH:HH:HH:HH* parameter specifies the MAC-address of the MIP linktrace destination.

The **target-mep** *mepid* parameter specifies the ID of the linktrace destination.

The **timeout** *timeout* parameter specifies the timeout used to wait for linktrace reply. The default value is 1-30 seconds.

The **ttl** *TTL* parameter specifies the initial TTL field value in the range 1-64. The default is 8 seconds.

Syntax: [no] cfm loopback domain *name* **ma** *ma-name* **scr-mep** *mep-id* { **target-mip** *HH:HH:HH:HH:HH:HH* | **target-mep** *mep-id* } [**number** *number*] [**timeout** *timeout*]

The **domain** *name* parameter specifies the maintenance domain to be used for a linktrace message. The *name* attribute is case-sensitive.

The **ma** *ma-name* parameter specifies the maintenance association to be used for a linktrace message. The *ma-name* attribute is case-sensitive.

The **src-mep** *mep-id* parameter specifies the Source ID in the range 1-8191.

The **target- mip** *HH:HH:HH:HH:HH:HH* parameter specifies the MAC address of the MIP linktrace destination.

The **target-mep** *mep-id* parameter specifies the Destination ID in the range 1-8191.

The **number** *number* parameter specifies the number of loopback messages to be sent.

The **timeout** *timeout* parameter specifies the timeout used to wait for linktrace reply.

You have to configure MAs with different MD levels to monitor the different endpoints with different

> NOTE
> You have to configure MAs with different MD levels to monitor the different endpoints with different VLAN IDs in the same VPLS instance.

Syslog message

If CFM is configured, a syslog message will be generated when remote MEPs change their states or if there are service cross connections.

Sample Syslog Messages

```
device#
SYSLOG: Jan  7 11:22:55:<9>Router2, DOT1AG: Remote MEP 4 in Domain VPLS-SP, MA ma_1 aged out
SYSLOG: Jan  7 11:23:13:<9>Router2, DOT1AG: Remote MEP 4 in Domain VPLS-SP, MA ma_1 become UP state
```

When a failure is detected within a VPLS cloud, use LSP Ping and Traceroute. Refer to LSP ping and traceroute on page 97 for additional information.

Support for IEEE 802.1ag CFM for Provider Bridges (PB) and Provider Backbone Bridges (PBB)

The device support the following single tagging and double tagging cases:

- MEP (up/down) and MIP on C-VLANs
- MEP (up/down) and MIP on S-VLANs - The ability to change tag-type 88a8 to S-VLANs

The CES 2000 Series supports both of the above capabilities in the following scenarios:

- MEP (up/down) and MIP on C-VLANs
- MEP (up/down) and MIP on S-VLANs -The ability to change tag-type 88a8 to S-VLANs
  - MEP on C-VLANs (extended to both default ESI and non-default ESI)

    NOTE
    The C-VLAN may be a child of another ESI or could be "stand-alone".

- – MEP on S-VLANs in an ESI

    NOTE
    The S-VLAN may be a child of another ESI or could be "stand-alone"

- MIP on standalone C-VLANs and stand-alone S-VLANs on a device (i.e. C-VLANs that are not a client of another ESI or S-VLANs that are not a client of another ESI).

The following configurations are not supported on CES 2000 Series devices:

- Handling MIP on a C-VLAN that is a client of an S-VLAN
- Handling MIP/MEP on a B-VLAN
- Handling MIP on an S-VLAN that is a client of a B-VLAN

802.1ag over PBB sub-second timer support

NOTE
The sub-second timer functionality is not supported on NetIron CES devices.

The **ccm-interval** command sets the time interval between two successive Continuity Check messages (CCMs) that are sent by MEPs in the specified Maintenance Domain. The default value is 10 seconds. There is support for sub-second timers 3.3 ms, 10 ms and 100 ms. The ISID CFM can be used in PBB networks and CFM monitoring between Backbone Edge Bridges.

The sub-second CCM interval is supported for the following scenarios:

- 802.1ag over Regular Vlan
- 802.1ag over ESI VLAN
- 802.1ag over VPLS

The following messages are supported for sub-second CCM interval:

- Loop Back Message and Reply (LBM, LBR)
- Link Trace Message and Reply (LTM, LTR)
- Delay Measurement Message and Reply (DMM, DMR)

Sub-second timer sample configuration

CER1

```
device(config)#cfm-enable
device(config-cfm)#domain-name customer level 7
device(config-cfm-md-customer)#ma-name admin vlan 100 priority 7
device(config-cfm-md-customer-ma-admin)#ccm-interval 100-ms
device(config-cfm-md-customer-ma-admin)#mep 1 down port ethe 1/13
```

CER2

```
device(config)#cfm-enable
device(config-cfm)#domain-name customer level 7
device(config-cfm-md-customer)#ma-name admin vlan 100 priority 7
device(config-cfm-md-customer-ma-admin)#ccm-interval 100-ms
```

CER3

```
device(config)#cfm-enable
device(config-cfm)#domain-name customer level 7
device(config-cfm-md-customer)#ma-name admin vlan 100 priority 7
device(config-cfm-md-customer-ma-admin)#ccm-interval 100-ms
device(config-cfm-md-customer-ma-admin)#mep 1 down port ethe 1/13
```

IEEE 802.3ah EFM-OAM

The IEEE 802.3ah Ethernet in the First Mile (EFM) is supported on the NetIron devices.

The IEEE 802.3ah Ethernet in the First Mile (EFM) standard specifies the protocols and Ethernet interfaces for using Ethernet over access links as a first-mile technology and transforming it into a highly reliable technology.

Using the Ethernet in the First Mile solution, the user will gain broadcast Internet access, in addition to services, such as Layer 2 transparent LAN services, Voice services over Ethernet Access networks, and Video and multicast applications, reinforced by security and Quality of Service control in order to build a scalable network.

The in-band management specified by this standard defines the operations, administration and maintenance (OAM) mechanism needed for the advanced monitoring and maintenance of Ethernet links in the first mile. The OAM capabilities facilitate network operation and troubleshooting. Basic 802.3 frames convey OAM data between two ends of the physical link. EFM OAM is optional and can be disabled on each physical port.

OAM initiatives are classified into three layers: transport, connectivity and service. The transport layer is the collection of forwarding entities and interconnected segments that form a multi-hop Ethernet network, and provide connectivity between devices. The transport layer OAM is specified by the IEEE 802.3ah (Clause 57) and provides single-link OAM capabilities. When OAM is present, two connected OAM sub-layers exchange protocol data units (OAMPDUs). OAM PDUs are standard-size frames that can be sent at a maximum rate of 10 frames per second. This limitation is necessary for reducing the impact on the usable bandwidth. It is possible to send each frame several times in order to increase the probability of reception. A combination of the destination MAC address, the Ethernet type/length field and Subtype allow distinguishing OAM PDU frames from other frames.

OAM functionality is designed to provide reliable service assurance mechanisms for both provider and customer networks.

The IEEE 802.3ah EFM standard offers an opportunity to create the operations, OAM sub-layer within the data-link layer of the OSI protocol stack. The data-link layer provides utilities for monitoring and troubleshooting Ethernet links.

Possible applications

The data-link layer OAM is targeted at last-mile applications and service providers can use it for demarcation point OAM services.

Ethernet Last Mile applications require robust infrastructure that is both passive and active. 802.3ah OAM aims to solve validation and testing problems in such an infrastructure.

Using the Ethernet demarcation service providers can additionally manage the remote device without utilizing an IP layer. This can be done by using link-layer SNMP counters, request and reply, loopback testing and other techniques.

EFM-OAM protocol

The functionality of the EFM-OAM can be summarized under the following categories:

- Discovery: Discovery is the mechanism to detect the presence of an OAM sub-layer on the remote device. During the discovery process, information about OAM entities, capabilities, and configurations are exchanged.
- Link monitoring: This process is used to detect link faults and to provide information about the number of frame errors and coding symbol errors.
- Remote fault detection: Provides a mechanism for an OAM entity to convey error conditions to its peer by way of a flag in the OAMPDUs.
- Remote loopback: This mechanism is used to troubleshoot networks and to isolate problem segments in a large network by sending test segments.

Discovery

Discovery is the first phase of EFM-OAM. At this phase, EFM-OAM identifies network devices along with their OAM capabilities. The Discovery process relies on the Information OAMPDUs. During discovery, the following information is advertised through the TLVs within periodic information OAMPDUs:

- OAM configuration (capabilities): Advertises the capabilities of the local OAM entity. Using this information, a peer can determine what functions are supported and accessible (for example, loopback capability).
- OAM mode: The OAM mode is conveyed to the remote OAM entity. The mode can be either active or passive, and can also be used to determine a device's functionality.
- OAMPDU configuration: This configuration includes the maximum OAMPDU size to delivery. In combination with the limited rate of 10 frames per second, this information can be used to limit the bandwidth allocated to OAM traffic.

Timers

- Two configurable timers control the protocol, one determining the rate at which OAMPDUs are to be sent, and the second controlling the rate at which OAMPDUs are to be received to maintain the adjacency between devices.
- An additional 1-second non-configurable timer is used for error aggregation, which is necessary for the Link Monitoring Process to generate link quality events.
- The timer should generate PDUs in the range of 1 - 10 seconds. The default value is 1 second.
- The Hold timer assumes the peer is dead if no packet is received for a period of 1 - 10 seconds. The default value is 5 seconds.

Flags

Included in every OAMPDU is a flags field, which contains, besides other information, the status of the discovery process. There are three possible values for the status:

- Discovering: Discovery is in progress.
- Stable: Discovery is completed. Once aware of this, the remote OAM entity can start sending any type of OAMPDU.

- Unsatisfied: When there are mismatches in the OAM configuration that prevent OAM from completing the discovery, the discovery process is considered unsatisfied and cannot continue.

Process overview

The discovery process allows local Data Terminating Entity (DTE) to detect OAM on a remote DTE. Once OAM support is detected, both ends of the link exchange state and configuration information (such as mode, PDU size, loopback support, and so on). If both DTEs are satisfied with the settings, OAM is enabled on the link. However, the loss of a link or a failure to receive OAMPDUs for five seconds may cause the discovery process the start over again.

DTEs may be in either active or passive mode. Active mode DTEs instigate OAM communications and can issue queries and commands to a remote device. Passive mode DTEs generally wait for the peer device to instigate OAM communications and respond to, but do not instigate, commands and queries. Rules of what DTEs in active or passive mode can do are discussed in the following sections.

Rules for active mode

A DTE in active mode:

- Initiates the OAM Discovery process
- Sends information PDUs
- May send event notification PDUs
- May send variable request or response PDUs
- May send loopback control PDUs

Exceptions

- A DTE in active mode does not respond to variable request PDUs from DTEs in passive mode
- A DTE in active mode does not react to loopback control PDUs from DTEs in passive mode

Rules for passive mode

A DTE in Passive mode:

- Waits for the remote device to initiate the Discovery process
- Sends Information PDUs
- May send Event Notification PDUs
- May respond to Variable Request PDUs
- May react to received Loopback Control PDUs
- Is not permitted to send Variable Request or Loopback Control OAMPDUs

Link monitoring process

The Link Monitoring Process is used for detecting and indicating link faults under a variety of circumstances. Link monitoring uses the Event Notification OAMPDU, and sends events to the remote OAM entity when there are problems detected on the link. The error events defined in the standard are:

- Errored Symbol Period (errored symbols per second): the number of symbol errors that occurred during a specified period exceeded a threshold. These are coding symbol errors (for example, a violation of 4B/5B coding).

- Errored Frame (errored frames per second): the number of frame errors detected during a specified period exceeded a threshold.
- Errored Frame Period (errored frames per N frames): the number of frame errors within the last N frames has exceeded a threshold.
- Errored Frame Seconds Summary (errored secs per M seconds): the number of errored seconds (one second intervals with at least one frame error) among the last M seconds has exceeded a threshold.

Since 802.3ah OAM does not guarantee the delivery of OAMPDUs, the Event Notification OAMPDU can be sent multiple times to reduce the probability of losing notifications. A sequence number is used to recognize duplicate events. The Link Monitoring Process operates for all the links on which EFM OAM is enabled.

Remote failure indication

Faults in Ethernet that are caused by slowly deteriorating quality are more difficult to detect than completely disconnected links. A flag in the OAMPDU allows an OAM entity to send failure conditions to its peer. The failure conditions are defined as follows:

- Link Fault: The Link Fault condition is detected when the receiver loses the signal. This condition is sent once per second in the Information OAMPDU.
- Dying Gasp: This condition is detected when the receiver goes down. The Dying Gasp condition is considered as unrecoverable. Conditions for dying gasp:
  - Reload or reset from MP
  - Interface disable (admin shutdown)
  - Link-OAM disable on interface (deconfiguration)
  - Crash on the box
- Device power down (incidental or deliberate).
- Critical Event: When a critical event occurs, the device is unavailable as a result of malfunction, and it is to be restarted by the user. The critical events can be sent immediately and continually.

When the dying gasp or critical event occurs, the device driver calls a special API in the EFM-OAM implementation.

The link fault applies only when the physical sublayer is capable of independent transmission and reception.

When a link receives no signal from its peer at the physical layer (for example, if the peer's laser is malfunctioning), the local entity sets this flag to let the peer know that its transmit path is inoperable. The link-down API will be called by the device driver in order to notify the remote device of the link fault.

Because the failure conditions are severe, when they are set in the flag, the OAMPDU is not subject to the normal rate limiting policy.

Remote loopback

An OAM entity can put its remote entity into loopback mode using a loopback control OAMPDU. This helps you ensure quality of links during installation or when troubleshooting. In loopback mode, each frame received is transmitted back on that same port except for OAMPDUs and pause frames. The periodic exchange of OAMPDUs must continue while in the loopback state to maintain the OAM session. The loopback command is acknowledged by responding with an information OAMPDU with the loopback state indicated in the state field.

Enabling and disabling EFM-OAM

The **link-oam** command, in Protocol Configuration mode, enables and disables the EFM-OAM protocol and enters into the EFM-OAM Protocol Configuration mode. The **link-oam disable** and **link-oam enable** commands reset all link-oam parameters to default values.

By default, EFM-OAM is disabled.

To enable EFM-OAM, enter a command such as the following:

```
device(config)link-oam
device(config-link-oam)#enable
```

Syntax: [no] link-oam

Syntax: enable

The **no** form of the command sets the 802.3ah EFM-OAM to the disabled state.

Specifying the timeout value

The timeout command is a hold down timer that specifies the number of seconds before it declares that the other side has stopped sending OAMPDUs.

```
device(config-link-oam) #timeout 10
```

Syntax: [no] timeout *value*

The **no** form of the command restores the default value of 5 OAMPDUs.

The *value* parameter specifies the number of seconds before declaring the remote as down. in the range of 1-10.

Specifying the PDU rate

To set the number of PDUs to be transmitted per second, use the pdu-rate command. The default value is 1.

```
device(config-link-oam) #pdu-rate 10
```

Syntax: [no] pdu-rate *value*

The *value* parameter specifies the number of PDUs in the range of 1-10.

The **no** form of the command restores the default value of 1.

Enabling and disabling the EFM-OAM state on the specified interface

The **ethernet** *slot/port* command in interface configuration mode, enables and disables EFM-OAM on the specified interface and sets its mode to active or passive.

When both peers are in passive mode (abnormal configuration), the information from "Remote Status" is not updated anymore and it may be inaccurate. By default, port state is disabled.

```
device(config-link_oam)# ethernet 2/1 active
```

Syntax: [no] ethernet *slot/port* { **active** | **passive** | **remote-failure** }

When **active** mode is specified, the device can send OAMPDU packets over this port in order to initiate an EFM-OAM discovery process. For the discovery process to be initiated the EFM-OAM protocol must have been enabled.

When **passive** mode is specified, the device cannot use this port to send OAMPDU packets, but can respond if it receives OAMPDUs from remote.

When **remote-failure** mode is specified, the device will be set for the remote-failure action.

The **no** form of the command sets the 802.3ah EFM-OAM to the disabled state.

Enabling an interface to accept remote loopback

NOTE
OAM remote loopback is supported only on the NetIron CES and NetIron CER platform and not supported on NetIron XMR and NetIron MLX platforms.

The **ethernet** *slot/port* allow-loopback command, in interface configuration mode, is used to enable the interface to respond to a loopback request from the remote. This is used for loopback traffic analysis.

```
device(config-link-oam) #ethernet 2/1 allow-loopback
```

Syntax: [no] ethernet *slot/port* **allow-loopback**

The **no** form of the command doesn't allow the interface to respond to the loopback request from the remote.

Defining remote failure actions

By default, on receipt of a remote failure message, the device will only log the event. This can be changed to block an interface on receipt of a remote failure message. The following commands display the three events that the protocol supports.

```
device(config-link-oam)#ethernet 2/1 remote-failure dying-gasp action block-interface
device(config-link-oam)#ethernet 2/1 remote-failure critical-event action block-interface
device(config-link-oam)#ethernet 2/1 remote-failure link-fault action block-interface
```

Syntax: [no] ethernet *slot/port* **remote-failure** *dying-gasp* | *link-fault* | *critical-event* **action** *block-interface*

The **no** form of the command returns to the default state of logging.

Forcing the EFM-OAM remote interface into loopback

The **link-oam remote-loopback ethernet** *slot/port start/stop* command starts and stops the remote loopback on the remote node.

```
device# link-oam remote-loop-back ethernet 1/1 start
```

Syntax: [no]link-oam remote loopback ethernet *slot/portstart/stop*

Display information

The following show commands will display OAM information.

Displaying OAM information

To show OAM information on all OAM enabled ports, enter a command such as the following:

```
device#show link-oam info
Ethernet Link Status     OAM Status      Mode       Local Stable    Remote Stable
1/1      up              up              active     satisfied       satisfied
1/2      up              up              passive    satisfied       satisfied
1/3      up              up              active     satisfied       satisfied
1/4      up              init            passive    unsatisfied     unsatisfied
1/5      down            down            passive    unsatisfied     unsatisfied
1/6      down            down            passive    unsatisfied     unsatisfied
1/7      down            down            passive    unsatisfied     unsatisfied
```

Displaying detailed information from a specific port

To show detailed OAM information, enter a command such as the following:

Syntax: show link-oam info detail ethernet *all* | *slot/port*

To show detailed OAM information on a specific ethernet port, enter a command such as the following:

Syntax: show link-oam info detail [**all** | **ethernet** *slot/port*]

Displaying OAM statistics

To show OAM statistics, enter a command such as the following:

Syntax: show link-oam statistics

DIsplaying detailed OAM statistics

To show detailed OAM statistics, enter a command such as the following:

```
device#show link-oam statistics detail
OAM statistics for Ethernet port: 1/1
        Tx statistics
                information OAMPDUs:                    587
                loopback control OAMPDUs:               0
                variable request OAMPDUs:               0
                variable response OAMPDUs:              0
                unique event notification OAMPDUs:      0
                duplicate event notification OAMPDUs:   0
                oranization specific OAMPDUs:           0
                link-fault records:                     0
                critical-event records:                 0
                dying-gasp records:                     0
        Rx statistics
                information OAMPDUs:                    442
                loopback control OAMPDUs:               0
                variable request OAMPDUs:               0
                variable response OAMPDUs:              0
                unique event notification OAMPDUs:      0
                duplicate event notification OAMPDUs:   0
                oranization specific OAMPDUs:           0
                unsupported OAMPDUs:                    0
                link-fault records:                     0
                critical-event records:                 0
                dying-gasp records:                     0
                discarded TLVs:                         0
                unrecognized TLVs:                      0
OAM statistics for Ethernet port: 1/2
        Tx statistics
                information OAMPDUs:                    440
                loopback control OAMPDUs:               0
                variable request OAMPDUs:               0
                variable response OAMPDUs:              0
                unique event notification OAMPDUs:      0
                duplicate event notification OAMPDUs:   0
                oranization specific OAMPDUs:           0
                link-fault records:                     0
                critical-event records:                 0
                dying-gasp records:                     0
        Rx statistics
                information OAMPDUs:                    441
                loopback control OAMPDUs:               0
                variable request OAMPDUs:               0
                variable response OAMPDUs:              0
                unique event notification OAMPDUs:      0
                duplicate event notification OAMPDUs:   0
                oranization specific OAMPDUs:           0
                unsupported OAMPDUs:                    0
                link-fault records:                     0
                critical-event records:                 0
                dying-gasp records:                     0
```

```
discarded TLVs:                         0
unrecognized TLVs:                      0
```

To show detailed OAM statistics, enter a command such as the following:

Syntax: show link-oam statistics detail ports [**all** | **ethernet** *slot/port*]

This field...	Displays...
Ethernet Port	Indicates if the ethernet port that EFM-OAM is enabled on.
Link Status	Indicates if the physical link is operational or any fault is detected on the link.
OAM Status	Indicates the status of OAM on the link between the local and remote DTEs. The status is enabled if OAM client is satisfied with local and remote settings.
Mode	Indicates if the DTE is in active or passive modes. Active DTEs can start the discovery process and passive ones can only respond.
Local Stable	Indicates the reception of the remote DTE state information and is satisfied with the remote OAM settings.
Remote Stable	Indicates the reception of the local DTE state information at the remote DTE and is satisfied with the local OAM settings.

Ping

Ping is a tool that helps you to verify the Internet connectivity at the IP level. The **ping** command sends an Internet Control Message Protocol (ICMP) echo request to the IP address or selected hostname.

Executing ping

The **ping** command, in the (Enable) mode, pings another device from the device. The device supports IP ping, which you can use to test connectivity to remote hosts. Ping sends an echo request packet to an address and waits for a reply.

The device can execute multiple ping commands at the same time. If you can connect to the device via the console, or through an inbound telnet or SSH session, it should be possible to initiate a ping. This applies to all versions of the ping command described below. The device can also resolve multiple DNS queries simultaneously, which allows multiple ping commands with the **hostname** option to be executed at the same time.

To initiate the device to ping to a target device with the IP address of 10.22.2.33, enter a command such as the following.

```
device# ping 10.22.2.33
```

Syntax: ping ip *address* | *hostname* | **vrf** *instance-name* [**source** *ip address*] [**count** *num*] [**timeout** *msec*] [**ttl** *num*] [**size** *byte*] [**quiet**] [**numeric**] [**no-fragment**] [**verify**] [**data** *1-to-4 bytehex*] [**brief**]

The required parameter is the IP address or the host name of the device.

The **vrf** *instance-name* parameter specifies a VPN routing/forwarding instance as the origin of the ping packets.

The **source** *ip addr* parameter specifies an IP address to be used as the origin of the ping packets.

The **count** *num* parameter specifies how many ping packets the device sends. You can specify from *1 - 4294967296* . The default is 1.

The **timeout** *msec* parameter specifies how many milliseconds the device waits for a reply from the pinged device. You can specify a timeout from *1 - 4294967296* milliseconds. The default is 5000 (5 seconds).

The **ttl** *num* parameter specifies the maximum number of hops. You can specify a TTL from 1 - 255. The default is 64.

The **size** *byte* parameter specifies the size of the ICMP data portion of the packet. This is the payload and does not include the header. You can specify from *0 - 9170* . The default is 16.

The **no-fragment** parameter turns on the "do not fragment" bit in the IP header of the ping packet. This option is disabled by default.

The **quiet** parameter hides informational messages such as a summary of the ping parameters sent to the device and instead displays only messages indicating the success or failure of the ping. This option is disabled by default.

The **verify** parameter verifies that the data in the echo packet (the reply packet) is the same as the data in the echo request (the ping). By default the device does not verify the data.

The **data** *1 - 4 byte hex* parameter lets you specify a specific data pattern for the payload instead of the default data pattern, "abcd", in the packet's data payload. The pattern repeats itself throughout the ICMP message (payload) portion of the packet.

> NOTE
> For numeric parameter values, the CLI does not check that the value you enter is within the allowed range. If you exceed the range for a numeric value, the software rounds the value to the nearest valid value.

The **brief** keyword causes ping test characters to be displayed. The following ping test characters are supported:

! Indicates that a reply was received.

. Indicates that the network server timed out while waiting for a reply.

U Indicates that a destination unreachable error PDU was received.

I Indicates that the user interrupted ping.

Executing ping VRF

> NOTE
> The Ping utilities have been enhanced by adding the **ping vrf** command in release 02.1.00 to help with management of Layer 3 VPNs.

The **ping vrf** command lets you test a specific VPN connection. To use this option, enter the following command.

Syntax: ping vrf *vrf-name ip-address*

The *vrf-name* parameter is the name of the VRF that you want to conduct a ping to.

The *ip-address* parameter is the IP address containing the VRF that you want to conduct a ping to.

Executing ping IPv6

The **ping ipv6** command allows you to verify the connectivity from a device to an IPv6 device by performing an ICMP for IPv6 echo test. As with IPv4, multiple IPv6 ping commands can be executed simultaneously by the device.

For example, to ping a device with the IPv6 address of 2001:db8:847f:a385:34dd::45 from the device, enter the following command.

```
device# ping ipv6 2001:db8:847f:a385:34dd::45
```

Syntax: ping ipv6 *ipv6-address* | *hostname* | **vrf** *instance-name* [**outgoing-interface** [**eth** *slot/port* | **ve** *number*]] [**source** *ipv6-address*] [**count** *number*] [**timeout** *milliseconds*] [**ttl** *number*] [**size** *bytes*] [**quiet**] [**numeric**] [**no-fragment**] [**verify**] [**data** *1-to-4 bytehex*] [**brief**]

The required parameter is the IPv6 address or the host name of the device. The *ipv6-address* parameter specifies the address of the target device. You must specify this address in hexadecimal using 16-bit values between colons, or specify a host name using an ASCII string.

The **vrf** *instance-name* parameter specifies a VPN routing/forwarding instance as the origin of the ping packets.

The **outgoing-interface** keyword specifies a physical interface over which you can verify connectivity. If you specify a physical interface, such as an Ethernet interface, you must also specify the port number of the interface. If you specify a virtual interface, such as a VE, you must specify the number associated with the VE.

> NOTE
> This option is applicable only when the destination IPv6 address is a link local address.

Specify **ethernet** *slot/port*.

The **source** *ipv6-address* parameter specifies an IPv6 address to be used as the origin of the ping packets.

The **count** *number* parameter specifies how many ping packets the sends. You can specify from *1 - 4294967296* . The default is 1.

The **timeout** *milliseconds* parameter specifies how many milliseconds the waits for a reply from the pinged device. You can specify a timeout from *1 - 4294967296* milliseconds. The default is 5000 (5 seconds).

The **ttl** *number* parameter specifies the maximum number of hops. You can specify a TTL from 1 - 255. The default is 64.

The **size** *bytes* parameter specifies the size of the ICMP data portion of the packet. This is the payload and does not include the header. You can specify from *0 - 9150* . The default is 16.

The **no-fragment** keyword turns on the "don't fragment" bit in the IPv6 header of the ping packet. This option is disabled by default.

The **quiet** keyword hides informational messages such as a summary of the ping parameters sent to the device and instead only displays messages indicating the success or failure of the ping. This option is disabled by default.

The **verify** keyword verifies that the data in the echo packet (the reply packet) is the same as the data in the echo request (the ping). By default the device does not verify the data.

The **data** *1 - 4 byte hex* parameter lets you specify a specific data pattern for the payload instead of the default data pattern, "abcd", in the packet's data payload. The pattern repeats itself throughout the ICMP message (payload) portion of the packet.

> NOTE
> For parameters that require a numeric value, the CLI does not check that the value you enter is within the allowed range. Instead, if you do exceed the range for a numeric value, the software rounds the value to the nearest valid value.

The **brief** keyword causes ping test characters to be displayed. The following ping test characters are supported:

! Indicates that a reply was received.

. Indicates that the network server timed out while waiting for a reply.

U Indicates that a destination unreachable error PDU was received.

I Indicates that the user interrupted ping.

Trace route

The trace route tool works by sending ICMP echo packets with varying IP Time-to-Live (TTL) values to the destination.

You can use IP traceroute to identify the path that packets take through the network on a hop-by-hop basis. The command output displays all network layer devices, such as routers, through which the traffic passes on its way to the destination.

The device can execute simultaneous **traceroute** commands from multiple inbound telnet or SSH sessions. Multiple simultaneous traceroutes from Web and SNMP, however are not allowed. The device can also resolve multiple DNS queries simultaneously, which allows multiple **traceroute** commands with the *hostname* option to be executed at the same time.

NOTE
Traceroute commands in outbound telnet sessions run on the remote telnet server and not on the local device.

Executing traceroute

The **traceroute** command, in the (Enable) mode, displays the routing path from the routing switch to the destination IP address as soon as the information is received. Traceroute requests display all responses to a given TTL. In addition, if there are multiple equal-cost routes to the destination, the device displays up to three responses by default.

NOTE
When executed in IPv4, the traceroute command does not display the IP address of the GRE tunnel interface path.

```
device> traceroute 10.33.4.7 minttl 5 maxttl 5 timeout 5
```

Syntax: traceroute *host-ip-addr* [**maxttl** *value*] [**minttl** *value*] [**numeric**] [**timeout** *value*] [**source-ip** *ip addr*]

The **maxttl** *value* parameter is the maximum TTL (hops) value: Possible value is 1 - 255. The default is 30 seconds.

The **minttl** *value* parameter is the minimum TTL (hops) value: Possible value is 1 - 255. The default is 1 second.

The **numeric** parameter lets you change the display to list devices by IP address instead of by name.

The **timeout** *value* parameter specifies the possible values. Possible value range is 1 - 120. Default value is 2 seconds.

The **source-ip** *ip addr* parameter specifies an IP address to be used as the origin for the traceroute.

Executing traceroute VRF

In the (Enable) mode, the **traceroute vrf** command functions like the standard **traceroute** command but requires you to specify a VRF table name. The **traceroute vrf** command must be used when the route to the destination is associated with a VRF table.

```
device# traceroute vrf blue 10.10.10.10
```

Syntax: traceroute vrf *vrf-nameip-address*

The *vrf-name* parameter is the name of the VRF for you want are running the traceroute.

The *ip-address* parameter is the IP address containing the VRF that you want to conduct a traceroute to.

Executing traceroute IPv6

The **traceroute ipv6** command traces a path from a device that supports IPv6 to an IPv6 host.

The CLI displays trace route information for each hop as soon as the information is received. Traceroute requests display all responses to a minimum TTL of 1 second and a maximum TTL of 30 seconds. In addition, if there are multiple equal-cost routes to the destination, the device displays up to three responses.

To trace the path from the device to a host with an IPv6 address of 2001:db8:349e:a384::34, enter the following command.

```
device> traceroute ipv6 2001:db8:349e:a384::34
```

Syntax: traceroute ipv6 *ipv6-address*

The *ipv6-address* parameter specifies the address of an IPv6 host. You must specify this address in hexadecimal using 16-bit values between colons.

Trace-l2 protocol

Trace-l2 traces is a proprietary protocol that traces the traffic path to a specified device in a VLAN. Also, it can be used to probe all reachable paths to all devices in a VLAN. It does the following:

- Traces a particular IP, MAC or hostname in a VLAN.
- Probes the entire Layer 2 topology.
- Displays the input or output ports of each hop in the path.
- Displays the round trip travel time of each hop.
- Displays hops in a VLAN that form a loop.
- Displays each hop's Layer 2 protocol such as STP, RSTP, 802.1w, SSTP, metro ring, or route-only.

The resulting trace displays a report that provides information about a packet's path to a device, such as hop and port information and travel time. It also can locate any Layer 2 loop in a VLAN. The probed Layer 2 information is discarded when a new **trace-l2** command is issued again.

For each hop in the path, trace-l2 displays its input/output port, Layer 2 (L2) protocols of the input port, and the microsecond travel time between hop and hop. It also prints out the hops which form a loop, if any. Displaying L2 topology lets a user easily obtain information of all hops.

Configuration considerations

The configuration considerations are as follows:

- Trace-l2 is enabled on the Extreme devices. It can be used to trace traffic only to devices.
- The devices that will participate in the trace-l2 protocol must be assigned to a VLAN and all devices on that VLAN must be Extreme devices that support the trace-l2 protocol.
- Extreme devices, as well as other vendor devices, that do not support the trace-l2 protocol, simply forward trace-l2 packets without a reply. Hence, these devices are transparent to the trace-l2 protocol.
- The destination for the packet with the trace-l2 protocol must be a device that supports the trace-l2 protocol and the destination cannot be a client, such as a personal computer, or devices from other vendors.

Tracing a traffic path

The trace-l2 protocol is enabled on a VLAN. You can trace the traffic path of a packet by entering a command such as the following.

```
device(config)#trace-l2 vlan 10 2.2.2.2
```

Syntax: [no] trace-l2 vlan *vlan-id destination-address*

The *destination address* can be a MAC address, an IP address, or a host. You can enter the destination-address in one of the following formats:

- HHHH.HHHH.HHHH - Destination MAC address
- A.B.C.D - Destination IP address
- ASCII string - destination host name

If a destination address is not specified or the destination does not exists, trace-l2 collects L2 topology information which can be displayed by issuing a **trace-l2 show** command. The command displays the following information.

```
trace-l2 reply vlan 2 from e26, 10.1.1.2, total round trip = 814 microsec
hop input  output   IP and/or MAC address       microsec   comment
```

```
  1   e28    e25    10.1.1.4 0000.003f.c400         316     e28: ring 11
  2   e15    e13    10.1.1.1 0000.0057.0d00         235     e15: ring 11
  3   e27           10.1.1.2 0000.0057.2500         263     e27: ring 11
```

In the output above, the last hop is the destination. Because 10.1.1.2 and 10.2.2.2 are addresses of the same device, the device can use 10.1.1.2 in the reply.

In general, **trace-l2** first tries to use the IP address of the virtual routing interface that is associated with a VLAN. If the virtual routing interface is not available, it then uses the loopback address. If both addresses are not available, it displays MAC address only.

The *input* and *output* ports show the path of the hops. Hop 3 has no output port because it is the destination.

The *microsec* column is the round trip time (sum of the time) to and from the previous hop. For example, 316 microsec for hop 1 is the time from the source to hop 1 and from hop 1 to the source. One way time is not available because the tracel2 protocol does not synchronize the clocks between hops.

The **comment** column shows the Layer 2 protocol used on the input port. It could be the following:

- STP - spanning tree protocol
- RSTP - Rapid STP, 802.1w draft 3
- 802.1w - Rapid STP
- ring - Metro ring ID of input port.
- Single STP - Includes Single STP, Single RSTP and Single 802.1w
- STP port disabled - The **spanning-tree ethernet disabled** command is configured.
- route-only - This device has route-only configuration
- port route-only - The input port has route-only configuration

Displaying Layer 2 topology information

To display information about the Layer 2 topology, first issue a **trace-l2vlan** command, then enter the **trace-l2 show** command as in the following example.

```
device(config)#trace-l2 vlan 10
Vlan 10 L2 topology probed, use "trace-l2 show" to display
device(config)#trace-l2 show
Vlan 10 L2 topology was probed 6 sec ago, # of paths: 2
path 1 from e27, 1 hops:
hop input  output IP and/or MAC address       microsec comment
1   e13           10.1.1.1 0000.0057.0d00           383 802-1w
path 2 from e25, 2 hops:
hop input  output IP and/or MAC address       microsec comment
1   e27    e26    10.1.1.3 0000.0052.ea00           657 802-1w
2   e28           10.1.1.4 0000.003f.c400           296 route-only
```

The **trace-l2 show** command does not display a path if the path is a subset of another path; therefore, the number of paths displayed could be fewer than the number of devices.

If the topology contains Layer 2 loops, a message such as the following is displayed.

```
*** Warning! The following 3 hops form a loop in vlan 2
  hop input  output IP and/or MAC address       microsec comment
  1   e25           10.1.2.2 0000.0057.2500
  2   e28           10.4.100.1 0000.003f.c400
  3   e29           10.1.1.1 0000.0057.0d00
```

Syntax: trace-l2 show

IPv6 Traceroute over an MPLS network

NOTE
IPv6 MPLS traceroute not supported on the BR-MLX-10Gx24-DM 24-port 10GbE module.

IPv6 traceroute behavior is similar to IPv4 traceroute. However, unlike IPv4 traceroute, IPv6 traceroute has a new 6PE label added during each hop across the MPLS cloud. Based on the IP header value, the node devices differentiate if the Internet Control Message Protocol version 6 (ICMPv6) echo request is from an IPv6 or IPv4 source device.

When the traceroute sends ICMPv6 echo request packets with a TTL value (hop limit) value of 1, the first router in the path replies with the *ttl-exceeded* error message to the source. The next packet has a TTL (hop limit) value of 2 and the second router replies with the *ttl-exceeded* error message. This process continues till the destination host receives the packets and returns an ICMPv6 Echo Reply message.

Based on the *ttl-exceeded* messages or the ICMPv6 Echo Reply messages received during the traceroute operation, the source device obtains details such as the hop sequence, total hops taken to complete the path, and the IPv4 or IPv6 addresses of devices that it passed during the path. For each hop, the traceroute gathers information about the hop number, best hop time, and the TTL value.

Tracing an IPv6 route through an MPLS domain

The following figure shows an MPLS-enabled provider network consisting of four LSRs. PE1 is the ingress PE Label Edge Router (LER), P1 and P2 are transit LSRs, and PE2 is the egress provider edge LER. CE1 and CE2 are CE devices located in different geographical locations.

FIGURE 8 IPv6 Traceroute in an MPLS cloud

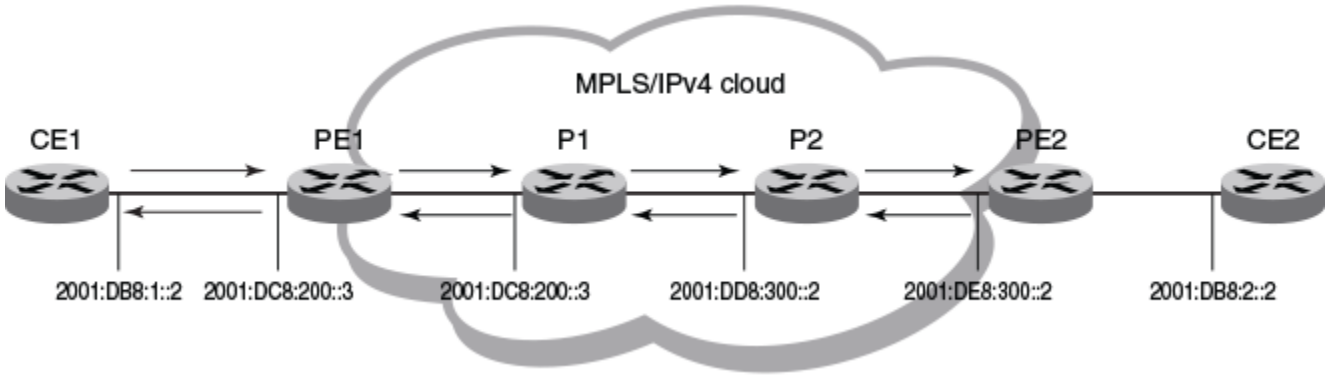

To understand the IPv6 traceroute behavior in an MPLS domain, assume the following:

- Customer traffic is tunneled through a MPLS VPN network, and traffic within the MPLS core is forwarded by label-switching only.
- The CE1 router sends UDP packets from CE1 router towards the CE2 router.
- Traceroute is configured to generate ICMPv6 messages per ICMP extensions and to use LSPs to forward these messages. Refer to Configuring IPv6 Traceroute over MPLS on page 96 for more information.
- The PE routers are aware of the source and destination IPv6 addresses while the transit LSRs have no such knowledge.
- The **traceroute** command is issued from CE1 to CE2 and reports the following information:

```
device# traceroute ipv6 2001:DB8:2::2
Type Control-c to abort
```

```
Tracing the route to IPv6 node 2001:DB8:2::2 from 1 to 30 hops
  1    <1 ms   <1 ms   <1 ms 2001:DB8:1::2
  2    <1 ms   <1 ms   <1 ms 2001:DC8:200::3
        MPLS Label=1026 Exp=0 TTL=1 S=0
        MPLS Label=794624 Exp=0 TTL=1 S=1
  3    <1 ms   <1 ms   <1 ms 2001:DD8:300::2
        MPLS Label=1029 Exp=0 TTL=1 S=0
        MPLS Label=794624 Exp=0 TTL=2 S=1
  4    <1 ms   <1 ms   <1 ms 2001:DE8:300::2
  5    <1 ms   <1 ms   <1 ms 2001:DB8:2::2
```

NOTE

The traceroute output reports information on a traceroute packet only when its TTL equals 1. Label stack information associated with subsequent routing of the ICMP message along the LSPs to the destination and back to the source is not displayed.

In the the previous scenario, the traceroute operation can be described as follows:

1. CE1 sends a traceroute probe with a TTL of 1 to its peer, CE2, with the destination IP address of 2001:DB8:2::2. PE1 decrements the packet's TTL by one and drops the expired packet. It generates a *ttl-exceeded* ICMPv6 message, and sends it back to CE1 with the source IPv6 address embedded in the IPv6 header of the expired packet. Traceroute reports the PE1 IPv6 address at hop 1, but there is no label information.

   ```
   1.     <1 ms   <1 ms   <1 ms 2001:DB8:1::2
   ```

2. CE1 sends a second traceroute probe to CE2, with an incremented TTL value of 2. PE1 decrements the TTL value to 1, and adds the 6PE label and the Label Distribution Protocol (LDP) label onto the packet to route it to CE2 by way of the transit router P1. PE1 also copies the TTL value from the IP header into the TTL field of the labels (recall that TTL propagation must be enabled on the ingress PE).

   The transit router P1 decrements the TTL, drops the expired packet since the TTL value is 0, and generates a *ttl-exceeded* ICMPv6 message. Before dropping the packet, and using the ICMPv6 extension mechanism, P1 copies the packet's label stack plus its IP header and appends both to the ICMPv6 message. Though the message destination is CE1, P1 cannot return the ICMPv6 message directly to CE1. It uses label-switching to forward the encapsulated ICMP response in the direction of the original traceroute probe along the configured LSPs and back to CE1. P1 sets the maximum TTL value of 255 to ensure that the message can reach its destination before it times out.

   Traceroute reports the IP address of P1, plus the label stack that was pushed onto the traceroute packet by PE1 and received by P1 when the packet's TTL was 1.

   ```
   2    <1 ms   <1 ms   <1 ms 2001:DC8:200::3
          MPLS Label=1026 Exp=0 TTL=1 S=0
          MPLS Label=794624 Exp=0 TTL=1 S=1
   ```

3. The third traceroute probe (TTL=3) is forwarded until it expires at the transit router P2. P2 (the Penultimate Hop Popping (PHP) LSR) generates the ICMPv6 message, appends the label stack from the expired traceroute packet, and passes it on to PE2 without imposing a label. PE2 forwards the ICMPv6 message back to CE1 along the return LSP.

   Traceroute reports the IP address of P2, plus the label stack which P2 received with the traceroute packet from P1 when the packet's TTL was 1.

   ```
   3    <1 ms   <1 ms   <1 ms 2001:DD8:300::2
          MPLS Label=1029 Exp=0 TTL=1 S=0
          MPLS Label=794624 Exp=0 TTL=2 S=1
   ```

4. The fourth traceroute probe (TTL=4) is forwarded until it expires at the egress provider edge device PE2. PE2 drops the packet and generates a *ttl-exceeded* ICMPv6 message without label stack extension since there is no label stack to report.

   Traceroute reports only the IP address of PE2. The transit router P2 popped the outer label before passing the traceroute packet on to the egress PE2 and PE2 pops the VPN label before sending the ICMPv6 message back to the customer source device CE1.

   ```
   4    <1 ms   <1 ms   <1 ms 2001:DE8:300::2
   5    <1 ms   <1 ms   <1 ms 2001:DB8:2::2
   ```

5. The fifth traceroute probe (TTL=5) has a TTL large enough for the packets to reach the customer destination device CE2. CE2 generates an ICMPv6 *port unreachable* message, which CE2 sends back to CE1.

   Traceroute reports only the IP address of the destination device CE2. No label extension is added because the received packet is not labeled. The *port unreachable* message is label-switched back to the customer source device CE1, as a normal data packet.

   ```
   5    <1 ms   <1 ms   <1 ms 2001:DB8:2::2
   ```

Configuring IPv6 Traceroute over MPLS

The **ipv6icmp mpls-response** command configures the behavior of the traceroute operation by controlling both the ICMPv6 message format (use ICMPv6 label stack extensions or not) and the manner in which the ICMPv6 messages are **forwarded through an MPLS domain** (by way of IP routing table lookup or through label-switching using LSPs).

MPLS response is enabled by default. To enable the MPLS response after it was disabled, enter the following command:

```
device(config)# ipv6 icmp mpls-response
```

You can use this version of the command if the traceroute is over an IPv6-aware MPLS core. In such a case, IPv6 traceroute uses the default option of using the routing tables to forward packets. The IPv6 link local addresses should not be used to send the ICMPv6 packet. At the same time, you can still use the **ipv6 icmp mpls-responseuse-lsp** command to use the configured LSPs.

To specify using LSP to forward the ICMPv6 messages with MPLS label extensions, enter the following command:

```
device(config)# ipv6 icmp mpls-response use-lsp
```

Use this version of the command if the MPLS core is non IPv6-aware, because the IPv6 forwarding will not work.

To specify generating ICMPv6 messages without MPLS label extensions, enter the following command:

```
device(config)# ipv6 icmp mpls-response no-label-extensions
```

To disable the IPv6 Traceroute over MPLS feature, enter the following command:

```
device(config)# no ipv6 icmp mpls-response
```

Syntax: [no] ipv6 icmp mpls-response [**use-lsp**] [**no-label-extension**]

The **mpls-response**parameter enables the ICMPv6 traceroute response in default mode. The feature is enabled by default and configured to use IP routing to forward ICMP messages.

The **use-lsp** parameter enables forwarding of ICMPv6 error messages along the LSPs configured for the MPLS domain. By default, using configured LSPs use is disabled.

The **no-label-extension** parameter disables the use of label stack information in the ICMPv6 error messages.

The **no** option disables the ICMPv6 traceroute response configuration. When the ICMP traceroute feature is disabled, standard traceroute using IPv6 forwarding is used to trace a traffic path through an MPLS domain.

NOTE
The **ipv6 icmp mpls-response** command supports TTL expiry for IPv6 packets only.

The output of the **show ipv6 traffic** command displays counts for ICMPv6 *ttl-exceeded* error reply packets.

LSP ping and traceroute

Overview

The LSP Ping and Traceroute feature provides Operation, Administration, and Maintenance (OAM) functionality for MPLS networks based up RFC 4379 (Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures).

The LSP ping and traceroute functions provide a mechanism to detect MPLS data plane failure. LSP ping is used to detect data plane failure and to check the consistency between the data plane and the control plane. LSP traceroute is used to isolate the data plane failure to a particular router and to provide LSP path tracing. They are implemented using MPLS echo request and reply messages which are sent as UDP packets to a well-known UDP port 3503. This section provides the details of LSP Ping and Traceroute operation

LSP ping operation

An MPLS echo request (described in MPLS echo request on page 97) is sent from the ingress to the egress LSR. At the transit LSRs, the ping packet is label switched (the same as a regular MPLS data packet) without any control plane intervention. Upon arriving at the egress LSR, the echo request is sent to the control plane for processing based on the IP Router Alert option and the well-known destination UDP port 3503. An echo reply (described in MPLS echo reply on page 98) is sent back as a UDP packet with an appropriate return code that depends on the result of the FEC stack validation.

LSP traceroute operation

An MPLS echo request (described in MPLS echo request on page 97) is sent from the ingress LSR with the TTL of the outermost label set to an incremental value that starts with a TTL value of 1. This request causes the MPLS echo request to be forwarded to the control plane for processing at each transit LSR, based on the MPLS TTL expiration value. An echo reply (described in MPLS echo reply on page 98) is sent back with a return code indicating that it is the transit LSR for the FEC specified in the echo request. This process repeats until the echo request arrives at the egress LSP. The echo request is then forwarded to the control plane for processing, based on the IP Router Alert option. An echo reply is sent back as a UDP packet with an appropriate return code that depends on the result of the FEC stack validation.

MPLS echo request

The MPLS echo request is sent from the ingress LSR as a labeled UDP packet (except for single-hop LSP). The echo request has the following characteristics.

IP/UDP header information:

- Source address = user-input or LSR ID.
- Destination address = user-input or 127.0.0.1.
- UDP source port = 3503.

- UDP destination port = 3503.
- IP TTL = 1
- Router Alert option is set.

By default, the reply mode is set to 2 (reply by way of an IPv4 UDP packet), and you can set it to 1 (no reply) or 3 (reply by way of an IPv4 UDP packet with Router Alert option).

The sender handle is set to an internally-generated, 32-bit number that is assigned to each ping or traceroute session when the ping or traceroute operation begins. This sender handle is sent back in the echo reply, which is used to locate the appropriate ping or traceroute session.

The sequence number is a running number associated with each ping or traceroute session. It is starts with a value of 1.

The TTL for the outermost label is set to 255 for a ping. For traceroute, it is 1, 2, 3, and so on.

You can configure a timeout when starting the ping or traceroute command. The default value is 5 seconds.

MPLS echo reply

The MPLS echo reply is sent by the transit (for traceroute) or egress (for ping and traceroute) LSR as a regular IPv4 UDP packet or an IPv4 UDP packet with Router Alert option depending on the reply-mode field of the echo request. If reply with Router Alert option is chosen, the user should make sure that all intermediate routers are capable of handling MPLS echo reply. If a reply is sent with Router Alert option and the reply is sent over a tunnel interface, the MPLS Router Alert label (label value 1) will be the topmost label for the packet. A reply with a Router Alert option should be used if and only if the normal IP return path is deemed unreliable.

The echo reply has the following characteristics.

IP/UDP header information:

- Source address = LSR ID
- Destination address = source IP address from the echo request
- UDP source port = 3503
- UDP destination port = UDP source port from the echo request
- IP TTL = 255
- Router Alert option set if and only if reply-mode field of the echo request set to 3.

The sender handle is copied from echo request message

The sequence number is copied from echo request message

LSP ping TLVs

Table 14 lists the TLVs defined in RFC 3479 that are included in an echo request and reply.

TABLE 14 Show Cfm output descriptions

TLV type	TLV name	TX in echo request	TX in echo reply
1	Target FEC stack	Yes	No
2	Downstream mapping	Yes if the **dsmap** option is set	Yes for transit LSRs only if downstream mapping TLV is included in the MPLS Echo request.
3	Pad	Depend on the size option	Yes (if value = 2)
7	Interface and Label Stack	N/A	Yes if the I flag in DS mapping is set

TABLE 14 Show Cfm output descriptions (continued)

TLV type	TLV name	TX in echo request	TX in echo reply
9	Errored TLV	N/A	Yes (if error is detected)
10	Reply TOS bytes	Yes if reply-tos option is set	TLV is not sent back. Just copy TOS byte into IP header.

The Extreme devices support sending and receiving downstream mapping TLVs without multipath information (where the multipath type is always set to 0). Note that the detailed multipath information can be used by the ingress LSR to ping or traceroute through all ECMP paths at the transit LSR. Currently, the Extreme devices do not support LDP LSPs with ECMP. Consequently, the multipath type of non-zero is not relevant in these operations.

LSP FEC types

For LDP LSPs, the LDP IPv4 prefix sub-TLV (sub-type = 1) is encoded in the target FEC stack of the echo request. For RSVP LSPs, the RSVP IPv4 LSP sub-TLV (sub-type = 3) is encoded in the target FEC stack.

NOTE
Static RSVP LSPs are no longer supported, so a ping or traceroute for a static LSP is not supported.

Redundant RSVP LSPs

For RSVP LSPs with redundant paths, ping or traceroute on a LSP is performed on the currently active path. For example, if the secondary path is the active path for an LSP, the MPLS echo request packets are sent out on the secondary path's interface.

If the active path changes while a ping or traceroute is in progress, the echo request continues to be sent out on the old active path. This implies that the echo request that was sent after path switchover times out. The user subsequently needs to restart the ping or traceroute.

One-to-one Fast ReRoute (FRR) LSPs

Similar to the redundant LSPs, a ping or traceroute on a one-to-one FRR LSP is performed on the active path. If a path switchover occurs while a ping or traceroute is in-progress, the echo request continues to be sent out on the old active path. This implies that the echo request sent after path switchover will time out.

A user can ping or trace the route of the ingress-originated detour of a one-to-one FRR LSP by specifying the detour parameter. The operation is started only if the detour is operationally up.

FRR bypass LSPs

The LSP ping and traceroute facilities support FRR bypass LSPs. You can ping or trace the protected LSP and bypass tunnel separately.

You can ping or trace the ingress-originated or transit-originated bypass tunnel by specifying either the name of bypass LSP (as you would any regular LSP name) or the entire RSVP session ID (including the tunnel endpoint, the tunnel ID, and the extended tunnel ID).

NOTE
In the current facility backup implementation, the bypass LSP name must be unique in the system (for example, the name cannot be the same as the regular LSP name).

The traceroute output of a backup tunnel depends on the setting of the **propagate-ttl** and **label-propagate-ttl-** options. If both **propagate-ttl** and **label-propagate-ttl** options are turned on, the traceroute output shows the detail of the bypass path. If both options are turned off, the bypass path is shown as a single hop. The options should be either both ON or both OFF.

To trace the route of a backup path, the TTL of the bypass and protected labels (if they are not implicit NULL labels) are set as in the following example:

- Both **propagate-ttl**and **label-propagate-ttl** are ON: TTL = 1, 2, 3, and so on, are set for both labels.
- Otherwise: bypass label TTL is set to 255. Protected label TTL is set to 1, 2, 3, and so on.

IP TTL is set to topmost label TTL. Otherwise, it is set to 255.

Transit-originated detour

The user can initiate a ping or traceroute operation on a transit-originated, detour LSP. Because the session name does not uniquely identify a session on a transit LSR, the user needs to specify the entire session ID (including the tunnel endpoint, tunnel ID, and extended tunnel ID) for the detour LSP to which the LSP ping or traceroute command is applied.

LSP reoptimization

If LSP reoptimization happens while the ping or traceroute is operating, the echo request is still sent out on the current LSP instance until the new instance is created. This avoids displaying partial information from the old and new paths if they are different; particularly for a traceroute. Similarly, if the ping or traceroute operation is started while LSP reoptimization is occurring, the LSP label, out interface, and other parameters from the currently up instance will be used.

PHP behavior

Ping is transparent to the penultimate LSR. MPLS and IP TTL operations performed on a ping packet are the same as tor a regular data packet. In the default case where the MPLS TTL is copied into the IP TTL, the echo request packet can arrive at the egress LSR with an IP TTL value greater than 1. Consequently, in this situation, the IP Router Alert option is used to direct the echo request packet to the control plane for ping processing.

For a traceroute operation, if the echo request is received with a downstream mapping TLV, the Implicit Null label is encoded in the Downstream label in the echo reply just like any other label.

Since a Extreme device advertises an implicit Null label to its upstream LSR for both LDP and RSVP LSPs, packets that arrive at the egress LSR do not have the tunnel label. For a single-hop LSP, the echo request is sent out from ingress LSR as an unlabeled UDP packet.

Using the LSP ping and traceroute commands

The following sections described operation of the LSP Ping and Traceroute command:

- Executing LDP LSP ping on page 100
- Executing RSVP LSP ping on page 101
- Executing LDP LSP traceroute on page 102
- Executing RSVP LSP traceroute on page 103

Executing LDP LSP ping

The LDP LSP ping command, sends an MPLS echo request from the ingress to the egress LSR.

To perform the LDP LSP ping operation, use the following command.

```
device# ping mpls ldp 10.22.22.22
Send 5 80-byte MPLS Echo Requests for LDP FEC 10.22.22.22/32, timeout 5000 msec
Type Control-c to abort
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max=0/1/1 ms.
device)#
```

Syntax: ping mpls ldp *ip-address* | *ip-address/mask-length* [**count** *num*] [**destination** *ip-address*] [**detail**] [**reply-mode no-reply** | **reply-mode router-alert**] [**reply-tos** *num*] [**size** *bytes*] [**source** *ip-address*] [**timeout***msec*] [**nexthop** *ipv4address*]

The **ldp** *ip-address* and *ip-address/mask-length* variables specify the LDP IPv4 destination prefix and mask length. If the **mask-length** is not specified, the default value is 32.

The **count** option with the *num* variable specifies the number of echo requests to send. The default value is 5.

The **destination** option specifies an IP address within the 127/8 subnet. The default address is 127.0.0.1

The **detail** option displays the details of the echo request and reply messages. By default, the display is in the brief mode.

The **reply-mode** option species the reply mode field in the echo request if and only if the user does not want the reply to be sent as an IPv4 UDP packet.

The **no-reply** option can be used to test one-way connectivity.

The **router-alert** option is used when the normal IP return path is unreliable. This option indicates that the reply should be sent as an IPv4 UDP packet with the Router Alert option. This option requires extra overhead processing at each LSR along the return path.

The **reply-tos** option specifies a TOS value between 0 and 254 to be included in the Reply-TOS-byte TLV. This value will be copied to the IP header TOS byte of the echo reply. By default, the reply-tos TLV is not included in the echo request.

> NOTE
> The last bit of the TOS byte is always 0.

The **size** option specifies that the size of the echo request including the label stack to be sent will be the value of the variable *bytes*. The pad TLV is used to fill the echo request message to the specified size. The minimum size is 80 byte for an LDP echo request. The maximum size is the size of the LSP MTU.

The **source** option specifies the IP address of any interface. This address is used as the destination address for the echo reply address. The default address is the LSR ID.

The **timeout** option specifies an interval in milliseconds for the echo request message. The default timeout is 5 seconds. The maximum timeout value is 5 minutes.

The **nexthop** specifies the nexthop IPv4 address to send the OAM request to. If an address that does not match the outgoing path for the tunnel is given, following error message appears as response:

```
Ping fails: LDP next-hop does not exist.
```

Executing RSVP LSP ping

The RSVP ping command in the (enable) mode, sends an MPLS echo request from the ingress to the egress LSR.

To perform the RSVP LSP ping operation, use the following command.

```
device# ping mpls rsvp lsp toxmr2frr-18
Send 5 92-byte MPLS Echo Requests over RSVP LSP toxmr2frr-18, timeout 5000 msec
Type Control-c to abort
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max=0/1/5 ms.
device)#
```

Syntax:ping mpls rsvp lsp *lsp-name*| **session** *tunnel-source-address tunnel-destination-addresstunnel-id* [**count** *num*] [**destination** *ip-address*] [**detail**] [**detour**] [**reply-mode no-reply** | **reply-mode router-alert**] [**reply-tos** *num*] [**size** *bytes*] [**source** *ip-address*] [**standby**] [**timeout** *msec*]

The **rsvp lsp** option specifies the name of the RSVP IPv4 LSP in the **lsp-name** variable.

The **rsvp session** option specifies the session ID. The **tunnel-source-address**,**tunnel- destination-address** and *tunnel-id* variables must all be specified to form a valid session ID.

The **count** option with the *num* variable specifies the number of echo requests to send. The default value is 5.

The **destination** option specifies an IP address within the 127/8 subnet. The default address is 127.0.0.1

The **detail** option displays the details of the echo request and reply messages. By default, the display is in the brief mode.

The **detour** option specifies a ping detour path. For a detour originated on the ingress LSR, you can ping the detour path using either the LSP name or session ID with the **detour** option specified.

The **reply-mode** option species the reply mode field in the echo request if and only if the user does not want the reply to be sent as an IPv4 UDP packet.

The **no-reply** option can be used to test one-way connectivity.

The **router-alert** option is used when the normal IP return path is unreliable. This option indicates that the reply should be sent as an IPv4 UDP packet with the Router Alert option. This option requires extra overhead processing at each LSR along the return path.

The **reply-tos** option specifies a TOS value between 0 and 254 to be included in the Reply-TOS-byte TLV. This value will be copied to the IP header TOS byte of the echo reply. By default, the reply-tos TLV is not included in the echo request.

> NOTE
> The last bit of the TOS byte is always 0.

The **size** option specifies that the size of the echo request including the label stack to be sent will be the value of the variable *bytes*. The pad TLV is used to fill the echo request message to the specified size. The minimum size is 92 bytes for an MPLS Echo request. The maximum size is the size of the LSP MTU.

The **source** option specifies the IP address of any interface. This address is used as the destination address for the echo reply address. The default address is the LSR ID.

The **standby** option directs the ping operation to the secondary path of a redundant LSP that is operationally up.

The **timeout** option specifies an interval in milliseconds for the echo request message. The default timeout is 5 seconds. The maximum timeout value is 5 minutes.

Executing LDP LSP traceroute

The LDP LSP traceroute command in the (enable) mode, sends and MPLS echo request from the ingress to the egress LSR.

To perform the LDP LSP traceroute operation, use the following command.

```
device# traceroute mpls ldp 10.22.22.22
Trace LDP LSP to 10.22.22.22/32, timeout 5000 msec, TTL 1 to 30
Type Control-c to abort
 1 10ms 10.22.22.22 return code 3(Egress)
device)#
```

Syntax: traceroute mpls ldp *ip-address/mask-length*[**destination** *ip-address*] [**dsmap**] [**min-ttl** *min-num*] [**max-ttl** *max-num*] [**reply-mode router-alert**] [**reply-tos** *num*] [**size** *bytes*] [**source** *ip-address*] [**timeout** *msec*] [**nexthop** *ipv4address*]

The **ldp** *ip-address/mask-length* variable specifies the LDP IPv4 destination prefix and mask length. If the **mask-length** is not specified, the default value is 32.

The **destination** option specifies an IP address within the 127/8 subnet. The default address is 127.0.0.1

The **dsmap** option enables the Downstream (DS) mapping TLV in the echo request for traceroute operation. The DS mapping TLV is used to instruct the transit LSR to include the next-hop interface and label information in the echo reply. By default, the DS TLV is not included in the echo request.

The **min-ttl** option specifies a minimum value in the *min-num* variable for the outermost label in traceroute operation. The default minimum TTL value is 1. Acceptable values that can be configured are: 1 - 255.

The **max-ttl** option specifies a maximum value in the *max-num* variable for the outermost label in traceroute operation. The default maximum TTL value is 30. Acceptable values that can be configured are: 1 - 255.

The **reply-moderouter-alert** option is used when the normal IP return path is unreliable. This option indicates that the reply should be sent as an IPv4 UDP packet with the Router Alert option. This option requires extra overhead processing at each LSR along the return path.

The **reply-tos** option specifies a TOS value between 0 and 254 to be included in the Reply-TOS-byte TLV. This value will be copied to the IP header TOS byte of the echo reply. By default, the reply-tos TLV is not included in the echo request.

NOTE
The last bit of the TOS byte is always 0.

The **size** option specifies that the size of the echo request including the label stack to be sent will be the value of the variable *bytes*. The pad TLV is used to fill the echo request message to the specified size. The minimum size is 92 bytes for an MPLS Echo request. The maximum size is the size of the LSP MTU.

The **source** option specifies the IP address of any interface. This address is used as the destination address for the echo reply address. The default address is the LSR ID.

The **timeout** option specifies an interval in milliseconds for the echo request message. The default timeout is 5 seconds. The maximum timeout value is 5 minutes.

The **nexthop** specifies the nexthop IPv4 address to send the OAM request to. If an address that does not match the outgoing path for the tunnel is given, following error message appears as response:

```
Traceroute fails: LDP next-hop does not exist.
```

Executing RSVP LSP traceroute

The RSVP LSP traceroute command in the (enable) mode, sends and MPLS echo request from the ingress to the egress LSR.

To perform the RSVP LSP traceroute operation, use the following command.

```
device# traceroute mpls rsvp lsp toxmr2frr-18
Trace RSVP LSP toxmr2frr-18, timeout 5000 msec, TTL 1 to 30
Type Control-c to abort
 1 1ms 10.22.22.22 return code 3(Egress)
device#
```

Syntax: traceroute mpls rsvp lsp *lsp-name* | **session** *tunnel-source-addresstunnel-destination-addresstunnel-id* [**destination***ip-address*] [**dsmap**] [**detour**] [**min-ttl** *min-num*] [**max-ttl** *max-num*] [**reply-mode router-alert**] [**reply-tos** *num*] [**size** *bytes*] [**source***ip-address*] [**standby**] [**timeout** *msec*]

The **rsvp lsp** option specifies the name of the RSVP IPv4 LSP in the **lsp-name** variable.

The **rsvp session** option specifies the session ID. The*tunnel-source-address, tunnel- destination-address* and *tunnel-id* variables must all be specified to form a valid session ID.

The **destination** option specifies an IP address within the 127/8 subnet. The default address is 127.0.0.1

The **dsmap** option enables the Downstream (DS) mapping TLV in the echo request for traceroute operation. The DS mapping TLV is used to instruct the transit LSR to include the next-hop interface and label information in the echo reply. By default, the DS TLV is not included in the echo request.

The **detour** option specifies a traceroute detour path. For a detour originated on the ingress LSR, you can ping the detour path using either the LSP name or session ID with the **detour** option specified.

The **reply-moderouter-alert** option is used when the normal IP return path is unreliable. This option indicates that the reply should be sent as an IPv4 UDP packet with the Router Alert option. This option requires extra overhead processing at each LSR along the return path.

The **reply-tos** option specifies a TOS value between 0 and 254 to be included in the Reply-TOS-byte TLV. This value will be copied to the IP header TOS byte of the echo reply. By default, the reply-tos TLV is not included in the echo request.

NOTE
The last bit of the TOS byte is always 0.

The **size** option specifies that the size of the echo request including the label stack to be sent will be the value of the variable *bytes*. The pad TLV is used to fill the echo request message to the specified size. The minimum size is 92 bytes for an MPLS Echo request. The maximum size is the size of the LSP MTU.

The **source** option specifies the IP address of any interface. This address is used as the destination address for the echo reply address. The default address is the LSR ID.

The **standby** option directs the traceroute operation to the secondary path of a redundant LSP that is operationally up.

The **timeout** option specifies an interval in milliseconds for the echo request message. The default timeout is 5 seconds. The maximum timeout value is 5 minutes.

Displaying LSP ping and traceroute statistics

You can use the **show mpls statistics oam** command to display the following LSP ping and traceroute counters:

- Ping and traceroute requests that are issued by the user
- Echo requests sent
- Echo requests received
- Echo request time-outs
- Echo replies sent
- Echo replies received
- Echo replies with error return codes

To display the LSP ping and traceroute counters use the **show mpls statistics oam** command, as shown in the following.

```
device # show mpls statistics oam
User ping request processed: 8
User traceroute request processed: 3
Echo requests: sent(102658), received(2865), timeout(0)
Echo replies: sent(2865), received(102628)
Echo reply return code distribution:        TX         RX
  Egress(3)                         :       0          102628
  Transit(8)                        :       0          0
  No return code(0)                 :       0          0
  Malformed request(1)              :       0          0
  Unsupported TLV(2)                :       2865       0
  No FEC mapping(4)                 :       0          0
  DS map mismatch(5)                :       0          0
  Unknown upstream intf(6)          :       0          0
  Reserved return code(7)           :       0          0
```

```
Unlabeled output intf(9)          :        0          0
FEC mapping mismatch(10)          :        0          0
No label entry(11)                :        0          0
Rx intf protocol mismatch(12)     :        0          0
```

Premature LSP termination(13) : 0 0

Syntax: show mpls statistics oam

When the detail option is specified, the echo reply is shown with a error return code based on the error codes listed in RFC 4379.

Clearing the LSP ping and traceroute counters

You can use the **clear mpls statistics oam** command to clear the LSP ping and traceroute counters as shown in the following.

```
device# clear mpls statistics oam
```

Syntax: clear mpls statistics oam

CFM monitoring for ISID

- ISID is configured in edge devices (BEB) of a PBB network.
- CFM is configured for ISID in a BEB and is monitored between BEBs.
- The CCM interval for the sub-second timer is supported for CER with PBIF version 0x56 and greater.
- Loopback, Link trace, and delay measurement messages are supported for ISID.
- MIP functionality is not applicable for ISID.

Configuring CFM monitoring for ISID

The following PBB configuration is mandatory to configure CFM ISID.

1. Configure ESI for B-VLAN and VLAN under the ESI.
2. Add ports into the configured B-VLAN.
3. Configure ESI for ISID and ISID under the ESI.
4. Associate ISID ESI as client ESI with B-VLAN ESI.

   Use the following commands for each step in the CFM configuration for ISID.

Sample configuration

```
device(config)#interface eth 1/1
device(config-if-e1000-1/1)#enable
device(config-if-e1000-1/1)#port-type backbone-network
device(config)#esi isid_1 encapsulation isid
device(config-esi-isid_esi_1)#isid 2000
device(config)# esi bvlan_1 encapsulation bvlan
device(config-esi-bvlan_1)#vlan 200
device(config-esi-bvlan_1-vlan-200)#tagged eth1/1 device(config-esi-bvlan_1)#esi-client isid_1

device(config)#interface ethernet 1/2
device(config-if-e1000-1/2)#enable
device(config-if-e1000-1/2)#port-type backbone-network
device(config)#esi isid_1 encapsulation isid
device(config-esi-isid_esi_1)#isid 2000
```

```
device(config)# esi bvlan_1 encapsulation bvlan
device(config-esi-bvlan_1)#vlan 200
device(config-esi-bvlan_1-vlan-200)#tagged ethernet 1/2
device(config-esi-bvlan_1)#esi-client isid_esi_1

device(config)#tag-value tag1 88A8
device(config)#interface ethernet 1/1
device(config-if-e1000-1/1)#enable
device(config)#interface ethernet 1/2
device(config-if-e1000-1/2)#enable
device(config)#vlan 200
device(config-vlan-200)#tagged eth 1/1
device(config-vlan-200)#tagged eth 1/2
```

Sample configuration for ISID CFM

The following configuration shows the sample configuration for ISID CFM.

```
device(config)#cfm-enable
device(config-cfm)#domain-name ISID_domain level 7
device(config-cfm-md-ISID_domain)#ma-name ISID_2000 esi isid_1 isid 2000 priority 7
device(config-cfm-md-ISID_domain-ma-ISID_2000)#ccm-interval 1-second
device(config-cfm-md-ISID_domain-ma-ISID_2000)#mep 1 down port eth 1/1
device(config-cfm-md-ISID_domain-ma-ISID_2000)#

device(config-cfm)#domain-name ISID_domain level 7 device(config-cfm-md-ISID_domain)#ma-name ISID_2000 esi
isid_1 isid 2000 priority 7
device(config-cfm-md-ISID_domain-ma-ISID_2000)#ccm-interval 1-second
device(config-cfm-md-ISID_domain-ma-ISID_2000)#mep 2 down port eth 1/2 device(config-cfm-md-ISID_domain-ma-
ISID_2000)#
```

Show commands for CFM monitoring for ISID

The following **show** commands provide output for each component of the sample configuration.

Show cfm

Use the **show cfm** command to display the cfm configuration.

Syntax: show cfm

```
device#show cfm
Domain: ISID_domain
Index: 1
Level: 3
Maintenance association: ISID_2000
Ma Index: 1
CCM interval: 1000 ms
ESI isid_1 ISID : 2000
Priority: 7
ETH-AIS TX: DISABLED
ETH-AIS RX: DISABLED
ETH-AIS Interval: 10 sec
MEP  Direction MAC            PORT       PORT-STATUS-TLV
==== ========= ============== ========== ===============
1    DOWN      0000.0011.86d1 ethe 1/1   N

device#show cfm
Domain: ISID_domain
Index: 1
Level: 3
Maintenance association: ISID_2000
Ma Index: 1
CCM interval: 1000 ms
ESI isid_1 ISID : 2000
Priority: 7
ETH-AIS TX: DISABLED
```

```
ETH-AIS RX: DISABLED
ETH-AIS Interval: 10 sec
MEP  Direction MAC             PORT       PORT-STATUS-TLV
==== ========= ============== ========== ===============
2    DOWN      0000.00ef.2a0b ethe 1/2   N
```

Show cfm connectivity

Use the **show cfm connectivity** command to display the cfm connectivity configuration.

Syntax: show cfm connectivity

```
device#show cfm connectivity
Domain: ISID_domain Index: 1
Level: 3
Maintenance association: ISID_2000
MA Index: 1
CCM interval: 1000 ms
ESI: isid_1 ISID: 2000
Priority: 7
ETH-AIS TX: DISABLED
ETH-AIS RX: DISABLED
ETH-AIS Interval: 10 sec
RMEP MAC            ISID  AGE  PORT    SLOTS STATE AIS_STATE
==== ============== ===== ==== ======= ===== ===== ======
2    0000.00ef.2a0b 2000  231  1/1     1     OK    None
device#show cfm connectivity
Domain: ISID_domain Index: 1
Level: 3
Maintenance association: ISID_2000
MA Index: 1
CCM interval: 1000 ms
ESI: isid_1 ISID: 2000
Priority: 7
ETH-AIS TX: DISABLED
ETH-AIS RX: DISABLED
ETH-AIS Interval: 10 sec
RMEP MAC            ISID  AGE  PORT    SLOTS STATE AIS_STATE
==== ============== ===== ==== ======= ===== ===== ======
1    0000.0011.86d1 2000  317  1/2     1     OK    None
```

Loopback messages

CFM loopback

Use the **cfm loopback** command to display loopback messages.

Syntax: cfm loopback domain *domain-name* **ma** *ma-name* **src-mep** *ID* **target-mep** *ID*

The following output shows the Loopback messages.

```
device#cfm loopback domain ISID_domain ma ISID_2000 src-mep 1 target-mep 2
DOT1AG: Sending 10 Loopback to 0000.00ef.2a0b, timeout 10000 msec
Type Control-c to abort
Reply from 0000.00ef.2a0b: time=1ms
Reply from 0000.00ef.2a0b: time<1ms
Reply from 0000.00ef.2a0b: time<1ms
Reply from 0000.00ef.2a0b: time<1ms
Reply from 0000.00ef.2a0b: time<1ms
Reply from 0000.00ef.2a0b: time<1ms
Reply from 0000.00ef.2a0b: time<1ms
Reply from 0000.00ef.2a0b: time<1ms
Reply from 0000.00ef.2a0b: time<1ms
Reply from 0000.00ef.2a0b: time<1ms
sent = 10 number = 10 A total of 10 loopback replies received.
Success rate is 100 percent (10/10), round-trip min/avg/max=0/0/1 ms.
```

CFM linktrace

Use the **cfm linktrace** command to display linktrace messages.

Syntax: cfm linktrace domain *domain-name* **ma** *ma-name* **src-mep** *ID* **target-mep** *ID*

The following output shows the linktrace messages.

```
device#cfm linktrace domain ISID_domain ma ISID_2000 src-mep 1 target-mep 2
Linktrace to 0000.00ef.2a0b on Domain ISID_domain, level 3: timeout 10ms, 8 hops
------------------------------------------------------------------
Hops MAC Ingress Ingress Action Relay Action
Forwarded Egress Egress Action Nexthop
------------------------------------------------------------------
Type Control-c to abort
1 0000.00ef.2a0b 1/2 IgrOK RLY_HIT
Not Forwarded
Destination 0000.00ef.2a0b reached
```

Delay-Measurement

CFM delay_measurement

Use the **cfm delay_measurement** command to display the delay measurement and delay variation using ISID.

Syntax: cfm delay_measurement domain *domain-name* **ma** *ma-name* **src-mep** *ID* **target-mep** *ID*

The following output shows the delay measurement and delay variation using ISID.

```
device#cfm delay_measurement domain ISID_domain ma ISID_2000 src-mep 1 target-mep 2
Y1731: Sending 10 delay_measurement to 0000.00ef.2a0b, timeout 1000 msec tras=0
Type Control-c to abort
Reply from 0000.00ef.2a0b: time= 35.295 us
Reply from 0000.00ef.2a0b: time= 35.400 us
Reply from 0000.00ef.2a0b: time= 35.115 us
Reply from 0000.00ef.2a0b: time= 35.265 us
Reply from 0000.00ef.2a0b: time= 35.040 us
Reply from 0000.00ef.2a0b: time= 35.265 us
Reply from 0000.00ef.2a0b: time= 35.190 us
Reply from 0000.00ef.2a0b: time= 35.325 us
Reply from 0000.00ef.2a0b: time= 35.280 us
Reply from 0000.00ef.2a0b: time= 35.205 us
sent = 10 number = 10 A total of 10 delay measurement replies received.
Success rate is 100 percent (10/10)
================================================================================
Round Trip Frame Delay Time : min = 35.040 us avg = 35.238 us max = 35.400 us
Round Trip Frame Delay Variation : min = 45 ns avg = 146 ns max = 285 ns
================================================================================
```

Link MA

Link MA can be used to monitor connectivity between any two Links in the network. It can be configured between any links since it is independent of the VLAN.

- The CCM interval for a sub-second timer is supported for CER with PBIF Support.
- Loopback and delay measurement messages are supported for Link MA.

Configuring Link MA

The below step captures the CFM configuration for Link MA

1. Domain configuration.

```
device(config-cfm)#domain-name d7 level 7
```

Syntax: domain-name *name* **level** *value*

2. MA configuration.

```
device(config-cfm-md-d7)#ma-name link link-ma priority 7
```

Syntax: ma-name *name* **link-ma priority** *value*

3. MEP configuration.

```
device(config-cfm-md-d7-ma-link)#mep 1 down port eth 1/1
```

Syntax: mep *ID dir* **port** *portID*

4. Individual -link monitor configuration.

```
device(config-cfm-md-erp-ma-ma-erp)#individual-link-monitoring
```

Syntax: [no] individual-link-monitor

Sample Link MA configuration

The following sample configuration shows the Link Monitoring between DUT1 and DUT2. It also shows the Link Monitoring between DUT2 and DUT3.

DUT1

```
device(config)#cfm-enable
device(config-cfm)#domain-name d7 level 7
device(config-cfm-md-d7)#ma-name link link-ma priority 7
device(config-cfm-md-d7-ma-link)#ccm-interval 1-second
device(config-cfm-md-d7-ma-link)#mep 1 down port eth 1/1
```

DUT2

```
device(config)#cfm-enable DUT_2(config-cfm)#domain-name d7 level 7
device(config-cfm-md-d7)#ma-name link link-ma priority 7
device(config-cfm-md-d7-ma-link)#ccm-interval 1-second
device(config-cfm-md-d7-ma-link)#mep 3 down port eth 1/1
deviceconfig-cfm-md-d7-ma-link)#mep 4 down port eth 1/2
```

DUT3

```
device(config)#cfm-enable
device(config-cfm)#domain-name d7 level 7
device(config-cfm-md-d7)#ma-name link link-ma priority 7
device(config-cfm-md-d7-ma-link)#ccm-interval 1-second
device(config-cfm-md-d7-ma-link)#mep 2 down port eth 1/2
```

Show commands

The following **show** commands provide output for each component of the sample configuration.

Show cfm

Use the **show cfm** command to display the cfm configuration.

Syntax: show cfm

```
device#show cfm
Domain: d7
Index: 1
Level: 7
Maintenance association: link
Ma Index: 1
CCM interval: 1000 ms
LINK MA ID: N/A
Priority: 7
ETH-AIS TX: DISABLED
ETH-AIS RX: DISABLED
ETH-AIS Interval: 10 sec
MEP  Direction MAC             PORT       PORT-STATUS-TLV
==== ========= ============== ========== ===============
2    DOWN      0000.0011.86d1 ethe 1/1   N

device#show cfm
Domain: d7
Index: 1
Level: 7
Maintenance association: link
Ma Index: 1
CCM interval: 1000 ms
LINK MA ID: N/A
Priority: 7
ETH-AIS TX: DISABLED
ETH-AIS RX: DISABLED
ETH-AIS Interval: 10 sec
MEP  Direction MAC             PORT       PORT-STATUS-TLV
==== ========= ============== ========== ===============
3    DOWN      0000.0011.6351 ethe 1/1   N
4    DOWN      0000.0011.634b ethe 1/2   N

device#show cfm
Domain: d7
Index: 1
Level: 7
Maintenance association: link
Ma Index: 1
CCM interval: 1000 ms
LINK MA ID: N/A
Priority: 7
ETH-AIS TX: DISABLED
ETH-AIS RX: DISABLED
ETH-AIS Interval: 10 sec
MEP  Direction MAC             PORT       PORT-STATUS-TLV
==== ========= ============== ========== ===============
1    DOWN      0000.00ef.2a0b ethe 1/2   N DUT_3#
```

Show cfm connectivity

Use the **show cfm connectivity** command to display the cfm connectivity configuration.

Syntax: show cfm connectivity

```
device#show cfm connectivity
Domain: d7 Index: 1
Level: 7
Maintenance association: link
MA Index: 1
CCM interval: 1000 ms
LINK MA ID: N/A
Priority: 7
```

```
ETH-AIS TX: DISABLED
ETH-AIS RX: DISABLED
ETH-AIS Interval: 10 sec
RMEP  MAC            VLAN/PEER AGE   PORT   SLOTS STATE  AIS_STATE
====  ============== ========= ===== ====== ===== ====== =========
3     0000.0011.6351 N/A       696   1/1    1     OK     None

device#show cfm connectivity
Domain: d7 Index: 1
Level: 7
Maintenance association: link
MA Index: 1
CCM interval: 1000 ms
LINK MA ID: N/A
Priority: 7
ETH-AIS TX: DISABLED
ETH-AIS RX: DISABLED
ETH-AIS Interval: 10 sec
RMEP  MAC            VLAN/PEER AGE   PORT   SLOTS STATE  AIS_STATE
====  ============== ========= ===== ====== ===== ====== =========
1     0000.00ef.2a0b N/A       799   1/1    1     OK     None
2     0000.0011.86d1 N/A       799   1/2    1     OK     None


device#show cfm connectivity
Domain: d7 Index: 1
Level: 7
Maintenance association: link
MA Index: 1
CCM interval: 1000 ms
LINK MA ID: N/A
Priority: 7
ETH-AIS TX: DISABLED
ETH-AIS RX: DISABLED
ETH-AIS Interval: 10 sec
RMEP  MAC            VLAN/PEER AGE   PORT   SLOTS STATE  AIS_STATE
====  ============== ========= ===== ====== ===== ====== =========
4     0000.0011.634b N/A       869   1/2    1     OK     None
```

Loop back messages

CFM loopback

Use the **cfm loopback** command to display loopback messages.

Syntax: cfm loopback domain *domain-name* **ma** *ma-name* **src-mep** *ID* **target-mep** *ID*

The following output shows the Loopback messages.

```
device#cfm loopback domain d7 ma link src-mep 2 target-mep 3
DOT1AG: Sending 10 Loopback to 0000.0011.6351, timeout 10000 msec
Type Control-c to abort
Reply from 0000.0011.6351: time=1ms
Reply from 0000.0011.6351: time<1ms
Reply from 0000.0011.6351: time<1ms
Reply from 0000.0011.6351: time<1ms
Reply from 0000.0011.6351: time<1ms
Reply from 0000.0011.6351: time<1ms
Reply from 0000.0011.6351: time<1ms
Reply from 0000.0011.6351: time<1ms
Reply from 0000.0011.6351: time<1ms
Reply from 0000.0011.6351: time<1ms
sent = 10 number = 10 A total of 10 loopback replies received.
Success rate is 100 percent (10/10), round-trip min/avg/max=0/0/1 ms.
```

CFM linktrace

Use the **cfm linktrace** command to display linktrace messages.

Syntax: cfm linktrace domain *domain-name* **ma** *ma-name* **src-mep** *ID* **target-mep** *ID*

```
The following output shows the linktrace messages.
device#cfm linktrace domain d7 ma link src-mep 2 target-mep 3
Link trace functionality is not supported on Link-MA.
```

Port status TLV

- Port status TLV is carried in every CCM message and it carries the state of transmitting port
- The state can be either 1 or 2
  - 2 - Port state is Forwarding
  - 1 - Port state other than Forwarding
- Port status TLV is supported for sub-second timers from PBIF version 0x56 onwards
- Port status TLV is supported for all type of VLANs
  - CVLAN, SVAN, ISID and BVLAN
- Port status TLV is not applicable for Link MA

Configuring port status TLV

Port status TLV is optional and will be carried in a CCM message only if it is enabled in the MEP configuration. Port Status TLV for a specified MEP can be enabled using the following command.

Syntax: [no] mep *id dir* **tlv-type port-status-tlv port** *portId*

Sample configuration of port status TLV

The following configuration shows the process of enabling port status TLV at the MEP level.

```
device(config)#cfm-enable
device(config-cfm)#domain-name customer level 7
device(config-cfm-md-customer)#ma-name admin vlan-id 100 priority 7
device(config-cfm-md-customer-ma-admin)#ccm-interval 1-second
device(config-cfm-md-customer-ma-admin)#mep 1 down tlv-type port-status-tlv port eth 1/1
device(config)#cfm-enable
device(config-cfm)#domain-name customer level 7
device(config-cfm-md-customer)#ma-name admin vlan-id 100 priority 7
device(config-cfm-md-customer-ma-admin)#ccm-interval 1-second
device(config)#cfm-enable
device(config-cfm)#domain-name customer level 7
device(config-cfm-md-customer)#ma-name admin vlan-id 100 priority 7
device(config-cfm-md-customer-ma-admin)#ccm-interval 1-second
device(config-cfm-md-customer-ma-admin)#mep 2 down tlv-type port-status-tlv port ethe 1/2
```

Show commands

The following commands are used to display the port status tlv at MEP.

Show cfm

Use the **show cfm** command to display the cfm configuration.

Syntax: show cfm

Show cfm connectivity

Use the **show cfm connectivity** command to display the cfm connectivity configuration.

show cfm connectivity

The following commands display the received port status tlv state at RMEP.

Remote defect indication

Remote Defect Indication (RDI) is a single bit, is carried by CCM to convey the MEPs in MA about reception of CCM messages by receiving MEPs (RMEP)

- The absence of RDI in a CCM indicates that the transmitting MEP is receiving CCMs from all remote MEPs
- The presence of RDI indicates that transmitting MEP is not receiving CCM from one or more RMEPs (one or more RMEP failed is in state) attached to the MEP.
- RDI is supported for all type of VLANs
- CVLAN, SVAN, ISID and BVLAN
- RDI is supported for regular and sub-second CCM intervals

Limitations

- UPMEP and MIP on C-VLAN ESI is not supported if it is a client of S-VLAN ESI (in Provider Edge).
- UPMEP and MIP on S-VLAN ESI is not supported if it is a client of ISID ESI (in Backbone Edge).
- Sub-second CCM interval is not supported for CES.
- RDI is not applicable for Link MA.

Sample configuration of Remote Defect Indication

The following sample configuration shows the RDI configuration.

DUT1

```
device(config)#cfm-enable
device(config-cfm)#domain-name customer level 7
device(config-cfm-md-customer)#ma-name admin vlan-id 100 priority 7
device(config-cfm-md-customer-ma-admin)#ccm-interval 1-second
device(config-cfm-md-customer-ma-admin)#mep 1 down port eth 1/1
DUT2
device(config)#cfm-enable
device(config-cfm)#domain-name customer level 7
device(config-cfm-md-customer)#ma-name admin vlan-id 100 priority 7
device(config-cfm-md-customer-ma-admin)#ccm-interval 1-second
DUT3
device(config)#cfm-enable
device(config-cfm)#domain-name customer level 7
device(config-cfm-md-customer)#ma-name admin vlan-id 100 priority 7
device(config-cfm-md-customer-ma-admin)#ccm-interval 1-second
device(config-cfm-md-customer-ma-admin)#mep 2 down port ethe 1/2
DUT4
device(config)#cfm-enable
device(config-cfm)#domain-name customer level 7
device(config-cfm-md-customer)#ma-name admin vlan-id 100 priority 7
device(config-cfm-md-customer-ma-admin)#ccm-interval 1-second
device(config-cfm-md-customer-ma-admin)#mep 3 down port eth 1/3
```

Show commands

The following **show** commands provide output for each component of the sample configuration.

Show cfm connectivity

Assume link between DUT 2 and 4 goes down. RMEP(DUT4's MEP) will get failed in DUT1 and DUT3. At this time DUT1 and 2 will start transmitting CCM with RDI bit set since RMEP has failed.

```
device#show cfm connectivity
Domain: customer Index: 1
Level: 7
Maintenance association: admin
MA Index: 1
CCM interval: 1000 ms
VLAN ID: 100
Priority: 7
ETH-AIS TX: DISABLED
ETH-AIS RX: DISABLED
ETH-AIS Interval: 10 sec
RMEP MAC            VLAN/PEER  AGE   PORT SLOTS  STATE  AIS_STATE
==== ============== ========== === ===== =====  ======  =========
2    0000.00ef.2a0b    100     799  1/1    1    OK       None
3    0000.0011.86d1    100     400  1/1    1   FAILED    None
device#
```

Frame Loss Measurement

The Frame Loss Measurement feature (ETH-LM) maintains counters of received and transmitted data frames between a pair of MEPs. These counters are used to calculate the frame loss ratio.

Only single-ended ETH-LM, which is used for on-demand OAM, is supported. An MEP sends frames with an ETH-LM request information to its peer MEP and receives frames with ETH-LM reply information from its peer MEP to perform loss measurement. Frames which carry the Loss Measurement Message (LMM) PDU are called LMM frames. Frames which carry the Loss Measurement Reply (LMR) PDU are called LMR frames.

When the Loss Measurement Message (LMM) is configured the Frame Loss Measurement is enabled.

Device considerations

Frame loss measurement, one-way delay measurement, and synthetic loss measurement are not supported on the following Extreme NetIron CES and Extreme NetIron CER Series device models:

- BR-CER-2024C-4X-RT-AC
- BR-CER-2024C-4X-RT-DC
- BR-CER-2024F-4X-RT-AC
- BR-CER-2024F-4X-RT-DC
- BR-CES-2024C-4X-AC
- BR-CES-2024C-4X-DC
- BR-CES-2024F-4X-AC
- BR-CES-2024F-4X-DC

LMM over VLAN

Frame Loss Measurement can be done over VLAN where Connectivity Fault Management (CFM) is configured. In this use case, CFM should be enabled and down MEP should be configured on the VLAN end-points which should be monitored. LMM can be configured on the end-points for periodic measurements irrespective of the CFM connectivity. Ensure CFM connectivity is UP and running before the LMM session actually get started. Otherwise, an error will be thrown.

LMM over VPLS

Frame Loss Measurement can be done over VPLS and VLL where Connectivity Fault Management (CFM) is configured. In this use case, CFM should be enabled and UP MEP should be configured on the VPLS end-points which should be monitored. LMM can be configured on the end-points for periodic measurements irrespective of the CFM connectivity. Ensure CFM connectivity is UP and running before the LMM session actually get started. Otherwise, an error will be thrown.

Configuration considerations and limitations

As the Frame Loss Measurement feature uses ACL for getting data packet counters, it will be affected as follows:

- When there is an active LMM session and an L2 ACL is getting bounded, there will be some drop or frame loss expected as the LMM ACL is getting re-programmed.
- The responder should be started first before starting the initiator. Otherwise, the LMM packets will be dropped at the responder and no ACLs will be programmed, which may lead to inaccurate results.
- During termination, stop the initiator before the responder. Stopping the responder first may lead to inaccurate results as mentioned in the previous point.
- Only one LMM session will be active per source MEP per priority. This means eight active sessions per source MEP, one active for each priority.
- Maximum of 32 LMM sessions can be created per source MEP (irrespective of the priority).
- Maximum of 100 LMM sessions can be activated per system at any given point of time irrespective of the MD, MA, and MEP.
- LMM functionality not guaranteed if there exists multiple VPLS end points sharing the single peer for that VPLS instance. There should be a single VPLS end point.
- As the measurement is performed in the LP, LMM functionality is not supported over LAG, if member ports are from multiple slots. Loss will be measured only the ports on the same slot.
- If any ACLS are dropped on the same port or vport, the packets matching those ACLs will not be counted or taken into account as the LMM ACLs will be listed below the layer 2 ACLs.
- Protocol packets or packets trapped to CPU are not counted.
- To measure frame loss on untagged endpoints in VPLS, cos 8 should be used which covers all the priorities, as there is no priority carried in the untagged packet. This feature not supported via SNMP as the priority range supported is 0 to 7.
- LMM initiator and responder should monitor on the same priority, otherwise the packet will be discarded on the responder or initiator side which leads to inaccurate results.
- If cos 8 is configured on the source MEP, no other session with different priority is supported as cos 8 already counts all the priorities. Cos 8 not supported via SNMP as it is additional and not as per the standard MIB.
- If the start time is configured without the daily option, it will be shown in the running-config until it is explicitly removed by the "stop now" command.
- LMM over VLL is not supported.

- For Layer 3 traffic, with VPLS the incoming priority in the data packet gets modified by DSCP bits and gets changed in the egress side. As the ingress and egress priorities are different in VPLS data traffic, only cos 8 should be used which monitors on all the priorities.
- For individual packet priority monitoring with VPLS L3 traffic, VLAN PCP and DSCP bits should be the same in the ingress traffic.

Supported configurations

The following functionalities are common for both VPLS and VLAN endpoints.

Monitor LMM on demand

The LMM can be started immediately whenever required and can be stopped after some period of time. The frame loss ratio will be calculated after every measurement interval configured and can be viewed whenever required. This use case will be useful whenever the administrator wants to measure immediately (on demand).

Monitor LMM for a fixed interval of time

The LMM can be configured to start at any fixed time (more than the current time) and can be stopped after some period of time (more than the start time). The frame loss ratio will be calculated after every measurement interval configured and can be viewed whenever required. This use case will be useful whenever the administrator wants to measure at particular time interval.

Monitor LMM after some relative time

The LMM can be configured to start after any relative time and can be stopped after some period of time (more than the start time). The frame loss ratio will be calculated after every measurement interval configured and can be viewed whenever required. This use case will be useful whenever the administrator wants to trigger the measurement after some duration.

Monitor LMM daily for fixed interval of time

The LMM can be configured to start daily at any fixed time and stop after some period of time (more than the start time). The frame loss ratio will be calculated after every measurement interval configured and can be viewed whenever required. This use case will be useful whenever the administrator wants to measure daily at particular time interval.

LMM configurations common for VLAN and VPLS

Before configuring Loss Measurement Message (LMM), Connectivity Fault Management (CFM) must be configured for the VLAN or VPLS. Refer to OAM chapter for the procedures to configure CFM for VLAN or VPLS.

The configuration of Loss Measurement Message (LMM) is the same process for both VLANs and VPLS.

LMM initiator session configuration

Use the following procedure to configure the LMM initiator session.

1. LMM initiator session creation.

   Create the Loss Measurement Message (LMM) session.

   ```
   device(config-cfm)#loss-measurement lmm initiator 1
   device(config-cfm-loss-measurement-lmm-initiator-1)#
   ```

   Syntax: lmm initiator *session_id*

2. LMM Initiator session configuration.

   Configure the LMM session.

   ```
   device(config-cfm-loss-measurement-lmm-initiator-1)#domain md1 ma ma1 src-mep 1 target-mep 2
   ```

   Syntax: domain *name* **ma** *name* **src-mep** *id* **target-mep** *id*

3. LMM session CoS configuration.

   ```
   device(config-cfm-loss-measurement-lmm-initiator-1)#Cos 1
   device(config-cfm-loss-measurement-lmm-initiator-1)#
   ```

   Syntax: Cos *value*

4. LMM session Tx-interval configuration.

   ```
   device(config-cfm-loss-measurement-lmm-initiator-1)#Tx-interval 10
   device(config-cfm-loss-measurement-lmm-initiator-1)#
   ```

   Syntax: Tx-interval *timer_value*

5. LMM session measurement-interval configuration.

   ```
   device(config-cfm-loss-measurement-lmm-initiator-1)#Measurement-interval 10
   device(config-cfm-loss-measurement-lmm-initiator-1)#
   ```

   Syntax: Measurement-interval *timer_value*

6. LMM session threshold configuration

   ```
   device(config-cfm-loss-measurement-lmm-initiator-1)#threshold forward average 5000 maximum 10000
   device(config-cfm-loss-measurement-lmm-initiator-1)#threshold backward average 5000 maximum 10000
   ```

   Syntax: threshold forward | **backward average** *value* **maximum** *value*

LMM responder session configuration

Use the following procedure to configure the LMM responder session.

1. LMM responder session creation.

   ```
   device(config-cfm)#loss-measurement lmm responder 1
   device(config-cfm-loss-measurement-lmm-responder-1)#
   ```

   Syntax: loss-measurement lmm responder *session_id*

2. LMM responder session configuration.

   ```
   device(config-cfm-loss-measurement-lmm-responder-1)#domain md1 ma ma1 src-mep 2 target-mep 1
   device(config-cfm-loss-measurement-lmm-responder-1)#
   ```

   Syntax: domain *name* **ma** *name* **src-mep** *id* **target-mep** *id*

3. LMM session CoS configuration.

```
device(config-cfm-loss-measurement-lmm-responder-1)#Cos 1
device(config-cfm-loss-measurement-lmm-responder-1)#
```

Syntax: Cos *value*

Starting LMM session responder

Use the **start** command to start the session responder.

```
device(config-cfm-loss-measurement-lmm-responder-1)#start now
```

Syntax: start { **now** | **after** *HH:MM:SS* | *HH:MM:SS* [**daily**] }

now starts the session immediately.

after HH:MM:SS starts the session after the indicated time interval.

HH:MM:SS starts the session at the indicated time.

HH:MM:SS daily starts the session at the indicated time every day.

Starting LMM Session Initiator

Use the **start** command to start the session initiator.
device(config-cfm-loss-measurement-lmm-initiator-1)#start after 01:10:00

Syntax:start{**now**|**after***HH:MM:SS*|*HH:MM:SS*[**daily**]}

now starts the session immediately.

after HH:MM:SS starts the session after the indicated time interval.

HH:MM:SS starts the session at the indicated time.

HH:MM:SS daily starts the session at the indicated time every day.

No configuration changes are supported once the session is started or triggered. Only the "Stop now" configuration is allowed which stops the session.

Session will not start if the target MEP not available. Session will be started, only if the target MEP is in FAILED state or OK state.

Stopping LMM Session Responder

Use the **stop** command to stop the session responder.

```
device(config-cfm-loss-measurement-lmm-responder-1)#stop now
```

Syntax: stop { **now** | **after** *HH:MM:SS* | *HH:MM:SS* [**daily**] }

now stops the session immediately.

after HH:MM:SS stops the session after the indicated time interval.

HH:MM:SS stops the session at the indicated time.

HH:MM:SS daily stops the session at the indicated time every day.

Stopping LMM Session Initiator

Use the **stop** command to stop the session initiator.

```
device(config-cfm-loss-measurement-lmm-initiator-1)#stop now
```

Syntax: stop { **now** | **after** *HH:MM:SS* | *HH:MM:SS* [**daily**] }

now stops the session immediately.

after HH:MM:SS stops the session after the indicated time interval.

HH:MM:SS stops the session at the indicated time.

HH:MM:SS daily stops the session at the indicated time every day.

Configuration examples

Configuration example for LMM over VLAN

CE-1 configuration

```
device(config)# cfm
device(config-cfm)#loss-measurement lmm initiator 1
device(config-cfm-loss-measurement-lmm-initiator-1)#domain md1 ma ma1 src-mep 3 target-mep 4
device(config-cfm-loss-measurement-lmm-initiator-1)#Cos 2
device(config-cfm-loss-measurement-lmm-initiator-1)#tx-interval 10
device(config-cfm-loss-measurement-lmm-initiator-1)#measurement-interval 10
```

CE-2 configuration

```
device(config)# cfm
device(config-cfm)# loss-measurement lmm responder 1
device(config-cfm-loss-measurement-lmm-responder-1)#domain md1 ma ma1 src-mep 4 target-mep 3
device(config-cfm-loss-measurement-lmm-responder-1)#Cos 2
```

Configuration example for VPLS tagged endpoints

PE-1 configuration (Initiator)

```
device(config-cfm)# loss-measurement lmm initiator 1
device(config-cfm-loss-measurement-lmm-initiator-1)#domain md1 ma ma1 src-mep 3 target-mep 4
device(config-cfm-loss-measurement-lmm-initiator-1)#Cos 2
device(config-cfm-loss-measurement-lmm-initiator-1)#tx-interval 10
device(config-cfm-loss-measurement-lmm-initiator-1)#measurement-interval 10
```

PE-2 configuration (Responder)

```
device(config-cfm)# loss-measurement lmm responder 1
device(config-cfm-loss-measurement-lmm-responder-1)#domain md1 ma ma1 src-mep 4 target-mep 3
device(config-cfm-loss-measurement-lmm-responder-1)#Cos 2
```

Configuration example for VPLS untagged endpoints

PE-1 configuration (Initiator)

```
device(config-cfm)# loss-measurement lmm initiator 1
device(config-cfm-loss-measurement-lmm-initiator-1)#domain md1 ma ma1 src-mep 3 target-mep 4
device(config-cfm-loss-measurement-lmm-initiator-1)#Cos 8
device(config-cfm-loss-measurement-lmm-initiator-1)#tx-interval 10
device(config-cfm-loss-measurement-lmm-initiator-1)#measurement-interval 10
```

PE-2 configuration (Responder)

```
device(config-cfm)# loss-measurement lmm responder 1
device(config-cfm-loss-measurement-lmm-responder-1)#domain md1 ma ma1 src-mep 4 target-mep 3
device(config-cfm-loss-measurement-lmm-responder-1)#Cos 8
```

Configuration example for VPLS tagged and untagged endpoints

PE-1 configuration (Initiator)

```
device(config-cfm)# loss-measurement lmm initiator 1
device(config-cfm-loss-measurement-lmm-initiator-1)#domain md1 ma ma1 src-mep 3 target-mep 4
device(config-cfm-loss-measurement-lmm-initiator-1)#Cos 8
device(config-cfm-loss-measurement-lmm-initiator-1)#tx-interval 10
device(config-cfm-loss-measurement-lmm-initiator-1)#measurement-interval 10
```

PE-2 configuration (Responder)

```
device(config-cfm)# loss-measurement lmm responder 1
device(config-cfm-loss-measurement-lmm-responder-1)#domain md1 ma ma1 src-mep 4 target-mep 3
device(config-cfm-loss-measurement-lmm-responder-1)#Cos 8
```

Starting LMM Sessions

Start the responder before starting the initiator.

CE-2 configuration

```
device(config)# cfm
device(config-cfm)# loss-measurement lmm responder 1
device(config-cfm-loss-measurement-lmm-responder-1)#start now
```

CE-1 configuration

```
device(config)# cfm
device(config-cfm)# loss-measurement lmm initiator 1
device(config-cfm-loss-measurement-lmm-initiator-1)#start now
```

Stopping LMM sessions

Stop the initiator before stopping the responder.

CE-1 configuration

```
device(config)# cfm
device(config-cfm)# loss-measurement lmm initiator 1
device(config-cfm-loss-measurement-lmm-initiator-1)#stop now
```

CE-2 configuration

```
device(config-cfm)# loss-measurement lmm responder 1
device(config-cfm-loss-measurement-lmm-responder-1)#stop now
```

Clearing history statistics per session

```
device(config-cfm)# loss-measurement lmm initiator 1
device(config-cfm-loss-measurement-lmm-initiator-1)#clear-stat
```

Clearing history statistics globally

```
device(config-cfm)# loss-measurement clear-stat
```

Syslog messages

Syslogs will be raised for the following cases:

- When the LMM session is started.
- When the LMM session is stopped.
- When the Average Frame Loss Ratio is greater than the Threshold Average Frame Loss Ratio.
- When the Maximum Frame Loss Ratio is greater than the Threshold Maximum Frame Loss Ratio.

Syslog message display output

The following are the Syslog message outputs displayed for various cases:

When the LMM session started

```
<Syslog>: Y.1731: The LMM session started for MA index 1, MD index 1, MEP id 2 Session index 1
```

When the LMM session stopped

```
<Syslog>: Y.1731: The LMM session started for MA index 1, MD index 1, MEP id 2 Session index 1
```

When the Average Frame Loss Ratio greater than Threshold Average Frame Loss Ratio

```
<Syslog>: Y.1731: The LMM session for MA index 1, MD index 1, MEP id 2 Session index 1 has crossed the
forward average threshold value, with value 35000.
```

When the Maximum Frame Loss Ratio greater than Threshold Maximum Frame Loss Ratio

```
<Syslog>: Y.1731: The LMM session for MA index 1, MD index 1, MEP id 2 Session index 1 has crossed the
forward maximum threshold value, with value 60000.
```

One-way Delay Measurement

One-way delay measurement can be used for on-demand or proactive OAM to measure frame delay and frame delay variation. Frame delay and frame delay variation measurements are performed by sending periodic frames with Ethernet Delay Measurement information to the peer MEP and receiving frames with Ethernet Delay Measurement information from the peer MEP during proactive measurement session and/or the diagnostic interval. Each MEP may perform frame delay and frame delay variation measurement.

When a MEP is enabled to generate frames with one-way delay measurement information, it periodically sends frames with one-way delay measurement information to its peer MEP in the same ME. When a MEP is enabled to generate frames with one-way delay

measurement information, it also expects to receive frames with one-way delay measurement information from its peer MEP in the same ME.

A MIP is transparent to the frames with one-way delay measurement information and therefore does not require any information to support one-way delay measurement functionality.

A MEP transmits frames with one-way delay measurement information with the following information element:

- TxTimeStampf: Timestamp at the transmission time of one-way delay measurement frame

The receiving MEP can compare this value with the RxTimef, the time at the reception of a one-way delay measurement frame and calculate the one-way frame delay as:

- Frame Delay = RxTimef - TxTimeStampf

Configuration considerations

- Only one one-way delay measurement session will be active per source MEP per priority.
- Maximum of 32 one-way delay measurement sessions can be created per source MEP.
- Maximum of 100 one-way delay measurement sessions can be activated per system at any given point of time.
- There can be maximum 16 one-way delay measurement sessions (8 Initiator sessions and 8 Receiver sessions) which can be active per MEP.
- The one-way delay measurement receiver session should be started before starting the initiator session. Otherwise, the one-way delay measurement packets will be dropped at the receiver, which may lead to inaccurate results.
- The NTP should be disabled and the system clock should be set explicitly through CLI when the one-way delay has to be measured between a Extreme device and another vendor device.

One-way Delay Measurement

In this case, each MEP sends frame with one-way Ethernet Delay Measurement information to its peer MEP to facilitate one-way frame delay and/or one-way frame delay variation measurements at the peer MEP.

One-way Delay Measurement transmission

When configured for one-way delay measurement, a MEP periodically transmits one-way delay measurement frames with the TxTimeStampf value.

One-way Delay Measurement reception

When configured for one-way delay measurement, a MEP, upon receiving a valid one-way delay measurement frame, uses the following values to make one-way frame delay measurement. A one-way delay measurement frame with a valid MEG level and a destination MAC address equal to the receiving MEP's MAC address is considered to be a valid one-way delay measurement frame. These values serve as input to the one-way frame delay variation measurement:

- One-way delay measurement frame's TxTimeStampf value
- RxTimef, which is the time at reception of the one-way delay measurement frame
- Frame Delayone-way = RxTimef - TxTimeStampf

Use cases

The following use cases are supported for one-way delay measurement.

One-way Delay Measurement over VLAN

One-way delay measurement can be done over VLAN where CFM is configured. In this use case, CFM should be enabled and the down MEP should be configured on the VLAN end-points (tagged ports) for periodic measurements irrespective of the CFM connectivity. Verify CFM connectivity is up and running before the one-way delay measurement session is actually started. Otherwise, this may cause an error.

NOTE
The one-way delay measurement should be configured over CFM, where CFM should be configured over the VLAN and the down MEPs should be configured only on the tagged ports.

One-way Delay Measurement over VPLS

One-way delay measurement can be done over VPLS where CFM is configured. In this use case, CFM should be enabled and the up MEP should be configured on the VPLS end-points which should be monitored. One-way delay measurement can be configured on the end-points for periodic measurements irrespective of the CFM connectivity. Ensure CFM connectivity is up and running before the one-way delay measurement session is actually started. Otherwise, this may cause an error.

NOTE
If the VPLS end is configured as an untagged port, then the one-way delay measurement packet will be considered as no priority and one-way delay measurement will be measured with priority 8. If priority 8 is configured for the one-way delay measurement session, then all the other priority one-way delay measurement sessions under the same MEP will not be allowed.

One-way Delay Measurement over VLL

One-way delay measurement can be done over VLL where CFM configured. In this use case, CFM should be enabled and the up MEP should be configured on the VLL end-points which should be monitored. One-way delay measurement can be configured on the end-points for periodic measurements irrespective of the CFM connectivity. Ensure CFM connectivity is up and running before the one-way delay measurement session is actually started. Otherwise, this may cause an error.

Supported configurations

The following are the additional supported configurations for monitoring one-way delay measurement based on different time intervals. The functionality discussed below are common for both VPLS and VLAN.

Monitor one-way delay measurement on demand

In this case, one-way delay measurement can be started immediately whenever required and can be stopped after a period of time. The one-way delay will be calculated after receiving each one-way delay measurement packet and delay statistics will be calculated for every measurement interval configured. It can be viewed whenever required. This use case is useful whenever the you want to measure immediately (on demand).

Monitor one-way delay fixed interval of time

In this case, the one-way delay measurement can be configured to start at any fixed time and can be stopped after a period of time. The one-way delay will be calculated after receiving each one-way delay measurement packet and delay statistics will be calculated for every measurement interval configured. It can be viewed whenever required. This use case is useful whenever the administrator wants to measure at particular time interval.

Monitor one-way delay after relative time

In this case, the one-way delay measurement can be configured to start after a relative time and can be stopped after a period of time. The one-way delay will be calculated after receiving each one-way delay measurement packet and delay statistics will be calculated every measurement interval configured. It can be viewed whenever required. This use case is useful whenever the administrator wants to trigger the measurement after some duration.

Monitor one-way delay daily for fixed interval of time

In this case, the one-way delay measurement can be configured to start daily at any fixed time and stop after some period of time. The one-way delay will be calculated after receiving each one-way delay measurement packet and delay statistics will be calculated for every measurement interval configured. It can be viewed whenever required. This use case is useful whenever the administrator wants to measure daily at particular time interval.

Configuration procedure

CFM configuration for VLAN

VLAN configuration

VLAN creation.

```
device(config)#vlan 20
device(config-vlan-20)#tagged ethernet 1/1
```

Syntax: vlan *id*

CFM configuration

1. Enabling CFM.

```
device(config)#cfm-enable
device(config-cfm)#
```

Syntax: cfm-enable

2. Domain configuration.

```
device(config)#cfm-enable
device(config-cfm)#domain-name md1 level 7
device(config-cfm-md-md1)#
```

Syntax: domain-name *md_name* [**id** *id*] **level** *level*

3. MA configuration.

```
device(config-cfm-md-md1)#ma-name ma1 vlan 20 priority 4
device(config-cfm-md-md1-ma-ma1)
```

Syntax: ma-name *ma_name* [**id** *id*] **vlan-id** *vlan* | **vpls-id** *vpls* **priority** *priority*

4. MEP configuration.

```
device(config-cfm-md-md1-ma-ma1)#mep 1 down port ethernet 1/1
```

Syntax: mep *id* { **down** | **up** } **port ethernet** *slot/port*

CFM configuration for VPLS and VLL

Creation of VPLS

```
device(config)#router mpls
device(config-mpls)#vpls vpls100 100
device(config-mpls-vpls-vpls100)#vlan 100
device(config-mpls-vpls-vpls100-vlan-10)#tagged Ethernet 1/1
```

Syntax: vpls *vpls-nameid*

Syntax: vlan *vlan-id*

Syntax: tagged ethernet *slot/port*

Creation of VLL

```
device(config)#router mpls
device(config-mpls)#vll vll100 100
device(config-mpls-vll-vll100)#vlan 100
device(config-mpls-vll-vll100-vlan-10)#tagged Ethernet 1/1
```

Syntax: vll *vll-nameid*

Syntax: vlan *vlan-id*

Syntax: tagged ethernet *slot/port*

CFM configuration

1. Enable CFM.

```
device(config)#cfm-enable
device(config-cfm)#
```

Syntax: cfm-enable

2. Configure the domain.

```
device(config)#cfm-enable
device(config-cfm)#domain-name md1 level 7
device(config-cfm-md-md1)#
```

Syntax: domain-name *md_name* **id** *id* **level** *level*

3. Configure MA.

```
device(config-cfm-md-md1)#ma-name ma1 vpls-id 100 priority 4
device(config-cfm-md-md1-ma-ma1)#
```

Syntax: ma-name *ma_name* [**id** *id*] **vlan-id** *vlan* | **vpls-id** *vpls* **priority** *priority*

Ma_name - Maintenance Association Name

4. Configure MEP.

```
device(config-cfm-md-md1-ma-ma1)#mep 1 up vlan 100 port ethernet 1/1
```

Syntax: mep *id* **down** | **up vlan** *vlan* **port ethernet** *slot/port*

One-way delay measurement configuration

NOTE
The following configuration is common for common for VLAN, VPLS, and VLL.

One-way delay measurement initiator session configuration

1. One-way delay measurement initiator session creation

```
device(config)#cfm
device(config-cfm)# oneway-dm initiator 1
device(config-cfm-oneway-dm-initiator-1)# domain md1 ma ma1 src-mep 1 target-mep 101
```

Syntax: oneway-dm initiator *session-index*

session_index - Is used to configure the one-way delay measurement initiator session index (1-1000)

Syntax: domain *md_name* **ma** *ma_name* **src-mep** *id* **target-mep** *id*

Md_name - Domain Name

Ma_name - Maintenance Association Name

Src-Mep ID - Source MEP

Target-MEP ID - Destination MEP

2. One-way delay measurement initiator session configuration

```
device(config-cfm-oneway-dm-initiator-1)#  cos 4
device(config-cfm-oneway-dm-initiator-1)# tx-interval 10
```

Syntax: cos *value*

Syntax: tx-interval *sec*

Cos (value) - Priority Value (1-7) (optional - Default value 7)

Tx-interval options include {start | stop} {now | after <HH:MM:SS> | <HH:MM:SS> | daily}

One-way delay measurement receiver session configuration

1. One-way delay measurement receiver session creation

```
device(config)#cfm
device(config-cfm)# oneway-dm receiver 1
device(config-cfm-oneway-dm-receiver-1)# domain md1 ma ma1 src-mep 101 target-mep 1
```

Syntax: oneway-dm receiver *session-index*

Syntax: domain *md_name* **ma** *ma_name* **src-mep** *id* **target-mep** *id*

2. One-way delay measurement receiver session configuration

```
device(config-cfm-oneway-dm-receiver-1)#  cos 4
device(config-cfm-oneway-dm-receiver-1)# measurement-interval 10
```

Syntax: cos *value*

Syntax: measurement-interval *sec*

3. One-way delay measurement receiver session threshold configuration

```
device(config-cfm-oneway-dm-receiver-1)#  threshold max 50
device(config-cfm-oneway-dm-receiver-1)#  threshold average 25
```

Syntax: threshold max *value*

Syntax: threshold average *value*

Starting one-way delay measurement receiver session

A receiver session can be started immediately, after a specified amount of time, once at a specific time, or a specific time daily.

```
device(config-cfm-oneway-dm-receiver-1)#  start now
```

Syntax: start now | **after** *HH:MM:SS* | *HH:MM:SS* **daily**

1. Start the one-way delay measurement receiver session after a period of time.

```
device(config-cfm-oneway-dm-receiver-1)#  start after 01:30:00
```

The example above will start after 1 hour and 30 minutes.

2. Start 1DM Receiver Session exactly at given time.

```
device(config-cfm-oneway-dm-receiver-1)#  start 09:30:00
```

The example above will be started exactly at 09:30 AM.

3. Start 1DM Receiver Session daily at given time.

```
device(config-cfm-oneway-dm-receiver-1)#  start 09:30:00 daily
```

The example above will be started daily exactly at 09:30 AM.

Starting the one-way delay measurement session initiator

A session can be started immediately, after a specified amount of time, once at a specific time, or a specific time daily.

```
device(config-cfm-oneway-dm-initiator-1)#  start now
```

Syntax: start now | **after** *HH:MM:SS* | *HH:MM:SS* **daily**

Stopping the one-way delay measurement initiator session

A session can be stopped immediately, after a specified amount of time, once at a specific time, or a specific time daily.

```
device(config-cfm-oneway-dm-initiator-1)#  stop now
```

Syntax: stop now | **after** *HH:MM:SS* | *HH:MM:SS* **daily**

Stopping the one-way delay measurement receiver session

A receiver session can be stopped immediately, after a specified amount of time, or at a specific time.

```
device(config-cfm-oneway-dm-receiver-1)#  stop now
```

Syntax: stop now | **after** *HH:MM:SS* | *HH:MM:SS*

NOTE
The one-way delay measurement Receiver session should be started before starting the one-way delay measurement Initiator session. Also, the one-way delay measurement Initiator session should be stopped before stopping the one-way delay measurement Receiver session.

NOTE
Relative time is converted to absolute time otherwise it would not point to the expected time after a config-save and reboot. This case is applicable to both start and stop times.

Configuration examples

Sample configuration of one-way delay measurement over VLAN

VLAN configurations

CE-1 configuration

```
device(config)# interface ethernet 1/1
device(config-if-e10000-1/1)# enable
device(config-if-e10000-1/1)# exit
device(config)# vlan 10
device(config-vlan-10)# tagged ethernet 1/1
```

CE-2 configuration

```
device(config)# interface ethernet 1/1
device(config-if-e10000-1/1)# enable
device(config-if-e10000-1/1)# exit
device(config)# vlan 10
device(config-vlan-10)# tagged ethernet 1/1
```

CFM configurations

CE-1 configuration

```
device(config)# cfm-enable
device(config-cfm)# domain md1 level 7
device(config-cfm-md-md1)# ma ma1 vlan 10 priority 4
device(config-cfm-md-md1-ma-ma1)# mep 1 down port ethernet 1/1
```

CE-2 configuration

```
device(config)# cfm-enable
device(config-cfm)# domain md1 level 7
device(config-cfm-md-md1)# ma ma1 vlan 10 priority 4
device(config-cfm-md-md1-ma-ma1)# mep 101 down port ethernet 1/1
```

One-way delay measurement configurations:

CE-1 configuration

```
device(config)# cfm
device(config-cfm)# oneway-dm initiator 1
device(config-cfm-oneway-dm-initiator-1)#domain md1 ma ma1  src-mep 1 target-mep 2
device(config-cfm-oneway-dm-initiator-1)#tx-interval 10
```

CE-2 configuration

```
device(config)# cfm
device(config-cfm)# oneway-dm receiver 2
device(config-cfm-oneway-dm-receiver-2)# domain md1 ma ma1  src-mep 2 target-mep 1
device(config-cfm-oneway-dm-receiver-2)# measurement-interval 10
```

Starting 1DM sessions

CE-1 configuration

```
device(config-cfm-oneway-dm-initiator-1)# start now
```

CE-2 configuration

```
device(config-cfm-oneway-dm-receiver-2)# start now
```

Stopping 1DM sessions

CE-1 configuration

```
device(config-cfm-oneway-dm-initiator-1)# stop now
```

CE-2 configuration

```
device(config-cfm-oneway-dm-receiver-2)# stop now
```

Sample configuration one-way delay measurement over VPLS or VLL

VPLS configurations

PE-1 configuration

```
device(config)# interface ethernet 1/1
device(config-if-e10000-1/1)# enable
device(config-if-e10000-1/1)# exit
device(config)# router mpls
device(config-mpls)# vpls vpls100 100
device(config-mpls-vpls-vpls100)# vlan 10
device(config-mpls-vpls-vpls100-vlan-10)# tagged ethernet 1/1
device(config-mpls-vpls-vpls100-vlan-10)# end
```

PE-2 configuration

```
device(config)# interface ethernet 1/1
device(config-if-e10000-1/1)# enable
device(config-if-e10000-1/1)# exit
```

```
device(config)# router mpls
device(config-mpls)# vpls vpls100 100
device(config-mpls-vpls-vpls100)# vlan 10
device(config-mpls-vpls-vpls100-vlan-10)# tagged ethernet 1/1
device(config-mpls-vpls-vpls100-vlan-10)# end
```

VLL configurations

PE-1 configuration

```
device(config)# interface ethernet 1/1
device(config-if-e10000-1/1)# enable
device(config-if-e10000-1/1)# exit
device(config)# router mpls
device(config-mpls)# vll vll100 100
device(config-mpls-vll-vll100)# vlan 10
device(config-mpls-vll-vll100-vlan-10)# tagged ethernet 1/1
device(config-mpls-vll-vll100-vlan-10)# end
```

PE-2 configuration

```
device(config)# interface ethernet 1/1
device(config-if-e10000-1/1)# enable
device(config-if-e10000-1/1)# exit
device(config)# router mpls
device(config-mpls)# vll vpls100 100
device(config-mpls-vll-vll100)# vlan 10
device(config-mpls-vll-vll100-vlan-10)# tagged ethernet 1/1
device(config-mpls-vll-vll100-vlan-10)# end
```

CFM configurations

PE-1 configuration

```
device(config)# cfm-enable
device(config-cfm)# domain md1 level 7
device(config-cfm-md-md1)# ma ma1 vpls 100 priority 4
device(config-cfm-md-md1-ma-ma1)# mep 1 up vlan 10 port ethernet 1/1
```

PE-2 configuration

```
device(config)# cfm-enable
device(config-cfm)# domain md1 level 7
device(config-cfm-md-md1)# ma ma1 vpls 100 priority 4
device(config-cfm-md-md1-ma-ma1)# mep 101 up vlan 10 port ethernet 1/1
```

One-way delay measurement configurations (common for VPLS/VLL)

NOTE

The one-way delay measurement receiver session should be started first before starting the initiator session. Otherwise, the one-way delay measurement packets will be dropped at the receiver, which may lead to inaccurate results.

CE-1 configuration

```
device(config)# cfm
device(config-cfm)# oneway-dm initiator 1
device(config-cfm-oneway-dm-initiator-1)#domain md1 ma ma1  src-mep 1 target-mep 2
device(config-cfm-oneway-dm-initiator-1)#tx-interval 10
```

CE-2 configuration

```
device(config)# cfm
device(config-cfm)# oneway-dm receiver 2
```

```
device(config-cfm-oneway-dm-receiver-2)# domain md1 ma ma1  src-mep 2 target-mep 1
device(config-cfm-oneway-dm-receiver-2)# measurement-interval 10
```

Starting one-way delay measurement sessions

CE-1 configuration

```
device(config-cfm-oneway-dm-initiator-1)# start now
```

CE-2 configuration

```
device(config-cfm-oneway-dm-receiver-2)# start now
```

Stopping one-way delay measurement sessions

CE-1 configuration

```
device(config-cfm-oneway-dm-initiator-1)# stop now
```

CE-2 configuration

```
device(config-cfm-oneway-dm-receiver-2)# stop now
```

Show commands

The **show cfm oneway-dm** *session_index* command is used to display the session for a specified index. If a session index is not specified all available session indices will be displayed.

Syntax: show cfm oneway-dm *session_index*

```
device# show cfm oneway-dm 101
One Way DM Session Index  : 101
------------------------------
1DM Session Index         : 101
Status                    : Running
Session Type              : Receiver
Domain                    : MD4
MA                        : MA4.1
Source MEP                : 2
Target MEP                : 1
Cos                       : 2
Measurement-Interval(in M : 30
Start time                : 22:56:41
Start time type           : Immediate
Stop time                 : 22:56:22
Stop time type            : Immediate
Threshold Configuration
-----------------------
Threshold Average         : 0
Threshold Max             : 0
```

The **show cfm oneway-dm statistics** *session_index* command is used to display the latest 32 measurement statistics for a specified session index. If a session index is not specified the statistics for all available session indices will be displayed.

```
device# show cfm oneway-dm statistics
```

Syntax: show cfm oneway-dm statistics *session_index*

NOTE

The statistics command is valid only for receiver session Indices. An error will occur for initiator session indices.

The following information will be displayed in the show command output:

```
device# show cfm oneway-dm statistics
HISTORY TABLE :
Flag - S:Suspect, All measurements are in us unit.
-------------------------------------------------------------------------------
Index Flag Start    Elapsed   Avg Delay   Max Delay   Min Delay     FDV Avg    FDV Max        FDV Min
-------------------------------------------------------------------------------------------------
89    -    16:26:41 00:30:00   708306.712  710750.194  705415.159      99.324    3983.960      59.430
88    -    15:55:41 00:30:00   706484.225  709951.529  703322.004     121.515    6629.525      59.549
87    S    15:26:41 00:30:00  3121638.643 4002430.793  704518.559   18410.613 3297707.009      58.890
86    S    14:55:41 00:30:00  4003266.051 4010147.033 3997113.188     160.754    9650.315      59.085
85    S    14:26:41 00:30:00  4007927.518 4011404.633 4004814.353     125.277    6287.675      59.505
-------------------------------------------------------------------------------------------------
```

NOTE

If a one-way delay measurement is skipped for any one-way delay measurement packet within the measurement interval, then it will be marked as suspect.

The **show cfm oneway-dm statistics** *session_index* row-index *row-index* command is used to display details for a specific session index.

Syntax: show cfm oneway-dm statistics *session-index* **row-index** *row-index*

```
device# show cfm oneway-dm statistics 1 row-index 2
One Way DM Session Index  : 1
-------------------------------------------------------------
HISTORY ENTRY :
-------------------------------------------------------------
Row Index                 : 2
Flag                      : -
Start Time                : 18:27:39
Elapsed Time              : 00:00:11
Valid RX Count            : 10
Total RX Count            : 10
Avg Delay                 :     13.115
Max Delay                 :     13.287
Min Delay                 :     12.956
Avg Frame Delay Variation :      0.110
Max Frame Delay Variation :      0.218
Min Frame Delay Variation :      0.016
```

Syslog messages

The following are the Syslog message outputs displayed for various cases.

When the one-way delay measurement session is started.

```
SYSLOG:  <time>  Y.1731: The DM session started for MA index <ma index>, MD index <md index>, MEP id <med
id> Session index <id>
```

When the one-way delay measurement session is stopped.

```
SYSLOG:  <time>  Y.1731: The DM session stopped for MA index <ma index>, MD index <md index>, MEP id <med
id> Session index <id>
```

When the Average delay is greater than the Threshold Average delay.

```
SYSLOG: <timestamp> Y.1731: The DM session for MA index <ma index>, MD index <md index>, MEP id <med id>
Session index <id> has crossed the forward average threshold, with value <value>
```

When the Maximum delay is greater than the Threshold Maximum delay.

```
SYSLOG: <timestamp> Y.1731: The DM session for MA index <ma index>, MD index <md index>, MEP id <med id>
Session index <id> has crossed the forward maximum threshold, with value <value>
```

When the Destination MEP moves to, or is already in a FAILED state, when the session is Active.

```
<Syslog>: 1DM Session <Id> not started as the RMEP <Id> is in FAILED state.
```

Synthetic loss measurement

Synthetic loss measurement (SLM) is part of the ITU-T Y.1731 standard. It can be used to periodically measure Frame Loss, Forward Loss Ratio (FLR), and Frame delay between a pair of point to point MEPs. Measurements are made between two MEPs belonging to the same domain and MA.

The procedure involves a Sender MEP sending an SLM Protocol Data Unit (PDU) once per transmit interval (e.g. 1 second, 10 seconds, 1 minute). The Remote MEP responds with a Synthetic Loss Reply (SLR). The messages are used to collect the number of SLMs and SLRs transmitted and received by the two MEPs.

Configuration considerations

- An MEP instance must be configured before configuring Synthetic loss measurement (SLM).
- A Synthetic loss measurement instance cannot be started if the target MEP is not known. However, the session can start if the remote MEP is known but in a failed state.
- A maximum of 32 SLM sessions can be created per source MEP.
- History data generated after every measurement cycle for a particular SLM session overwrites the oldest entry after 32 history entries.
- Only one Synthetic loss measurement session will be active per source MEP per COS.
- At any point of time a maximum of 100 SLM sessions can be activated on a node. This number is shared across all Y1731 modules.
- A maximum of 1000 SLM sessions can be configured over a system. This number is shared across all Y1731 modules.
- Synthetic loss measurement functionality will not be accurate if VPLS is point to multipoint.
- Synthetic loss measurement support is currently not available for MLX and XMR devices.
- Configuration of tx-interval, measurement interval, threshold, and clear statistics is possible only under the initiator mode.
- The same set of attributes are available under both the initiator and the responder mode, but attribute configuration will be rejected if it does not apply for the selected mode.
- Synthetic loss measurement should not be configured over VLAN untagged ports in the case of a regular VLAN.
- When COS 8 is used on an initiator and responder, a cos value is randomly chosen between 0-7 before transmission of an Synthetic loss measurement (SLM) packet. On the responder side, all SLM packets for the target MEP are accounted for session 8 by ignoring the COS. Similar handling is present for Synthetic Loss Reply (SLR) processing. SLR packet uses the same cos which was present in the incoming SLM packet.
- When synthetic loss measurement is configured over VPLS untagged end-point, only cos 8 can be used.
- The initiator and responder for a particular SLM session should have the same cos configured on both ends.
- Other than an immediate case, the start and stop configuration will always be a part of the running configuration. It is persistent after a reload.
- The stop now command stops any running session. It cancels the start of any scheduled session. In addition, it also resets the start/stop time to "00:00:00" and type to "Immediate" for non-periodic sessions.
- Session configuration cannot be changed when it is running.

- Before configuring any SLM session, ensure the device is configured with the correct date and time. Use the show clock command to verify. Otherwise bring the clock to present time with the set clock hh:mm:ss mm-dd-yy command.
- Synthetic loss measurement (SLM) and Synthetic Loss Reply (SLR) packets are not transmitted or received over blocked ports.

Commands

The following commands are described for the initiator and the responder.

```
device(config-cfm)#loss-measurement slm initiator
```

Syntax: loss-measurement slm initiator | **responder** | **clear-stat**

Initiator - is used configure synthetic loss measurement parameters on Tx side.

Responder - is used to configure synthetic loss measurement parameters on Responder side.

Clear-stat - is used to clear the history logs globally.

```
device(config-cfm)#loss-measurement slm initiator 1
```

Syntax: loss-measurement slm initiator *session_index*

Session_index - is used to configure the session index in range. The acceptable range is 1 - 1000.

```
device(config-cfm-loss-measurement-slm-initiator-1)# domain
```

Syntax: domain

Domain - is used to configure the domain name

```
device(config-cfm-loss-measurement-slm-initiator-1)# cos 1
```

Syntax: cos *cos*

Cos- is used to configure the priority value. The acceptable range is 1 - 8. The default is 7.

```
device(config-cfm-loss-measurement-slm-initiator-1)# tx-interval 1
```

Syntax: tx-interval *interval*

Interval - is used to configure the Tx interval between SLM packets (default - 1sec).

```
device(config-cfm-loss-measurement-slm-initiator-1)# tx-interval 1
```

Syntax: measurement-interval *interval*

Interval - is used to configure SLM Measurement interval (default- 15min).

```
device(config-cfm-loss-measurement-slm-initiator-1)# threshold forward
```

Syntax: threshold [**forward** | **backward**] [**average** | **maximum**] *value*

Default values:

Threshold Forward Average 0xFFFFFFFF mili-percent

Threshold Backward Average 0xFFFFFFFF mili-percent

Threshold Forward maximum 0xFFFFFFFF mili-percent

Threshold Backward maximum 0xFFFFFFFF mili-percent

Configuration examples

Sample configuration of synthetic loss measurement over VLAN

VLAN Configurations:

DUT1 Configuration

```
device(config)# vlan 2
device(config-vlan-2)# tagged ethernet 1/1
```

DUT2 Configuration

```
device(config)# vlan 2
device(config-vlan-2)# tagged ethernet 1/1
```

CFM Configurations:

DUT1 Configuration

```
device(config)# cfm-enable
device(config-cfm)# domain md1 level 7
device(config-cfm-md-md1)# ma ma1 vlan 10 priority 4
device(config-cfm-md-md1-ma-ma1)# mep 3 down port ethernet 1/1
```

DUT2 Configuration

```
device(config)# cfm-enable
device(config-cfm)# domain md1 level 7
device(config-cfm-md-md1)# ma ma1 vlan 2 priority 4
device(config-cfm-md-md1-ma-ma1)# mep 4 down port ethernet 1/1
```

SLM Configurations:

DUT1 Configuration

```
device(config)# cfm
device(config-cfm)# loss-measurement slm initiator 1
device(config-cfm-loss-measurement-slm-initiator-1)#domain md1 ma ma1 src-mep 3 target-mep 4
device(config-cfm-loss-measurement-slm-initiator-1)#cos 2
device(config-cfm-loss-measurement-slm-initiator-1)#tx-interval 1
device(config-cfm-loss-measurement-slm-initiator-1)#measurement-interval 1
```

DUT2 Configuration

```
device(config)# cfm
device(config-cfm)# loss-measurement slm responder 1
device(config-cfm-loss-measurement-slm-responder-1)#domain md1 ma ma1 src-mep 4 target-mep 3
device(config-cfm-loss-measurement-slm-responder-1)#cos 2
```

Starting synthetic loss measurement sessions:

NOTE

Start the synthetic loss measurement (SLM) session on the responder side before the initiator.

DUT2 Configuration (Responder)

```
device(config)# cfm
device(config-cfm)# cfm loss-measurement slm responder 1
device(config-cfm-loss-measurement-slm-responder-1)#start now
```

DUT1 Configuration (Initiator)

```
device(config)# cfm
device(config-cfm)# cfm loss-measurement slm initiator 1
device(config-cfm-loss-measurement-slm-initiator-1)#start now
```

Stopping synthetic loss measurement sessions:

> NOTE
> Stop the initiator before stopping the responder.

DUT1 Configuration

```
device(config)# cfm
device(config-cfm)# cfm loss-measurement slm initiator 1
device(config-cfm-loss-measurement-slm-initiator-1)#stop now
```

DUT2 Configuration

```
device(config)# cfm
device(config-cfm)# cfm loss-measurement slm responder 1
device(config-cfm-loss-measurement-slm-responder-1)#stop now
```

Clearing loss statistics:

You can clear history statistics on the initiator side at any point of time using the following command.

```
device(config)# cfm
device(config-cfm)# cfm loss-measurement slm initiator 1
device(config-cfm-loss-measurement-slm-initiator-1)#clear-stat
```

You can clear the history statistics globally using the following command.

> NOTE
> When this command is executed, history logs will be cleared for all sessions in the system.

```
device(config)# cfm
device(config-cfm)# cfm loss-measurement slm clear-stat
```

Sample configuration - synthetic loss measurement over VPLS

VPLS Configurations:

LER1 Configuration

```
device(config)# router mpls
device(config-mpls)# vpls vpls100 100
device(config-mpls-vpls-vpls100)# vlan 10
device(config-mpls-vpls-vpls100-vlan-10)# tagged ethernet 1/1
device(config-mpls-vpls-vpls100-vlan-10)# end
```

LER2 Configuration

```
device(config)# router mpls
device(config-mpls)# vpls vpls100 100
device(config-mpls-vpls-vpls100)# vlan 10
device(config-mpls-vpls-vpls100-vlan-10)# tagged ethernet 1/1
device(config-mpls-vpls-vpls100-vlan-10)# end
```

CFM Configurations:

```
device Configuration
device(config)# cfm-enable
device(config-cfm)# domain md1 level 7
device(config-cfm-md-md1)# ma ma1 vpls 100 priority 4
device(config-cfm-md-md1-ma-ma1)# mep 3 up vlan 10 port ethernet 1/1
```

LER2 Configuration

```
device(config)# cfm-enable
device(config-cfm)# domain md1 level 7
device(config-cfm-md-md1)# ma ma1 vpls 100 priority 4
device(config-cfm-md-md1-ma-ma1)# mep 4 up vlan 10 port ethernet 1/1
```

Synthetic loss measurement configurations:

LER1 Configuration

```
device(config)# cfm
device(config-cfm)# loss-measurement slm initiator 1
device(config-cfm-loss-measurement-slm-initiator-1)#domain md1 ma ma1 src-mep 3 target-mep 4
device(config-cfm-loss-measurement-slm-initiator-1)#cos 2
device(config-cfm-loss-measurement-slm-initiator-1)#tx-interval 1
device(config-cfm-loss-measurement-slm-initiator-1)#measurement-interval 1
```

LER2 Configuration

```
device(config)# cfm
device(config-cfm)# loss-measurement slm responder 1
device(config-cfm-loss-measurement-slm-responder-1)#domain md1 ma ma1 src-mep 4 target-mep 3
device(config-cfm-loss-measurement-slm-responder-1)#cos 2
```

> NOTE
> Start the synthetic loss measurement session on the responder side before the initiator.

Starting synthetic loss measurement sessions:

LER2 Configuration (Responder)

```
device(config)# cfm
device(config-cfm)# cfm loss-measurement slm responder 1
device(config-cfm-loss-measurement-slm-responder-1)#start now
```

LER1 Configuration (Initiator)

```
device(config)# cfm
device(config-cfm)# cfm loss-measurement slm initiator 1
device(config-cfm-loss-measurement-slm-initiator-1)#start now
```

Stopping synthetic loss measurement sessions:

> NOTE
> Stop the initiator before stopping the responder.

LER1 Configuration

```
device(config)# cfm
device(config-cfm)# cfm loss-measurement slm initiator 1
device(config-cfm-loss-measurement-slm-initiator-1)#stop now
```

LER2 Configuration

```
device(config)# cfm
device(config-cfm)# cfm loss-measurement slm responder 1
device(config-cfm-loss-measurement-slm-responder-1)#stop now
```

Clearing loss statistics:

You can clear history statistics on the initiator side at any time using the following command.

```
device(config)# cfm
device(config-cfm)# cfm loss-measurement slm initiator 1
device(config-cfm-loss-measurement-slm-initiator-1)#clear-stat
```

You can clear the history statistics globally using the following command.

NOTE

When this command is executed, history logs will be cleared for all sessions in the system.

```
device(config)# cfm
device(config-cfm)# cfm loss-measurement slm clear-stat
```

Sample configuration - Synthetic loss measurement over VLL

VLL Configurations:

LER1 Configuration

```
device(config)# router mpls
device(config-mpls)# vll vll100 100
device(config-mpls-vll-vll100)# vlan 10
device(config-mpls-vll-vll100-vlan-10)# tagged ethernet 1/1
device(config-mpls-vll-vll100-vlan-10)# end
```

LER2 Configuration

```
device(config)# router mpls
device(config-mpls)# vll vll100 100
device(config-mpls-vll-vll100)# vlan 10
device(config-mpls-vll-vll100-vlan-10)# tagged ethernet 1/1
device(config-mpls-vll-vll100-vlan-10)# end
```

CFM Configurations:

LER1 Configuration

```
device(config)# cfm-enable
device(config-cfm)# domain md1 level 7
device(config-cfm-md-md1)# ma ma1 vll 100 priority 4
device(config-cfm-md-md1-ma-ma1)# mep 3 up vlan 10 port ethernet 1/1
```

LER2 Configuration

```
device(config)# cfm-enable
device(config-cfm)# domain md1 level 7
device(config-cfm-md-md1)# ma ma1 vll 100 priority 4
device(config-cfm-md-md1-ma-ma1)# mep 4 up vlan 10 port ethernet 1/1
```

Synthetic loss measurement configurations:

LER1 Configuration

```
device(config)# cfm
device(config-cfm)# loss-measurement slm initiator 1
device(config-cfm-loss-measurement-slm-initiator-1)#domain md1 ma ma1 src-mep 3 target-mep 4
device(config-cfm-loss-measurement-slm-initiator-1)#cos 2
device(config-cfm-loss-measurement-slm-initiator-1)#tx-interval 1
device(config-cfm-loss-measurement-slm-initiator-1)#measurement-interval 1
```

LER2 Configuration

```
device(config)# cfm
device(config-cfm)# loss-measurement slm responder 1
device(config-cfm-loss-measurement-slm-responder-1)#domain md1 ma ma1 src-mep 4 target-mep 3
device(config-cfm-loss-measurement-slm-responder-1)#cos 2
```

Starting synthetic loss measurement sessions:

NOTE

Start the synthetic loss measurement (SLM) session on the responder side before the initiator.

LER2 Configuration (Responder)

```
device(config)# cfm
device(config-cfm)# cfm loss-measurement slm responder 1
device(config-cfm-loss-measurement-slm-responder-1)#start now
```

LER1 Configuration (Initiator)

```
device(config)# cfm
device(config-cfm)# cfm loss-measurement slm initiator 1
device(config-cfm-loss-measurement-slm-initiator-1)#start now
```

Stopping synthetic loss measurement sessions:

NOTE

Stop the initiator before stopping the responder.

LER1 Configuration

```
device(config)# cfm
device(config-cfm)# cfm loss-measurement slm initiator 1
device(config-cfm-loss-measurement-slm-initiator-1)#stop now
```

LER2 Configuration

```
device(config)# cfm
device(config-cfm)# cfm loss-measurement slm responder 1
device(config-cfm-loss-measurement-slm-responder-1)#stop now
```

Clearing loss statistics:

You can clear history statistics on the initiator side at any time using the following command.

```
device(config)# cfm
device(config-cfm)# cfm loss-measurement slm initiator 1
device(config-cfm-loss-measurement-slm-initiator-1)#clear-stat
```

You can clear the history statistics globally using the following command.

NOTE

When this command is executed, history logs will be cleared for all sessions in the system.

```
device(config)# cfm
device(config-cfm)# cfm loss-measurement slm clear-stat
```

Show commands

The show cfm loss-measurement slm **session_index** command is used to display the configuration data for a specified indices.

Syntax: show cfm loss-measurement slm *sessionindex*

```
device# show cfm loss-measurement slm 1
-----------------------------------------------------
SLM Session Index         : 1
Status                    : Stopped
Session Type              : Initiator
Domain                    : d1
MA                        : m1
Source MEP                : 20
Target MEP                : 30
Cos                       : 0
Start time                : 16:29:57
Start time type           : Immediate
Stop time                 : 16:30:02
Stop time type            : Immediate
Tx-interval (in Sec)      : 1
Measurement-Interval      : 1
Forward Average           : 0
Forward Max               : 0
Backward Average          : 0
Backward Max              : 0
-----------------------------------------------------------
```

TABLE 15 show cfm loss-measurement slm output

Row	Definition
SLM Session Index	Session index value
Status	stopped or running
Session Type	initiator or responder
Domain	domain name
MA	ma name
Source MEP	source mep id
Target MEP	target mep id or RMEP
COS	data priority loss in which needs to be monitored
Start time	Configured start time
Start time type	Immediate, relative, fixed, periodic
Stop time	configured stop time
Stop time type	Immediate, relative, fixed, periodic
Tx interval (sec)	transmission interval in sec, only for initiator
Measurement interval	measurement-interval in minutes, only for initiator
Forward Average	configured forward average threshold, for initiator
Forward Max	configured forward maximum threshold, for initiator
Backward Average	configured backward average threshold, for initiator
Backward Max	configured backward maximum threshold, for initiator

Syntax: Show cfm loss-measurement slm statistics *sessionindex*

```
device# show cfm loss-measurement slm statistics 1
HISTORY TABLE :
Flag - S:Suspect
-------------------------------------------------------------------------------------------------
Index Flag    Start      Elapsed    TxFwd      RxFwd      TxBck      RxBck      FLR(ratio) BLR(ratio)
-------------------------------------------------------------------------------------------------
5     -       16:29:56   00:00:05   5          5          5          5          0.00000    0.00000
4     -       16:29:39   00:00:14   14         4          4          4          0.71428    0.00000
3     -       16:29:25   00:00:04   4          0          0          0          1.00000    0.00000
2     -       16:29:17   00:00:03   3          0          0          0          1.00000    0.00000
1     -       16:27:14   00:00:03   3          0          0          0          1.00000    0.00000
=================================================================================================
```

Syntax: show cfm loss-measurement slm statistics detailed *session_index rowindex*

```
device# show cfm loss-measurement slm statistics detailed 1 2
HISTORY TABLE :
Flag - S:Suspect
-------------------------------------------------
Index                         : 2
Flag                          : -
Start                         : 16:29:56
Elapsed                    : 00:00:05
TxFwd                      : 5
RxFwd                     : 5
TxBck                     : 5
RxBck                     : 5
FLR(ratio) Max       : 0.00000
FLR(ratio) Min          : 0.00000
FLR(ratio) Avg          : 0.00000
BLR(ratio) Max       : 0.00000
BLR(ratio) Min          : 0.00000
BLR(ratio) Avg          : 0.00000
-------------------------------------------------
```

Syntax: show cfm loss-measurement slm statistics detailed *session_index*

```
device# show cfm loss-measurement slm statistics detailed 1
HISTORY TABLE :
Flag - S:Suspect
-------------------------------------------------
Index                         : 1
Flag                          : -
Start                         : 16:29:22
Elapsed                    : 00:00:05
TxFwd                      : 5
RxFwd                     : 5
TxBck                     : 5
RxBck                     : 5
FLR(ratio) Max       : 0.00000
FLR(ratio) Min          : 0.00000
FLR(ratio) Avg          : 0.00000
BLR(ratio) Max       : 0.00000
BLR(ratio) Min          : 0.00000
BLR(ratio) Avg          : 0.00000
-------------------------------------------------
Index                         : 2
Flag                          : -
Start                         : 16:29:56
Elapsed                    : 00:00:05
TxFwd                      : 5
RxFwd                     : 5
TxBck                     : 5
RxBck                     : 5
FLR(ratio) Max       : 0.00000
FLR(ratio) Min          : 0.00000
FLR(ratio) Avg          : 0.00000
BLR(ratio) Max       : 0.00000
BLR(ratio) Min          : 0.00000
```

```
BLR(ratio) Avg          : 0.00000
-------------------------------------------------
```

Syslog messages

Syslogs will be raised for the following cases:

When the SLM session started

```
<Syslog>: SLM Session started for Session Index <id>
```

When the SLM session stopped

```
<Syslog>: SLM Session stopped for Session Index <id>
```

When the Average Frame Loss Ratio greater than Threshold Average Frame Loss Ratio for both forward and backward case.

```
<Syslog>: SLM Average FLR <value> greater than Threshold Average FLR <value>.
```

When the Maximum Frame Loss Ratio greater than Threshold Maximum Frame Loss Ratio for both forward and backward case.

```
<Syslog>: SLM Average FLR <value> greater than Threshold Average FLR <value>.
```

Port Mirroring

Mirroring and Monitoring

You can monitor traffic on Extreme device ports by configuring another port to "mirror" the traffic on the ports you want to monitor. By attaching a protocol analyzer to the mirror port, you can observe the traffic on the monitored ports.

Monitoring traffic on a port is a two-step process:

- Enable a port to act as the mirror port. This is the port to which you connect your protocol analyzer.
- Enable monitoring on the ports you want to monitor.

You can monitor input traffic, output traffic, or both.

Any port on a module can operate as a mirror port and you can configure more than one mirror port. You can configure the mirror ports on different modules and you can configure more than one mirror port on the same module.

Configuration guidelines for monitoring traffic

Use the following considerations when configuring mirroring for inbound and outbound traffic:

- Any port can be mirrored and monitored except for the management port.
- Only one inbound mirror port can be configured for any inbound monitor port.
- Only one outbound mirror port can be configured for any outbound monitor port.
- A LAG port can be configured as either an inbound or outbound monitor port.
- A LAG port cannot be configured as either an inbound or an outbound mirror port.
- Both input and output monitoring are supported.
- Monitoring for LAG ports is supported.
- sFlow and monitoring can be enabled concurrently on the same port.
- ACL-based inbound mirroring is supported.
- ACL-based inbound sFlow is not concurrently supported.
- On the CES 2000 Series, there can be at most one port configured as the mirror port per port region (aport region is 24-1GbE ports or 2 10-GbE ports). There is no limit on the number of monitor ports that can be configured per port region.

Assigning a mirror port and monitor ports

To configure ethernet port 3/1 for port mirroring, enter the following command.

```
device(config)# mirror-port ethernet 3/1
```

Syntax: [no] mirror-port ethernet slot/portnum

NOTE
If a port is configured as a mirror port, all traffic sent from that port will retain the encapsulation of the port being monitored and not add the encapsulation of the Egress port.

Enter the slot and port number of the port that will be the mirrored.

```
device(config)# interface ethernet 4/1
device(config-if-4/1)# monitor ethernet 3/1
```

Syntax: [no] monitor ethernet slot/portnum both | input | output

Enter the slot and port number of the port that will serve as the monitor port. This port cannot be the same as the mirror port.

NOTE
A mirror port must be an Ethernet port.

Specify input if the port will monitor incoming traffic, output to monitor outgoing traffic, or both to monitor both types of traffic.

NOTE
In VPLS, when an unknown unicast traffic is handled, it uses the corresponding VLAN Forwarding ID to flood the packets to the VLAN domain which contains both the monitored port as well as the mirroring port. But in VLL, there is no such flood handling mechanism and hence, there is a discrepancy in the output of the **show statistic brief** command in terms of the **Packet Transmit** count on the mirroring port.

Displaying mirror and monitor port configuration

To display the inbound and outbound traffic mirrored to each mirror port, enter the following command at any level of the CLI.

```
device# show monitor config
Monitored Port 3/1
    Input traffic mirrored to: 2/1
    Output traffic mirrored to: 1/1
Monitored Port 4/1
    Input traffic mirrored to: 1/2
    Output traffic mirrored to: 2/1
```

Syntax: show monitor config

To display the actual traffic mirrored to each mirror port, enter the following command at any level of the CLI.

```
device# show monitor actual
Monitored Port 3/1
    Output traffic mirrored to: 1/1
Monitored Port 4/1
    Input traffic mirrored to: 1/2
```

Syntax: show monitor actual

This output displays the output traffic mirrored to mirror port 1/1 from port 3/1 and input traffic mirrored to mirror port 1/2 from port 4/1, which are explicitly configured.

ACL-based inbound mirroring

The Multi-Service IronWare software supports using an ACL to select traffic for mirroring from one port to another. Using this feature, you can monitor traffic in the mirrored port by attaching a protocol analyzer to it.

Considerations when configuring ACL-based inbound mirroring

The following must be considered when configuring ACL-based inbound mirroring:

- Configuring a common destination ACL mirror port for all ports of a PPCR (see below)
- Support with ACL CAM sharing enabled (see below)
- The **mirror** and **copy-sflow** keywords are mutually exclusive on a per-ACL clause basis.
- ACL-based inbound mirroring and port-based inbound mirroring are mutually exclusive on a per-port basis.
- ACL-based mirroring must be configured at the LAG level for individual LAG member ports.
- Configuring ACL-based mirroring at the port level on the primary port of a LAG mirrors all traffic on that LAG to the monitor port.

Configuring a common destination ACL mirror port for all ports of a PPCR

All ports using the same PPCR must have a common destination ACL mirror port when configuring ACL-based inbound mirroring. For Example, where ports 4/1 and 4/2 belong to the same PPCR, the following configuration that configures them with different destination ACL mirror ports will fail and generate an error message as shown.

```
device(config)# interface ethernet 4/1
device(config-if-e10000-4/1)# acl-mirror-port ethernet 6/1
device(config-if-e10000-4/1)# interface ethernet 4/2
device(config-if-e10000-4/2)# acl-mirror-port ethernet 6/2
Error: 4/2 and 4/1 should have the same ACL mirror port
```

Support with ACL CAM sharing enabled

For ACL CAM sharing to function, either one of the following conditions must be true:

- All ports that belong to a PPCR have the **acl-mirror-port** command configured to direct mirrored traffic to the same port.
- None of the ports that belong to the PPCR have the **acl-mirror-port** command configured.

ACL CAM sharing cannot function with the configuration shown in the following example because port 4/1 has ACL port mirroring configured and port 4/2 does not.

```
device(config)# enable-acl-cam-sharing
device(config)# interface ethernet 4/1
device(config-if-e10000-4/1)# ip access-group 101 in
device(config-if-e10000-4/1)# acl-mirror-port ethernet 6/1
device(config-if-e10000-4/1)# interface ethernet 4/2
device(config-if-e10000-4/2)# ip access-group 101 in
```

Configuring ACL-based inbound mirroring

The following sections describe how to configure ACL-based Inbound Mirroring on a Extreme device:

- Creating an ACL with a mirroring clause
- Applying the ACL to an interface
- Specifying a destination mirror port
- Specifying the destination mirror port for physical ports
- Specifying the destination mirror port for a LAG
- Configuring ACL-based mirroring for ACLs bound to virtual interfaces
- Specifying the destination mirror port for IP receive ACLs

Creating an ACL with a mirroring clause

The **mirror** keyword in IPv4, Layer 2 and IPv6 ACL clauses directs traffic that matches the clause criteria to be mirrored to another port. In the following examples, the ACL is used to direct IP traffic to a mirror port.

ACL-based Mirroring Supported for IPv4 ACLs.

```
device(config)# access-list 101 permit ip any any mirror
device(config)# access-list 101 permit ip any any
```

ACL-based Mirroring supported for IPv6 Inbound ACLs.

```
device(config)# ipv6 access-list gem
device(config-ipv6-access-list gem)# permit tcp 2001:DB8::/64 2001:DB8::/64 mirror
device(config-ipv6-access-list gem)# permit udp 1000:1::/64 2000:1::/64 mirror
device(config-ipv6-access-list gem)# permit icmp 1000:1::/64 2000:1::/64 mirror
device(config-ipv6-access-list gem)# permit ipv6 any any
```

ACL-based Mirroring supported for Layer-2 Inbound ACLs.

```
device(config)# access-list 400 permit 0000.0000.0010
ffff.ffff.ffff 0000.0000.0020 ffff.ffff.ffff any mirror
device(config)# access-list 400 permit 0000.0000.0050
ffff.ffff.ffff 0000.0000.0020 ffff.ffff.ffff any mirror
device(config)#access-list 400 permit any any any
```

The **mirror** parameter directs selected traffic to the mirrored port. Traffic can only be selected using the **permit** clause. The mirror parameter is supported on rACLs.

> NOTE
> As with any ACL, the final clause must permit desired traffic to flow: be sure to add an appropriate **permit any any** clause to the end of any ACL intended to mirror (and not filter) traffic. Failure to include the **permit** clause will result in disruption of traffic through any interface to which the ACL is applied.

Applying the ACL to an interface

You must apply the ACL to an interface using the **ip access-group command** as shown in the following.

```
device(config)# interface ethernet 1/1
device(config-if-e10000-1/1)# ip access-group 101 in
```

Specifying the destination mirror port

You can specify physical ports or a LAG to mirror traffic from. The following sections describe how to perform each of these configurations.

Specifying the destination mirror port for physical ports

You must specify a destination port for traffic that has been selected by ACL-based Inbound Mirroring. This configuration is performed at the Interface Configuration of the port whose traffic you are mirroring. In the following example, ACL mirroring traffic from port 1/1 is mirrored to port 1/3.

```
device(config)# interface ethernet 1/1
device(config-if-e10000-1/1)# acl-mirror-port ethernet 1/3
```

You can also use the ACL-mirroring feature to mirror traffic from multiple ports to a single port using the Multiple Interface Configuration (MIF) mode as shown in the following example.

```
device(config)# interface ethernet 1/1 to 1/2
device(config-mif-e10000-1/1-1/2)# acl-mirror-port ethernet 1/3
```

Syntax: [no] acl-mirror-port ethernet [**slot/port**]

The [slot/port] variable specifies port that ACL-mirror traffic from the configured interface will be mirrored to.

Specifying the destination mirror port for a LAG

You can mirror the traffic that has been selected by ACL-based inbound mirroring from all ports in a LAG by configuring a destination (monitor) port for the LAG at the interface configuration level of the LAG's primary port. Configuring mirroring on the primary port of the LAG causes ACL-selected traffic from all ports in the LAG (including any ports subsequently added to the LAG dynamically on the XMR Series and MLX Series) to be mirrored to the monitor port. For example, in the following configuration all traffic on LAG "mylag" will be mirrored to port 10/4:

```
device(config)# lag mylag static
device(config-lag-mylag)# ports ethernet 10/1 to 10/3
device(config-lag-mylag)# primary-port 10/1
device(config-lag-mylag)# deploy
device(config-lag-mylag)# exit
device(config)# interface ethernet 10/1
device(config-if-e1000-10/1)# acl-mirror-port ethernet 10/4
```

Syntax: [no] acl-mirror-port ethernet *slot/port*

The **ethernet***slot/port* variable specifies the port that ACL-mirror traffic from the LAG will be mirrored to.

The following considerations apply when configuring ACL-based mirroring with LAGs:

- You must configure ACL-mirroring for an individual member port from the LAG configuration level. Attempting to configure ACL-mirroring at the interface level for an individual member port will fail and display the following message.

  ```
  Error: please use  config level to configure ACL based mirroring on  port.
  ```

- If an individual port is configured for ACL-based mirroring, you cannot add it to a LAG. If you want to add it to a LAG, you must remove it from ACL-based mirroring first. Then you can add it to a LAG. It can then be configured for either ACL-based LAG mirroring or for mirroring an individual port within a LAG.

  If you attempt to add a port that is configured for ACL-based mirroring to a LAG, the following message will display.

  ```
  ACL port is configured on port 2/1, please remove it and try again.
   transaction failed:  Config Vetoed
  ```

- When a LAG with ACL-based mirroring configured on it is deleted or not deployed, the ACL-based mirroring configuration is removed from each of the individual ports that made up the LAG, including the primary port.

Configuring ACL-based mirroring for ACLs bound to virtual interfaces

For configurations that have an ACL bound to a virtual interface, you must configure the **acl-mirror-port** command on a port for each PPCR that is a member of the virtual interface. For example, in the following configuration ports 4/1 and 4/2 share the same PPCR while port 4/3 uses another PPCR.

```
device(config)# vlan 10
device(config-vlan-10)# tagged ethernet 4/1 to 4/3
device(config-vlan-10)# router-interface ve 10
device(config)# interface ethernet 4/1
device(config-if-e10000-4/1)# acl-mirror-port ethernet 5/1
device(config)# interface ve 10
device(config-vif-10)# ip address 10.10.10.254/24
device(config-vif-10)# ip access-group 102 in
device(config)# access-list 101 permit ip any any mirror
```

In this configuration, the **acl-mirror-port** command is configured on port 4/1 which is a member of ve 10. Because of this, ACL-based mirroring will apply to VLAN 10 traffic that arrives on ports 4/1 and 4/2. It will not apply to VLAN 10 traffic that arrives on port 4/3 because that port uses a different PPCR than ports 4/1 and 4/2. To make the configuration apply ACL-based mirroring to VLAN 10 traffic arriving on port 4/3, you must add the following command to the configuration.

```
device(config)# interface ethernet 4/3
device(config-if-e10000-4/3)# acl-mirror-port ethernet 5/1
```

If the ve contains LAG ports, configuration of acl-mirror-port command on an individual LAG port will also apply to other LAG ports that are in the same PPCR. For example, in the following configuration the acl-mirror-port command is configured for LAG port 10/2, which is a member of ve.

```
device(config)# lag mylag static
device(config-lag-mylag)# ports ethernet 10/1 to 10/4
device(config-lag-mylag)# primary-port 10/1
device(config-lag-mylag)# deploy
device(config-lag-mylag)# acl-mirror-port ethe-port-monitored 10/2 ethe 11/3
device(config)# vlan 10
device(config-vlan-10)# tagged ethe 10/1 to 10/4
device(config-vlan-10)# router-interface ve 10
```

The ACL-based mirroring will apply to VLAN 10 traffic incoming on ports 10/1 and 10/2 since they are in the same PPCR and are members of a virtual interface. However it will not apply to VLAN 10 traffic incoming on 10/3 and 10/4 since they are in a different PPCR. To apply ACL-based mirroring on VLAN 10 traffic incoming on 10/3 and 10/4, you will have to additionally configure the acl-mirror-port ether-port-monitored 10/3 ethe 11/3 command under the LAG.

Specifying the destination mirror port for IP Receive ACLs

When specifying a destination port for IP Receive ACLs, you must configure the **acl-mirror-port** command on all ports supported by the same PPCR. For example, if you are using mirroring traffic for an rACL on a 4 x 10G interface module and you want to mirror traffic incoming on the first PPCR, you have to configure the **acl-mirror-port** command on both ports 1 and 2. If you want to mirror IP Receive ACL permit traffic incoming on all ports of the module, you have to configure the **acl-mirror-port** command on all ports of the module.

Telemetry Solutions

Telemetry Solutions overview

Telemetry Solutions provides a VLAN matching capability for IPv4 and IPv6 ACLs. Telemetry Solutions also includes new types of Policy-Based Routing (PBR) next-hop (network interface). You can create policies that classify network traffic into different categories based on the extended ACLs and forward each category of traffic differently, based on the configured policy. With Telemetry Solutions, the ACL match can be based on both VLAN ID and the existing Layer 3 or Layer 4 fields.

Telemetry Solutions improves the user experience with options to classify the network traffic (VLAN matching) and providing more choices for PBR forwarding. You can also utilize the **rule-name** field in the route-map to organize and extract information about PBR configurations.

Limitations

The ACL keyword *VLAN* is only intended to be used in PBR. For ACLs that contain the *VLAN* keyword and is used as standalone ACL, the following restrictions apply:

- An ACL that contains the *VLAN* keyword cannot be applied to Virtual Interfaces (VEs).
- The *VLAN* keyword will be ignored and will have no effect if the ACL is:
  - applied to a physical interface or LAG interface
  - applied to a management interface
  - used as an IP receive ACL
  - used in ACL-based rate-limiting
- If the **set interface** command exists in a route-map and the route-map is applied to a interface, it will only permit packets from the configured VLAN unless the command **allow-all-vlan pbr** is also configured on the interface.

Configuration examples

NOTE

Telemetry can also be configured from SNMP. Refer to the *Unified IP MIB Reference guide* for more information.

Configuration example 1

FIGURE 9 Configuration example 1

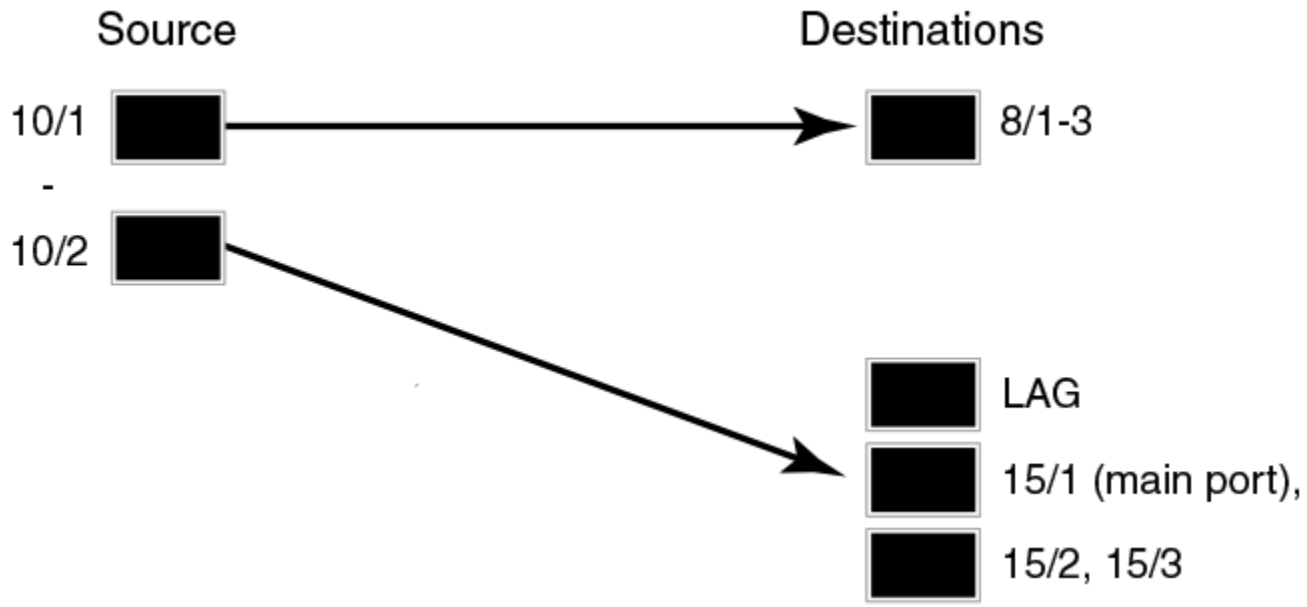

ACL Definition

```
ip access-list extended xGW_Filter1
 permit vlan 114 udp any eq 1066 any
ipv6 access-list xGW_Filter1
 permit vlan 112 ipv6 2001:db8:200::/48 any
ip access-list extended xGW_Filter2
 permit vlan 2405 ip host  any
 permit vlan 3000 ip any any
```

ACL Association and Path Naming

```
route-map xGW_map permit 1
 rule-name xGW_path1
 match ip address xGW_Filter1
 match ipv6 address xGW_Filter1
 set next-hop-flood-vlan 2 preserve-vlan
route-map xGW_map permit 2
 rule-name xGW_path2
 match ip address xGW_Filter2
 set interface ethernet 15/1 preserve-vlan
```

Associate Path Policy to ingress

```
interface ethernet 10/1
 ip policy route-map xGW_map
 ipv6 policy route-map xGW_map
 allow-all-vlan pbr
interface ethernet 10/2
 ip policy route-map xGW_map
 allow-all-vlan pbr
```

Egress Port Definition

```
vlan 2
 untag ethernet 8/1 to 8/3
lag iris_view
 ports ethernet 15/1 to 15/3
 primary port 15/1
 deploy
```

Configuration example 2

FIGURE 10 Configuration example 2

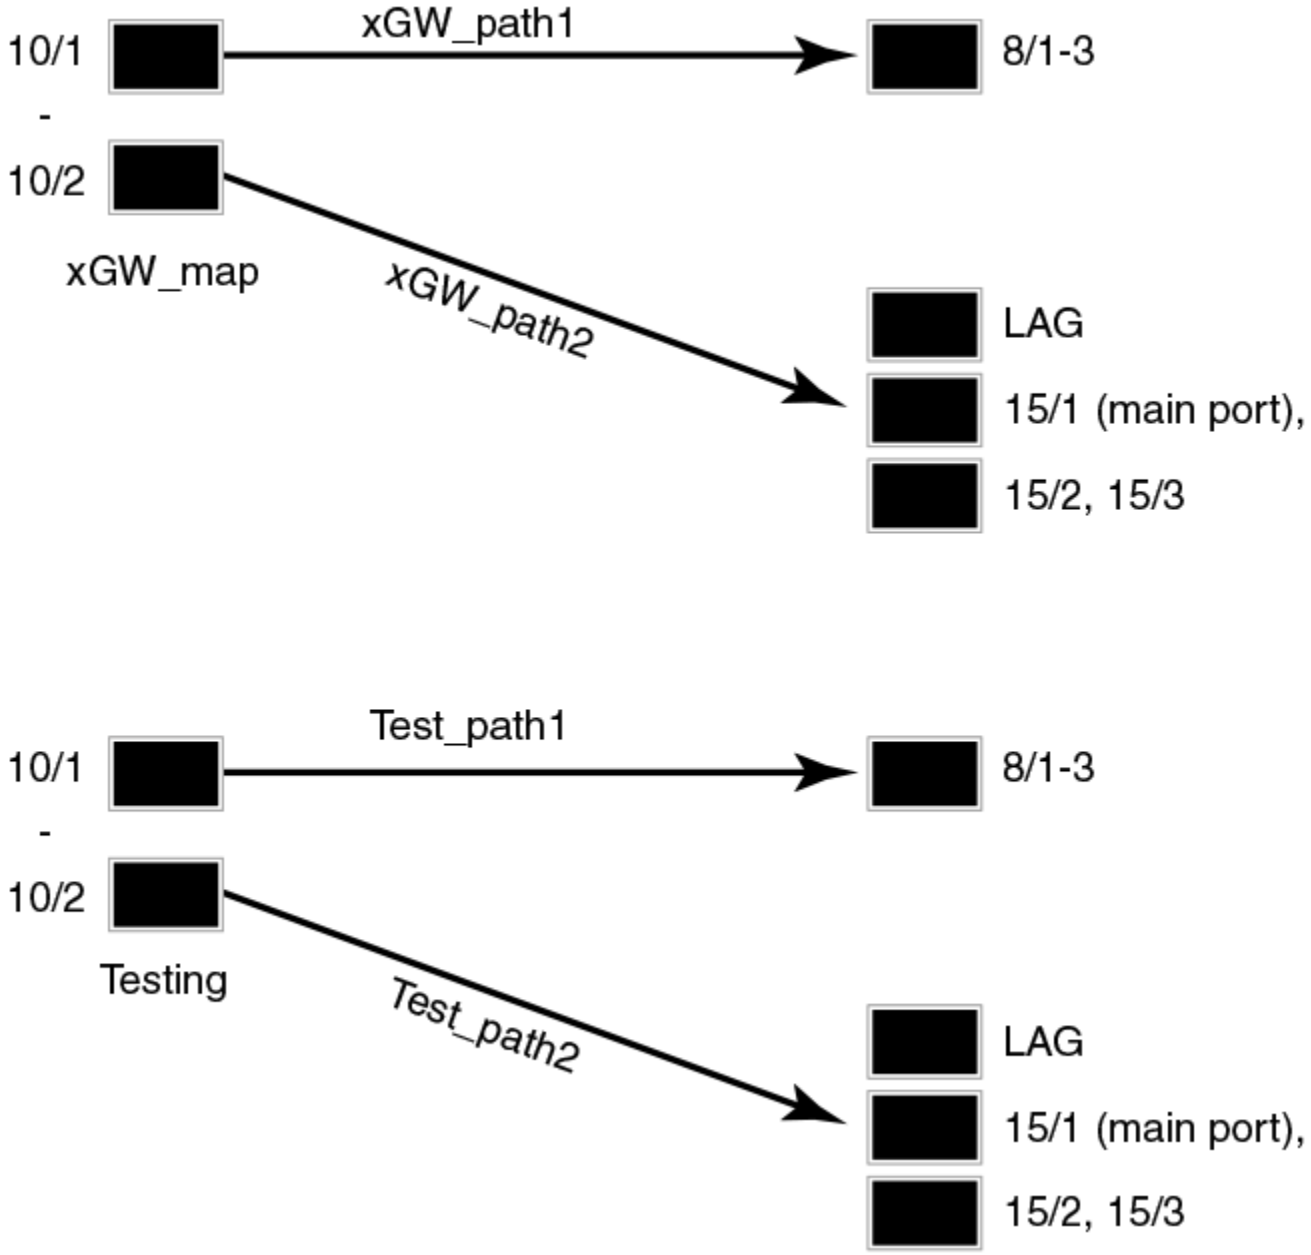

Define Test ACL configurations

```
ip access-list extended Test_filter1
 permit vlan 112 ip host 10.100.50.1 any
 permit vlan 114 udp any eq 2075 any
ip access-list extended Test_filter2
 deny vlan 2405 ip host 10.33.44.55 any
 permit vlan 3000 ip any any
```

Associate Test ACL with Test map/paths

```
route-map Testing permit 1
 rule-name Test_path1
 match ip address Test_filter1
 set next-hop-flood-vlan  2 preserve-vlan
route-map Testing permit 2
 rule-name Test_path2
 match ip address Test_filter2
 set interface ethernet 15/1 preserve-vlan
```

Apply new map to Source ports

```
interface ethernet 10/1
 ip policy route-map Testing
 allow-all-vlan pbr
interface ethernet 10/2
 ip policy route-map Testing
 allow-all-vlan pbr
```

Rebind ACLs

```
 ip rebind-acl all
```

Modify destination ports (if necessary)

```
vlan 2
 untag ethernet 8/1 to 8/3
lag iris_view
 ports ethernet 15/1 to 15/3
 primary-port 15/1
 deploy
```

Apply production map to Source ports

```
interface ethernet 10/1
 ip policy route-map xGW_map
 allow-all-vlan pbr
interface ethernet 10/2
 ip policy route-map xGW_map
 allow-all-vlan pbr
Rebind ACL's
 ip rebind-acl all
```

Modify destination ports (if necessary)

```
vlan 2
 untag ethernet 8/1 to 8/3
lag iris_view
 ports ethernet 15/1 to 15/3
 primary-port 15/1
 deploy
```

Configuration example 3

FIGURE 11 Configuration example 3

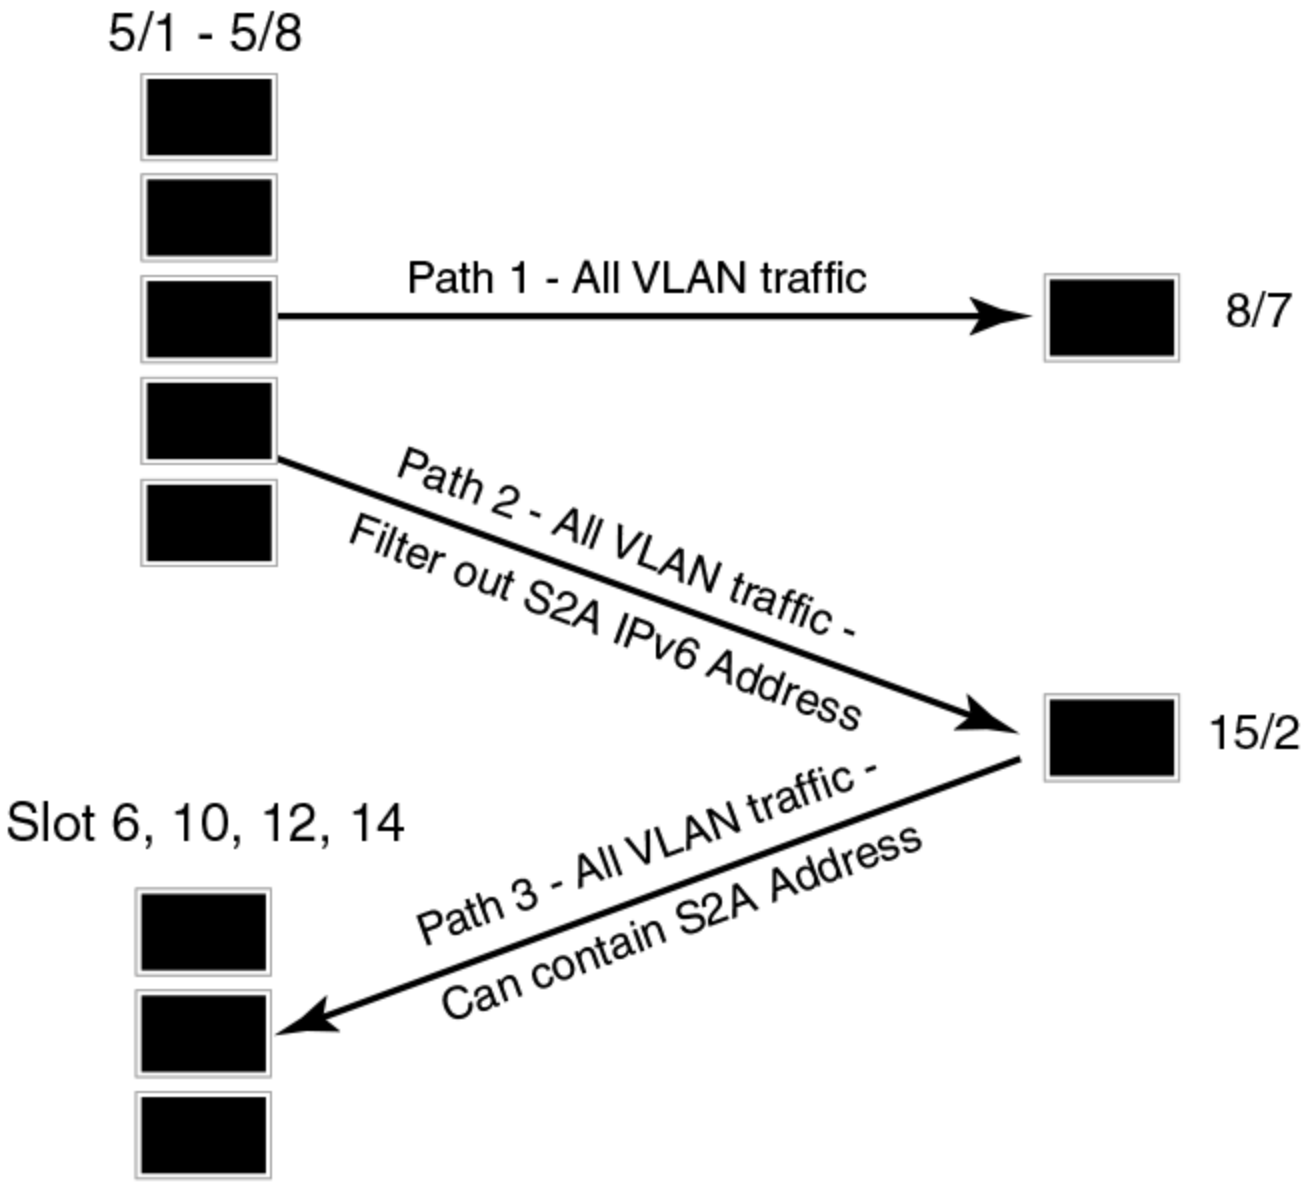

Define ACL configurations

```
ipv6 access-list S2A_traffic
 permit vlan 2011 ipv6 2001:db8:200:1001:194:200::/96 any
 permit vlan 2012 ipv6 2001:db8:200:1001:194:200::/96 any
 permit vlan 2015 ipv6 2001:db8:200:1001:194:200::/96 any
 permit vlan 2016 ipv6 2001:db8:200:1001:194:200::/96 any
 permit vlan 2405 ipv6 2001:db8:200:1001:194:200::/96 any
 permit vlan 2435 ipv6 2001:db8:200:1001:194:200::/96 any
ipv6 access-list Non_S2A_Traffic
 permit ipv6 any any
```

```
ip access-list extended Non_S2A_Traffic
 permit ip any any
```

Associate Traffic ACL with S2A map

```
route-map S2A permit 1
 rule-name S2A_Path
 match ipv6 address S2A_Traffic
 set interface ethernet 8/7 preserve-vlan
route-map S2A permit 2
 rule-name All-Traffic
 match ip address Non_S2A_Traffic
 match ipv6 address Non_S2A_Traffic
 set next-hop-flood-vlan  2 preserve-vlan
```

Apply S2A map to source ports

```
interface ethernet 5/1
 ip policy route-map S2A
 ipv6 policy route-map S2A
 allow-all-vlan pbr
interface ethernet 5/8
 ip policy route-map S2A
 ipv6 policy route-map S2A
 allow-all-vlan pbr
```

Configure destination ports

```
vlan 2
 untag ethernet 8/7 ethernet 15/2
```

With this construct, S2A traffic is explicitly allowed to 8/7 and all other traffic is also sent to 8/7 and 15/2.

Define ACL configurations

```
ipv6 access-list S2A_OtherVLAN
 permit vlan 2007 ipv6 any any
 permit vlan 2008 ipv6 any any
 permit vlan 2009 ipv6 any any
 permit vlan 2010 ipv6 any any
 permit vlan 2017 ipv6 any any
 permit vlan 2019 ipv6 any any
 permit vlan 2009 ipv6 any any
 permit vlan 2010 ipv6 any any
 permit vlan 2017 ipv6 any any
 permit vlan 2019 ipv6 any any
```

NOTE

This would include any S2A IP address packets from these VLANS.

Associate Test ACL with Test map/paths

```
route-map OtherSlot permit 1
 rule-name S2A_OtherVLANPath
 match ipv6 address S2A_OtherVLAN
 set interface ethernet 15/2 preserve-vlan
```

Apply other slot map to source ports on slot 6, 10, 12, 14

```
interface ethernet 6/1
 ipv6 policy route-map OtherSlot
 allow-all-vlan pbr
```

Configuring telemetry solutions

1. Configure IPv4/IPv6 ACLs to match desired traffic.
2. Configure PBR policies to redirect traffic to desired destinations.
3. Apply the PBR policies to interfaces (physical ports, LAG ports or Virtual interfaces).
4. Use the show commands to display information about PBR configurations and operations.

> NOTE
> If both IPv4 and IPv6 traffic need to be subjected to PBR, the IPv4 and IPv6 access lists must be created separately. In addition, both **ip policy route-map xGW_map** and **ipv6 policy route-map xGW_map** must be configured on the interface.

Truncating packets for analysis

Truncating egress packets truncates packets to a specific size across ports before being sent to an analyzer.

Egress packets are truncated to the size required by an analysis tool, while still allowing other tools access to the entire packet data. This allows analyzer tools to process less information or packet data, thereby increasing the capacity available for analysis.

The **egress-truncate** command enables truncation of egress packets across different enabled ports in the Extreme MLX devices.

Truncate egress packets

The **egress-truncate** command truncates egress packets across ports before being sent to an analyzer to increase analyzer processing performance.

The **egress-truncate-size** command must be configured globally to set the packet size.

Once configured globally, the **egress-truncate** command must be enabled on the specific egress ports. Once configured, the status of the **egress-truncate** command can be displayed using the **show interfaces ethernet** command.

Configuring the egress truncate command

The following example sets the size of the truncated egress packets to 64 bytes on slot 2.

```
device(config)#egress-truncate-size 64 slot 2
device(config-if-1/1)#egress-truncate
```

The following are some examples of alternate configuration.

This example shows the configuration of a truncated size of 200 on all slots.

```
device(config)#egress-truncate-size 200 slot all
device(config-if-e1000-1/1)#egress-truncate
```

This example shows the configuration of a truncated size of 100 on slot 1 device 2.

```
device(config)#egress-truncate-size 100 slot 1 2
device(config-if-e1000-1/1)#egress-truncate
```

This example shows the configuration of a truncated size of 150 to all PPCRs of slot 2.

```
device(config)#egress-truncate-size 150 slot 2
device(config-if-e1000-1/1)#egress-truncate
```

Checking truncate configuration

The configuration of the **egress-truncate** command can be checked using the following show commands.

- Show interface ethernet
- Show egress-truncate
- Show egress-truncate interface

802.1BR and VN-tag header processing

802.1BR and VN-tag header processing can be managed using the features discussed in this section. Once the headers are processed, the traffic can be forwarded to be processed by other analytic tools.

The 802.1BR header stripping feature performs the following on 802.1BR traffic .

- Identify 802.1BR traffic
- Strip 802.1BR tags
- Forward stripped packets to the next processing port for further filtering and forwarding.

The VN-tag header stripping feature performs the following on VN-tag traffic.

- Identify VN-tag traffic
- Strip VN-tags
- Forward stripped packets to the next processing port for further filtering and forwarding.

802.1BR header stripping

The feature enables the system to strip the 802.1BR header from ingress traffic to comply with analytic tools that do not understand the 802.1BR header.

As part of this feature, the system will identify packets with an 802.1BR header, strip the header, and send the packet to next processing port.

> NOTE
> The 802.1BR header stripping feature is available only on the following line-cards.
>
> - BR-MLX-40Gx4
> - BR-MLX-10Gx20
> - BR-MLX-100Gx2-CFP2

The syntax of the **strip-802-1br** command is discussed in the following examples.

Syntax: [no] strip-802-1br all | **slot** *slot-num* | **slot** *slot-num* **device** *device-id*

> NOTE
> To use the 802.1BR header stripping functionality, switch to **config-pkt-encap-proc** command line interface (CLI) mode.

Configuring 802.1BR header stripping on all modules

The **strip-802-1br all** command enables the 802.1BR header stripping feature on all the cards that support this feature.

```
device(sw)# config terminal
device(config)# packet-encap-processing
device(config-pkt-encap-proc)# strip-802-1br all
```

Syntax: [no] strip-802-1br all

Configuring 802.1BR header stripping on a specific module

The **strip-802-1br slot** command enables the 802.1BR header stripping feature on a specific card that supports this feature.

NOTE
Slot-num represents the module ID.

```
device(sw)# config terminal
device(config)# packet-encap-processing
device(config-pkt-encap-proc)# strip-802-1br slot 3
```

Syntax: [no] strip-802-1br slot *slot-num*

The following example enables the 802.1BR header stripping feature on a specific device (identified using a device-ID) of a card, that supports this feature.

NOTE
Slot-num represents the module ID and device-id represents the np-id.

```
device(sw)# config terminal
device(config)# packet-encap-processing
device(config-pkt-encap-proc)# strip-802-1br slot 2 device-id 1
```

Syntax: [no] strip-802-1br slot *slot-num* **device** *device-id*

VN-tag header stripping

As part of VN-tag header stripping feature, the system strips the VN-tag header from the ingress traffic as some analytic tools do not understand the VN-tag header.

NOTE
The VN-tag header stripping feature is available only on the following line-cards.

- BR-MLX-40Gx4
- BR-MLX-10Gx20
- BR-MLX-100Gx2-CFP2

The syntax of the **strip-vn-tag** command is discussed in the following examples.

Syntax: [no] strip-vn-tag all | **slot** *slot-num* | **slot** *slot-num* **device** *device-id*

NOTE
To use the VN-tag header stripping functionality, switch to **config-pkt-encap-proc** command line interface (CLI) mode.

Configuring VN-tag header stripping on all modules

The **strip-vn-tag all** command enables the VN-tag header stripping feature on all the cards that support this feature.

```
device(sw)# config terminal
device(config)# packet-encap-processing
device(config-pkt-encap-proc)# strip-vn-tag all
```

Configuring VN-tag header stripping on a specific module

The **strip-vn-tag all** command enables the VN-tag header stripping feature on a specific card that supports this feature.

NOTE
Slot-num represents the module ID.

```
device(sw)# config terminal
device(config)# packet-encap-processing
device(config-pkt-encap-proc)# strip-vn-tag slot 3
```

Syntax: [no] strip-vn-tag slot *slot-num*

The following example enables the VN-tag header stripping feature on a specific device (identified using a device-ID) of a card, that supports this feature.

NOTE
Slot-num represents the module ID and device-id represents the np-id.

```
device(sw)# config terminal
device(config)# packet-encap-processing
device(config-pkt-encap-proc)# strip-vn-tag slot 2 device-id 1
```

Syntax: [no] strip-vn-tag slot *slot-num* **device** *device-id*

Show packet encap processing commands

The status and configuration of 802.1BR and VN-tag header processing are available with the show packet-encap-processing commands.

show packet-encap-processing

The **show packet-encap-processing** command displays the status of the 802.1BR and VN-tag header processing features.

Syntax

show packet-encap-processing

Modes

Exec mode.

Command Output

The **show packet-encap-processing** command displays the following information:

Output field	Description
Slot ID	Slot ID
Dev ID	Device ID
802.1BR Strip	On or Off
802.1BR Bypass	Not supported
VN-tag Strip	On or Off
VN-tag Bypass	Not supported
NVGRE Strip	On or Off

NOTE

802.1BR Bypass and VN-Tag Bypass are not supported in NetIron R06.0.00a.

Examples

The following is an example of **show packet-encap-processing** command output.

```
device(config)# show packet-encap-processing
ON      : Feature is configured
 -       : Feature is not configured
 *       : Feature is not supported
 <Blank> : Slot is Empty
-----------------------------------------------------------
| Slot| Dev| 802.1BR| 802.1BR| VN-Tag | VN-Tag | NVGRE  |
| Id  | Id | Strip  | Bypass | Strip  | Bypass | Strip  |
-----------------------------------------------------------
| S1  | 1  | -      | -      | -      | -      | -      |
|     | 2  | -      | -      | ON     | -      | -      |
-----------------------------------------------------------
| S2  | 1  | *      | *      | *      | *      | *      |
|     | 2  | *      | *      | *      | *      | *      |
-----------------------------------------------------------
| S3  | 1  | *      | *      | *      | *      | *      |
|     | 2  | *      | *      | *      | *      | *      |
-----------------------------------------------------------
| S4  | 1  | -      | -      | -      | -      | -      |
|     | 2  | ON     | -      | -      | -      | -      |
-----------------------------------------------------------
```

History

Release version	Command history
6.0.00a	This command was introduced.

show packet-encap-processing strip-802-1BR

The **show packet-encap-processing strip-802-1BR** command displays the status of the 802.1BR header processing features.

Syntax

show packet-encap-processing strip-802-1BR

Modes

EXEC mode.

Command Output

The **show packet-encap-processing strip-802-1BR** command displays the following information:

Output field	Description
Slot	Slot number
Dev	Device number
VN-Tag Strip	On, - , or *

Examples

The following is an example of **show packet-encap-processing strip-802-1BR** command output.

```
device(config)# show packet-encap-processing strip-802-1BR
ON      : Feature is configured
 -       : Feature is not configured
 *       : Feature is not supported
 <Blank> : Slot is Empty
----------------------------
| Slot| Dev| 802.1BR Strip |
----------------------------
| S1  | 1  | -             |
|     | 2  | -             |
----------------------------
| S2  | 1  | *             |
|     | 2  | *             |
----------------------------
| S3  | 1  | *             |
|     | 2  | *             |
----------------------------
| S4  | 1  | -             |
|     | 2  | ON            |
----------------------------
```

History

Release version	Command history
6.0.00a	This command was introduced.

show packet-encap-processing strip-vn-tag

The **show packet-encap-processing strip-vn-tag** command displays the status of the strip-vn-tag header processing features.

Syntax

show packet-encap-processing strip-vn-tag

Modes

EXEC mode.

Command Output

The **show packet-encap-processing strip-vn-tag** command displays the following information:

Output field	Description
Slot	Slot number
Dev	Device number
VN-Tag Strip	On, - , or *

Examples

The following is an example of **show packet-encap-processing strip-vn-tag** command output.

```
device(config)# show packet-encap-processing strip-vn-tag
ON      : Feature is configured
 -       : Feature is not configured
 *       : Feature is not supported
 <Blank> : Slot is Empty
----------------------------
| Slot| Dev| VN-Tag Strip  |
----------------------------
| S1  | 1  | -             |
|     | 2  | ON            |
----------------------------
| S2  | 1  | *             |
|     | 2  | *             |
----------------------------
| S3  | 1  | *             |
|     | 2  | *             |
----------------------------
| S4  | 1  | -             |
|     | 2  | -             |
----------------------------
```

History

Release version	Command history
6.0.00a	This command was introduced.

show packet-encap-processing slot

The **show packet-encap-processing slot** command displays the status of the 802.1BR and VN-tag header processing features on a elected slot.

Syntax

show packet-encap-processing { **slot** *slot-num* }

Parameters

slot

Identifies the slot to be displayed.

slot-num

The slot number.

Modes

EXEC mode

Command Output

The **show packet-encap-processing slot** command displays the following information:

Output field	Description
Slot ID	Slot ID
Dev ID	Device ID
802.1BR Strip	On or Off
802.1BR Bypass	Not supported
VN-tag Strip	On or Off
VN-tag Bypass	Not supported
NVGRE Strip	On or Off

NOTE

802.1BR Bypass and VN-Tag Bypass are not supported in NetIron R06.0.00a.

Examples

The **show packet-encap-processing slot** command displays the following information:

```
device(config)# show packet-encap-processing slot 1
ON      : Feature is configured
 -       : Feature is not configured
 *       : Feature is not supported
 <Blank> : Slot is Empty
-----------------------------------------------------------
| Slot| Dev| 802.1BR| 802.1BR| VN-Tag | VN-Tag | NVGRE  |
| Id  | Id | Strip  | Bypass | Strip  | Bypass | Strip  |
-----------------------------------------------------------
| S1  | 1  | -      | -      | -      | -      | -      |
|     | 2  | -      | -      | ON     | -      | -      |
-----------------------------------------------------------
```

History

Release version	Command history
6.0.00a	This command was introduced.

show packet-encap-processing interface ethernet

The **show packet-encap-processing interface ethernet** command displays the status of the 802.1BR and VN-tag header processing features.

Syntax

show packet-encap-processing interface ethernet

Modes

EXEC mode.

Command Output

The **show packet-encap-processing interface ethernet** command displays the following information:

Output field	Description
Port State	Enabled or Disabled
Feature-Name	On or Off
Status	On or OFF

NOTE
802.1BR Preservation and VN-Tag Preservation are not supported in NetIron R06.0.00a.

Examples

The following is an example of **show packet-encap-processing interface ethernet** command output.

```
device(config)# show packet-encap-processing interface ethernet 1/1
------------------------------------------------------------
  Port State :                  Disabled
------------------------------------------------------------
  Feature-Name                  Status
  802.1BR Stripping             ON
  802.1BR Preservation          OFF
  VN-tag Stripping              OFF
  VN-tag Preservation           OFF
  NVGRE Stripping               OFF
------------------------------------------------------------
```

History

Release version	Command history
6.0.00a	This command was introduced.

show running-config

The show running-config output discussed in this section is unique to config-pkt-encap-proc mode.

Syntax

show running-config

Modes

config-pkt-encap-proc mode

Examples

The following is an example output when the functionality has been enabled on all ppcrs have been configured.

```
device(config-pkt-encap-proc)# strip-802-1br all
device(config-pkt-encap-proc)# show running-config
packet-enap-processing
    strip-802-1br all
```

The following is an example output when the functionality has been enabled on all ppcrs have been configured, and then disabled on a selective ppcr.

```
device(config-pkt-encap-proc)# strip-802-1br all
device(config-pkt-encap-proc)# no strip-802-1br slot 4 device-id 1
device(config-pkt-encap-proc)# show running-config
packet-enap-processing
      strip-802-1br slot 1
   strip-802-1br slot 2
   strip-802-1br slot 3
   strip-802-1br slot 4 device-id 2
```

History

Release version	Command history
6.0.00a	This command was introduced.

IP payload length based filtering using ACL

Using this feature a range of IP payload length can be configured to be used for filtering of traffic with ACL.

The IP payload length is the size of the data portion of the IP datagram. The IP payload length range can be configured for port per packet processor (PPCR) filtering. This range of IP payload length then can be used as a filter parameter with the Access control List (ACL). This feature is supported for both IP and IPv6 traffic. The IP payload length based filtering using ACL feature allows a user to filter ingress IP/IPv6 traffic based on IP payload length of packets.

- IP payload length is the size of data carried in IP packets.
- In IPv4 packets payload length is the total length excluding the IP header length.
- In IPv6 packets payload length is present in the IPv6 header.
- The range of IP payload length can be configured for both versions.
- IPv4 extended and IPv6 ACL filters can be configured with the match-payload- len clause.

- IP packets having a payload length inside the configured range will be filtered with using the ACL.
- IPv4 packets with option header and IPoMPLS transit traffic are not supported with feature.

NOTE
IPv4 packets with option header and IPoMPLS transit traffic are not supported with feature.

NOTE
In a case where the IP payload length range is not configured all the filters with match-payload-len will be ignored and packets are not matched.

Configuration steps

The IP payload length range can be updated on either the global, slot or PPCR level.

- Enable and configure the IP payload length for PPCR using the **ip match-payload-len** command.
- Enable the IP payload length check attribute in ACL filter.
- Apply ACL on the interface.

Enabling and configuring the IP payload length for PPCR

The ACL filter creation command supports the match payload length attribute.

```
device(config)# access-list 111 permit ip any any match-payload-len
```

Globally configure the IPv4 payload length

This command sets the IP payload length range [700, 1000] in each PPCR of all slots.

```
device(config)#ip match-payload-len slot all range 700 1000
```

Configure the IPv4 payload length on all PPCR on a given slot

This command will set the IP payload length equal to 800 for all PPCR in slot 2.

```
device(config)#ip match-payload-len slot 2 range 800 800
```

Configure the IPv4 payload length on a selected PPCR

This command will set the IP payload length less than or equal to 1000 for PPCR 1 of slot 1.

```
device(config)#ip match-payload-len slot 1 ppcr 1 range 0 1000
```

Removing IPv4 payload length configuration from a selected PPCR

This command will remove the IP payload length configuration from PPCR 1 of slot 2. When using the remove command the range attribute and values are not required.

```
device(config)#no ip match-payload-len slot 2 ppcr 1
```

Enabling and configuring the IPv6 payload length for PPCR

The ACL filter creation command supports the match payload length attribute.

```
device(config-ipv6-access-list payload)#permit ipv6 any any match-payload-len
```

Globally configure the IPv6 payload length

This command sets the IPv6 payload length range [700, 1000] in each PPCR of all slots.

```
device(config)#ipv6 match-payload-len slot all range 700 1000
```

Configure the IPv6 payload length on all PPCR on a given slot

This command will set the IPv6 payload length equal to 800 for all PPCR in slot 2.

```
device(config)#ipv6 match-payload-len slot 2 range 800 800
```

Configure the IPv6 payload length on a selected PPCR

This command will set the IPv6 payload length less than or equal to 1000 for PPCR 1 of slot 1.

```
device(config)#ipv6 match-payload-len slot 1 ppcr 1 range 0 1000
```

Removing IPv6 payload length configuration from a selected PPCR

This command will remove the IPv6 payload length configuration from PPCR 1 of slot 2. When using the remove command the range attribute and values are not required.

```
device(config)#no ipv6 match-payload-len slot 2 ppcr 1
```

show ip match-payload-len

This show command displays the configuration for all PPCRs on which IP payload length range is configured.

Syntax

show ip match-payload-len interface ethernet *slot / port*

Modes

EXEC mode.

Command Output

The **show ip match-payload-len** command displays the following information:

Output field	Description
Slot	Slot number
PPCR	PPCR number
Min-Payload-length	Minimum configured payload length
Max-Payload-length	Maximum configured payload length

Examples

The following is an example of the show command for the IP payload length system configuration.

```
device(config)#show ip match-payload-len
IP Match Payload Length Configuration
Slot        PPCR       Min-Payload-length        Max-Payload-length
1             1            0                         1000
1             2            700                       1000
2             2            800                       800
3             1            700                       1000
3             2            700                       1000
```

The following is an example of the show command for the IP payload length configuration on a specific interface.

```
device(config)#show ip match-payload-len interface ethernet 1/5
IP Match Payload Length Configuration
Slot        PPCR               Min-Payload-length               Max-Payload-length
1             2                        0                                        1000
```

History

Release version	Command history
6.0.00a	This command was introduced.

show ip match-payload-len interface ethernet

This show command displays the IP payload length configuration on a specific interface.

Syntax

show ip match-payload-len { **interface ethernet** *slot* | *port* }

Parameters

interface ethernet

Indicates a specific interface output to be displayed.

slot

The specific slot of the interface.

port

The specific port of the interface.

Modes

EXEC mode

Command Output

The **show ip match-payload-len interface ethernet** command displays the following information:

Output field	Description
Slot	Slot number
PPCR	PPCR number
Min-Payload-length	Minimum configured payload length
Max-Payload-length	Maximum configured payload length

Examples

The following is an example of the show command for the IP payload length configuration on a specific interface.

```
device(config)#show ip match-payload-len interface ethernet 1/5
IP Match Payload Length Configuration
Slot        PPCR          Min-Payload-length        Max-Payload-length
1             2               0                         1000
```

History

Release version	Command history
6.0.00a	This command was introduced.

show ipv6 match-payload-len

This show command displays the configuration for all PPCRs on which IPv6 payload length range is configured.

Syntax

show ipv6 match-payload-len

Modes

EXEC mode

Command Output

The **show ipv6 match-payload-len** command displays the following information:

Output field	Description
Slot	Slot number
PPCR	PPCR number
Min-Payload-length	Minimum configured payload length
Max-Payload-length	Maximum configured payload length

Examples

The following is an example of the show command for the IP payload length system configuration.

```
device(config)#show ipv6 match-payload-len
IPv6 Match Payload Length Configuration
Slot        PPCR         Min-Payload-length        Max-Payload-length
1             1             0                         1000
1             2             700                       1000
2             2             800                       800
3             1             700                       1000
3             2             700                       1000
```

History

Release version	Command history
6.0.00a	This command was introduced.

show ipv6 match-payload-len interface ethernet

This show command displays the IPv6 payload length configuration on a specific interface.

Syntax

show ipv6 match-payload-len { **interface ethernet** *slot* | *port* }

Parameters

interface ethernet

Indicates a specific interface output to be displayed.

slot

The specific slot of the interface.

port

The specific port of the interface.

Modes

EXEC mode

Command Output

The **show ipv6 match-payload-len interface ethernet** command displays the following information:

Output field	Description
Slot	Slot number
PPCR	PPCR number
Min-Payload-length	Minimum configured payload length
Max-Payload-length	Maximum configured payload length

Examples

The following is an example of the show command for the IPv6 payload length configuration on a specific interface.

```
device(config)#show ipv6 match-payload-len interface ethernet 1/5
IPv6 Match Payload Length Configuration
Slot        PPCR          Min-Payload-length         Max-Payload-length
1             2               0                          1000
```

History

Release version	Command history
6.0.00a	This command was introduced.

Remote Network Monitoring

This chapter describes the remote monitoring features available on Extreme products:

- **Remote Monitoring (RMON) statistics** - All Extreme products support RMON statistics on the individual port level. Refer to RMON support on page 176.
- **sFlow** - sFlow collects interface statistics and traffic samples from individual interfaces on a device and exports the information to a monitoring server.

Basic management

The following sections contain procedures for basic system management tasks.

Viewing system information

You can access software and hardware specifics for a device.

To view the software and hardware details for the system, enter the **show version** command.

```
device# show version
```

Syntax: show version

Viewing configuration information

You can view a variety of configuration details and statistics with the show option. The **show** option provides a convenient way to check configuration changes before saving them to flash.

The show options available will vary for the device and by configuration level.

To determine the available show commands for the system or a specific level of the CLI, enter the following command.

```
device# show ?
```

Syntax: show *option*

You also can enter "**show** " at the command prompt, then press the TAB key.

Viewing port statistics

Port statistics are polled by default every 10 seconds.

You can view statistics for ports by entering the following **show** commands:

- show interfaces
- show configuration

Viewing STP statistics

You can view a summary of STP statistics for the device. STP statistics are by default polled every 10 seconds.

To view spanning tree statistics, enter the **show span** command. To view STP statistics for a VLAN, enter the **span vlan** command.

Clearing statistics

You can clear statistics for many parameters with the clear option.

To determine the available **clear** commands for the system, enter the following command.

```
device# clear ?
```

Syntax: clear option

You also can enter "**clear** " at the command prompt, then press the TAB key.

> **NOTE**
> Clear commands are found at the Privileged EXEC level.

RMON support

The RMON agent supports the following groups. The group numbers come from the RMON specification (RFC 1757):

- Statistics (RMON Group 1)
- History (RMON Group 2)
- Alarms (RMON Group 3)
- Events (RMON Group 9)

The CLI allows you to make configuration changes to the control data for these groups, but you need a separate RMON application to view and display the data graphically.

Statistics (RMON group 1)

Count information on multicast and broadcast packets, total packets sent, undersized and oversized packets, CRC alignment errors, jabbers, collision, fragments and dropped events is collected for each port on a device.

No configuration is required to activate collection of statistics for the device. This activity is by default automatically activated at system start-up.

> **NOTE**
> The NetIron system provides limited MIB counters. Extreme uses "rmon_giant" to represent oversized packet, i.e 9216 and above.

You can view a textual summary of the statistics for all ports by entering the following CLI command.

```
device(config)# show rmon statistics
Ethernet statistics 1 is active, owned by monitor
 Interface 1/1 (ifIndex 1) counters
                   Octets          0
              Drop events          0                  Packets          0
           Broadcast pkts          0            Multicast pkts          0
     CRC alignment errors          0            Undersize pkts          0
            Oversize pkts          0                 Fragments          0
                  Jabbers          0                Collisions          0
```

```
          64 octets pkts          0     65 to 127 octets pkts          0
   128 to 255 octets pkts          0    256 to 511 octets pkts          0
   512 to 1023 octets pkts         0   1024 to 1518 octets pkts         0
```

Syntax: show rmon statistics [*num* | **ethernet** *slot/port* | **management***num* | | **begin** *expression* | **exclude** *expression* | **include** *expression*]

The *portnum* parameter specifies the port number. You can use the physical port number or the SNMP port number. The physical port number is based on the product.

- The ports are numbered according to slot and port. For example, the first port in slot 1 is 1/1. The third port in slot 7 is 7/3.

The SNMP numbers of the ports start at 1 and increase sequentially. For example, if you are using a Chassis device and slot 1 contains an 8-port module, the SNMP number of the first port in slot 2 is 9. The physical port number of the same port is 2/1.

This command shows the following information.

TABLE 16 Export configuration and statistics

This line...	Displays...
Octets	The total number of octets of data received on the network. This number includes octets in bad packets. This number does not include framing bits but does include Frame Check Sequence (FCS) octets.
Drop events	Indicates an overrun at the port. The port logic could not receive the traffic at full line rate and had to drop some packets as a result. The counter indicates the total number of events in which packets were dropped by the RMON probe due to lack of resources. This number is not necessarily the number of packets dropped, but is the number of times an overrun condition has been detected.
Packets	The total number of packets received. This number includes bad packets, broadcast packets, and multicast packets.
Broadcast pkts	The total number of good packets received that were directed to the broadcast address. This number does not include multicast packets.
Multicast pkts	The total number of good packets received that were directed to a multicast address. This number does not include packets directed to the broadcast address.
CRC alignment errors	The total number of packets received that were from 64 - 1518 octets long, but had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error). The packet length does not include framing bits but does include FCS octets.
Undersize pkts	The total number of packets received that were less than 64 octets long and were otherwise well formed. This number does not include framing bits but does include FCS octets.
Fragments	The total number of packets received that were less than 64 octets long and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error). It is normal for this counter to increment, since it counts both runts (which are normal occurrences due to collisions) and noise hits. This number does not include framing bits but does include FCS octets.
Oversize packets	The total number of packets received that were longer than 1518 octets and were otherwise well formed. This number does not include framing bits but does include FCS octets.
Jabbers	The total number of packets received that were longer than 1518 octets and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error).

TABLE 16 Export configuration and statistics (continued)

This line...	Displays...
	NOTE This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition where any packet exceeds 20 ms. The allowed range to detect jabber is between 20 ms and 150 ms. This number does not include framing bits but does include FCS octets.
Collisions	The best estimate of the total number of collisions on this Ethernet segment.
64 octets pkts	The total number of packets received that were 64 octets long. This number includes bad packets. This number does not include framing bits but does include FCS octets.
65 to 127 octets pkts	The total number of packets received that were 65 - 127 octets long. This number includes bad packets. This number does not include framing bits but does include FCS octets.
128 to 255 octets pkts	The total number of packets received that were 128 - 255 octets long. This number includes bad packets. This number does not include framing bits but does include FCS octets.
256 to 511 octets pkts	The total number of packets received that were 256 - 511 octets long. This number includes bad packets. This number does not include framing bits but does include FCS octets.
512 to 1023 octets pkts	The total number of packets received that were 512 - 1023 octets long. This number includes bad packets. This number does not include framing bits but does include FCS octets.
1024 to 1518 octets pkts	The total number of packets received that were 1024 - 1518 octets long. This number includes bad packets. This number does not include framing bits but does include FCS octets.

NOTE
The number of entries in a RMON statistics table directly corresponds to the number of ports on a system. For example, if the system is a 26 port device, there will be 26 entries in the statistics display.

History (RMON group 2)

All active ports by default will generate two history control data entries per active device interface. An active port is defined as one with a link up. If the link goes down the two entries are automatically be deleted.

Two history entries are generated for each device:

- a sampling of statistics every 30 seconds
- a sampling of statistics every 30 minutes

The history data can be accessed and displayed using any of the popular RMON applications

A sample RMON history command and its syntax is shown below.

```
device(config)# rmon history 1 interface 1 buckets 10 interval 10 owner nyc02
```

Syntax: rmon history *entry-number* **interface ethernet** *slot/port* | **management** *num* **buckets** *number***interval** *sampling-interval* **owner** *text-string*

You can modify the sampling interval and the bucket (number of entries saved before overwrite) using the CLI. In the above example, owner refers to the RMON station that will request the information.

NOTE
To review the control data entry for each port or interface, enter the **show rmon history** command.

Alarm (RMON group 3)

Alarm is designed to monitor configured thresholds for any SNMP integer, time tick, gauge or counter MIB object. Using the CLI, you can define what MIB objects are monitored, the type of thresholds that are monitored (falling, rising or both), the value of those thresholds, and the sample type (absolute or delta).

An alarm event is reported each time that a threshold is exceeded. The alarm entry also indicates the action (event) to be taken if the threshold be exceeded.

A sample CLI alarm entry and its syntax is shown below.

```
device(config)# rmon alarm 1 ifInOctets.6 10 delta rising-threshold 100 1 falling threshold 50 1 owner nyc02
```

Syntax: rmon alarm *entry-number MIB-object.interface-num sampling-time sample-type threshold- type threshold-value event-number threshold-type threshold-value event-number* **owner** *text-string*

The *sample-type* can be absolute or delta.

The *threshold-type* can be falling-threshold or rising-threshold.

Event (RMON group 9)

The two elements to the Event Group are the event control table and the event log table. The event control table defines the action to be taken when an alarm is reported. Defined events can be found by entering the CLI command, show event. The Event Log Table collects and stores reported events for retrieval by an RMON application.

A sample entry and syntax of the event control table is shown below.

```
device(config)# rmon event 1 description 'testing a longer string' log-and-trap public owner nyc02
```

Syntax: rmon event *event-entry* **description** *text-string* **log** | **trap** | **log-and-trap** | **owner** *rmon-station*

sFlow

sFlow is a system for collecting information about traffic flow patterns and quantities within and among a set of devices. You can configure a device to perform the following tasks:

- Sample packet flows
- Collect packet headers from sampled packets to gather ingress and egress information on these packets
- Compose flow sample messages from the collected information
- Relay messages to an external device known as a collector

Participating devices can also relay byte and packet counter data (counter samples) for ports to the collector.

The port connected to the collector forwards sFlow packets in management VRF and default VRF. The Extreme implementation of sFlow data collection supports AS path information in the following types of sFlow packets:

- Non-default VRF IPv4 sampled packets
- Non-default VRF IPv6 sampled packets
- Default VRF IPv4 sampled packets
- Default VRF IPv6 sampled packets

RFC 3176, "InMon Corporation's sFlow: A Method for Monitoring Traffic in Switched and Routed Networks" describes sFlow. Refer to this RFC to determine the contents of the sampled packet.

Extreme supports sFlow v5, which replaces the version outlined in RFC 3176.

sFlow event workflow

If the sFlow destination is IPv6 and the sFlow Agent IPv6, then an IPv6 agent will be selected from the configured interface. Otherwise, IPv4 will be selected from the configured interface ID from the **sFlow agent** command. If the sFlow agent is not configured, the router ID is used.

If the sFlow destination is IPv4, and the sFlow agent is configured, then an IPv4 agent will be selected from the configured interface. If the sFlow agent is not configured, the router-ID is used.

The Agent IP address selects the first IP address in the interface IP address list. The Agent IPv6 address is unspecified by default. Use the **show sflow** command to verify the interface IP address list.

The status of an IP-port (UP, DOWN) will not impact the sFlow source IP.

The adding or deleting of IP addresses on the interface upon which the sFlow agent interface is configured or a router ID change will trigger the following events:

1. Router ID event:

   If the sFlow agent is not configured, or has been configured but an IP interface does not contain an IP address, then the sFlow agent will use the current management VRF router ID (if any). If the management VRF has changed, then the sFlow agent will also update the agent IP address. However, if the management VRF is disabled or assigned to the default VRF (default behavior), then a router ID event will be applied for the global router ID. The sFlow agent will be updated accordingly.

2. Adding IP address event:

   Adding an IP address on an interface upon which the sFlow agent is configured on will impact an agent-IP based on the following scenarios:

   - If this IP address is the first IP address in the table then the sFlow agent selects it.
   - If the added IP address is positioned on the top of the IP table (due to IP address sequence order), then an agent IP will be reassigned to it. However, if it is not, then it will not impact the agent IP address.

3. Deleting IP address event:

   Deleting an IP address on an interface that the sFlow agent is configured on will impact an agent-IP based on the following scenarios:

   - If the deleted IP address is an equivalent to agent IP address then the next IP address on the same interface will be selected.
   - If no more IP addresses are found on that interface, then the agent IP address will use the router ID as the default behavior. sFlow agent IPv6 will be unspecified. Otherwise there is no action.

Configuration considerations

- Sample data is collected from inbound traffic on ports enabled for sFlow, but it does not collect the outbound traffic, even if the sFlow forwarding is enabled in the egress port. However, byte and packet counters that are sent to the collector include ingress and egress traffic statistics. The actual IP source address of the IP header is taken from the router port address of the best route to the sFlow collector IP address.
- Interface module processors directly forward sFlow packets to the specified sFlow collector. The sFlow collector is reachable by the way of ports on any of the Interface modules. Extreme requires sFlow collector to be connected to non-management port.
- For multicast traffic, sFlow sampling will display incorrect output for egress VLANs. In some configuration scenarios ingress VLAN may be incorrect.
- sflow is implemented in the default VRF only. Therefore, sflow data is only accessible by the sflow collector (sflow destination hosts) defined in the default VRF.

Source address

The sampled sFlow data sent to the collectors includes an agent_address field. This field identifies the IP address of the device that sent the data.

If the sFlow destination is IPv6 and the sFlow Agent IPv6, then an IPv6 agent will be selected from the configured interface. Otherwise, IPv4 will be selected from the configured interface ID from the sFlow agent command. If the sFlow agent is not configured, the router-ID is used.

If the sFlow Destination is IPv4, and the sFlow agent is configured, then an IPv4 agent will be selected from the configured interface. If the sFlow agent is not configured, the router-ID is used.

sFlow looks for an IP address in the following order, and uses the first address found:

- The router ID configured by the **ip router-id** command, in the interface IP address list. The Agent IPv6-address is unspecified by default. Use the **show sflow** command to verify the interface IP address list.
- The first IP address on the lowest-numbered loopback interface
- The first IP address on the lowest-numbered virtual interface
- The first IP address on any interface

NOTE
If an IP address is not already configured when you enable sFlow, the feature uses the source address 0.0.0.0. To display the agent_address, enable sFlow, and then enter the **show sflow** command. Refer to sFlow forwarding on page 189 and Displaying sFlow information on page 192.

NOTE
If you change the agent_address, you must disable and then re-enable sFlow to use the newly configured address.

Sampling rate

The sampling rate is the average ratio of the number of packets incoming on an sFlow-enabled port, to the number of flow samples taken from those packets. Device ports send only the sampled traffic to the CPU. sFlow sampling requires high LP CPU usage, which can affect performance in some configurations, especially if a high sampling rate is implemented.

Configured rate and actual rate

When you enter a sampling rate value, this value is the configured rate. The software rounds the value you enter to the next higher power of 2 to obtain the actual rate. This value becomes the actual sampling rate. For example, if the configured sampling rate is 1000, then the actual rate is 1024; and the hardware samples 1 in 1024 packets. If the configured sampling rate is 1025, then the actual rate is 2048.

NOTE
This behavior applies to the XMR Series and MLX Series platforms and does not apply to the CES 2000 Series and CER 2000 Series devices. In CES 2000 Series and CER 2000 Series devices, the system does not apply rounding.

Extended router information

Extended router information contains information for the next hop router. This information includes the next hop router's IP address and the outgoing VLAN ID. Extended router information also includes the source IP address prefix length and the destination IP address prefix length.

The prefix length of IPv4 source and destination IP addresses is collected only if you configure BGP on the devices.

Extended gateway information

Extended gateway information is included in an sFlow sampled packet if BGP is enabled. The extended gateway information includes the following BGP information about the packet's destination route:

- This router's autonomous system (AS) number
- The route's source IP AS number
- The route's source peer AS number
- The AS path to the destination

In BGP-configured routers, AS Path information is collected from each node traversed by the sFlow packets.

NOTE
AS communities and local preferences are not included in the sampled packets.

To obtain extended gateway information, use "struct extended_gateway" as described in RFC 3176

sFlow null0 sampling

This feature allows Extreme devices to sample null0 dropped packets. This is useful in cases such as DOS attack on a particular route.

Configuring steps

1. Enable sFlow.
2. Enable null0 sampling.
3. Configure null0 routes.

NOTE
Above commands can be performed in any order.

Feature characteristics

- IPv4, IPv4-VPN, IPv6 null0 routes can be sFlow sampled.
- Only explicitly configured null0 routes can be sFlow sampled. Implicit null0 drops cannot be sFlow sampled.
- By default, null0 sFlow sampling feature is disabled.

Limitations

- When this feature is enabled, due to sampling of more packets (discarded packets) than the usual number till now, the actual sampling rate for regular streams will be reduced.
- This feature does not support PBR related null0 drops.
- This feature does not support default null0 route drops.

Backward compatibility

The current sFlow functionalities and ACL based sFlow functionalities will co-exist with this feature. As the dropped packets hit TM, if mirroring is enabled on that port, these dropped packets will also get mirrored.

Enabling/disabling the null0 sFlow sampling

These commands include the enabling and disabling of the null0 sampling.

Enter the following command to enable sFlow sampling for null0 routes.

```
device(config)# sflow null0-sampling
```

To disable null0 sampling, enter the following command.

```
device(config)# no sflow null0-sampling
```

Syntax: [no] sflow null0-sampling

Configuring a null0 route

For configuring a route for null0 sampling, use the following command.

```
device(config)# ip route 10.10.10.100/32 null0
```

Syntax: [no] [**ip** | **ipv6**] **route** *ip-addr* **null0**

Displaying sFlow show command

This command will display the configuration for sFlow.

```
device(config)# show sflow
sFlow services are enabled.
sFlow management VRF is enabled.
sFlow management VRF name is default-vrf.
sFlow agent IP address: 55.55.55.56
sFlow agent IPV6 address: unspecified
sFlow source IP address: unspecified, UDP 8888
sFlow source IPv6 address: unspecified, UDP 8888
Collector IP 77.7.7.2, UDP 6343
Polling interval is 20 seconds.
Configured default sampling rate: 1 per 2048 packets.
0 UDP packets exported
124 sFlow samples collected.
133 sFlow management-vrf UDP packets dropped
0 ACL sFlow samples collected.
sFlow ports   Global Sample Rate   Port Sample Rate   Hardware Sample Rate
        1/5                2048               2048   port_down
        1/8                2048               2048   2048
sFlow Null-0 Sampling is Enabled.
```

Configuring sFlow statistics

When traffic is received in the sFlow enabled interface, packets are sent to the LP CPU. The packets are processed by sFlow module by adding sFlow header along with the packet header and thereafter sent to the sFlow collector. The statistics of sFlow samples are maintained in the sFlow collector.

Use this command in the sFlow module to display the total count per interface for both sFlow and ACL based samples in all the slots where sFlow is configured.

```
device(config)# show sflow statistics
Sflow Ports                Flow Samples count     Acl Samples Count
1/1                   800                          0
1/5                   0                          900
2/1                   600                          0
device(config)# show sflow statistics  ethernet 1/1
Sflow Ports                Flow Samples count     Acl Samples Count
1/1                   800                          0
```

Syntax: show sflow statistics

clear statistics sflow command clears all the statistics collected per interface.

```
device(config)# clear statistics sflow
device(config)# show sflow statistics  ethernet 1/1
Sflow Ports                Flow Samples count     Acl Samples Count
1/1                   0                          0
```

sFlow support for MPLS

In addition to the Layer 2 or Layer 3 information typically exported across devices, when sFlow sampling is configured on VPN endpoint interfaces, you can export MPLS or VPN information, such as VLL, VPLS, and VRF customer endpoint interfaces details. This functionality allows service providers to collect sFlow information from VPN customers.

For incoming packets to an endpoint interface sampled by sFlow, the following additional information is collected and exported in the sFlow packets:

- **MPLS VC information:** including the VC name, VC index, and VC label COS
- **MPLS tunnel information:** including the LSP tunnel name, the tunnel index as assigned by the router, and the tunnel COS used

  NOTE
  IP over MPLS (non-Layer 3 VPN or VRF) packets are not supported for sFlow processing.

sFlow with VPLS local switching

This feature allows sFlow to carry the original VLAN ID of the incoming traffic in scenarios where a VPLS instance has multiple endpoints and different endpoints with different VLAN IDs -- implementing automatic VLAN ID translation.

When VPLS CPU protection is enabled in conjunction with sFlow, hardware flooded with sFlow, hardware flooded with broadcast, multicast, and unknown unicast, packets are marked with a source VLAN ID of 0. The destination VLAN ID cannot be determined in such cases. This behavior applies to all VPLS traffic.

NOTE
You must configure MAs with different MD levels to monitor the different endpoints with different VLAN IDs in the same VPLS instance.

Configuring and enabling sFlow

To configure sFlow, you must specify the collector information. The collector is the external device to which you are exporting the sFlow data. Optionally, you can change the polling interval and the sampling rate. Next, you enable sFlow globally and then enable forwarding on individual interfaces.

Specifying the collector

sFlow exports traffic statistics to an external collector. You can specify up to four collectors. You can specify more than one collector with the same IP address if the UDP port numbers are unique. You can have up to four unique combinations of IP address and UDP port number.

To specify sFlow collectors, enter a command such as the following.

```
device(config)# sflow destination 10.10.10.1
```

This command specifies a collector with IP address 10.10.10.1, listening for sFlow data on UDP port 6343.

Syntax: [no] sflow destination ip-addr [**dest-udp-port**]

The *ip-addr* variable specifies the collector's IP address.

The *dest-udp-port* variable specifies the UDP port on which the sFlow collector will be listening for exported sFlow data. The default port number is 6343.

The sampled sFlow data sent to the collectors includes an agent_address field. This field identifies the device that sent the data. Refer to Source address on page 182.

Changing the polling interval

The polling interval defines how often sFlow byte and packet counter data for a port are sent to the sFlow collectors.

The default polling interval is 20 seconds. You can change the interval to a value from 1 to any higher value. The interval value applies to all interfaces on which sFlow is enabled. If you set the polling interval to 0, counter data sampling is disabled.

To change the polling interval, enter a command such as the following at the global CONFIG level of the CLI.

```
device(config)# sflow polling-interval 30
```

Syntax: [no] sflow polling-interval secs

The *secs* variable specifies the interval and can be from 1 to any higher value. The default is 20 seconds. If you specify 0, counter data sampling is disabled.

Changing the sampling rate

The sampling rate is the average ratio of the number of packets received on an sFlow-enabled port to the number of flow samples taken from those packets. By default, all sFlow-enabled ports use the default sampling rate, which is 2048. With a sampling rate of 1024, on average, 1 in every 1024 packets forwarded on an interface is sampled.

You can change the default (global) sampling rate. You also can change the rate on an individual port, overriding the default sampling rate.

> NOTE
> sFlow uses CPU resources to send sFlow samples to the collector. If you set a low sampling value on a high rate interface (for example 10 GbE), the interface module CPU utilization can become high.

Configuration considerations

The sampling rate is a fraction in the form 1/N, meaning that, on average, one out of every N packets will be sampled. The **sflow sample** command at the global level or port level specifies N, the denominator of the fraction. A higher denominator means a lower sampling rate because fewer packets are sampled. Likewise, a lower number for the denominator means a higher sampling rate because more packets are sampled. For example, if you change the denominator from 2,000 to 512, the sampling rate increases because four times as many packets will be sampled.

> NOTE
> It is recommended that you do not change the denominator to a value lower than the default. Sampling requires CPU resources. Using a low denominator for the sampling rate can cause high CPU utilization.

Changing the global rate

If you change the global sampling rate, the change is applied to all sFlow-enabled ports except those ports on which you have already explicitly set a sampling rate. For example, if you enable sFlow on ports 1/1, 1/2, and 5/1 and you configure the sampling rate on port 1/1 but leave the other two ports using the default rate, then changing the global sampling rate would apply to ports 1/2 and 5/1 but not port 1/1. sFlow uses the sampling rate you explicitly configured on the individual port even if you globally changed the sampling rate for the other ports.

Sampling rate for new ports

When you enable sFlow on a port, the port's sampling rate is set to the global default sampling rate. This also applies to ports on which you disable and then re-enable sFlow. The port retains the sampling rate it had when you disabled sFlow forwarding on the port, unless the sflow sampling rate is removed or moved to default rate.

Sflow sampling on CES 2000 Series and CER 2000 Series devices

NOTE
Sflow samples outbound traffic if the sflow enabled port is monitored by a mirror port.

On CES 2000 Series and CER 2000 Series devices, if mirrored Sflow packets are received in the LP CPU there is no option to distinguish them from regular Sflow packets.

Changing the default sampling rate

NOTE
The CES 2000 Series and the CER 2000 Series devices support sFlow sampling rate configuration on a per-port basis. The XMR Series and MLX Series devices support sFlow sampling rate configuration on a per-packet processor basis.

To change the default (global) sampling rate, enter a command such as the following at the global configuration level.

```
device(config)# sflow sample 1024
```

Syntax: [no] sflow sample *num*

The *num* variable specifies the average number of packets from which each sample will be taken. The sampling rate you configure is the actual sampling rate. You can enter a value from 512 through 1048576. The default is 2048.

Changing the sampling rate on a port

You can configure an individual port to use a different sampling rate than the global default sampling rate. This is useful in cases where ports have different bandwidths. For example, if you are using sFlow on 10/100 ports and Gigabit Ethernet ports, you may want to configure the Gigabit Ethernet ports to use a higher sampling rate (gathering fewer samples per number of packets) than the 10/100 ports.

To change the sampling rate on an individual port, enter a command such as the following at the configuration level for the port.

```
device(config-if-e10000-1/1)# sflow sample 8192
```

Syntax: [no] sflow sample num

The *num* variable specifies the average number of packets from which each sample will be taken. The software rounds the value you enter up to the next odd power of 2. The actual sampling rate becomes one of the values listed in Changing the default sampling rate on page 188.

Configuring the sFlow source interface

sFlow source interface is globally defined for all sFlow destinations. For detailed information about the sFlow agent, refer to sFlow event workflow on page 181.

```
device(config)#sflow source [ipv6] [[ethernet | loopback | ve | pos <interface-id>] |[null0]] [<udp-port-id>]
```

Syntax: [no] sflow source [**ipv6**] [[**ethernet** | **loopback** | **ve** | **pos** *interface-id*] | [**null0**]] [*udp-port-id*]

By default, the sFlow source interface is not specified, and the outgoing interface of an sFlow packet will be used as the source interface and address. The sFlow source port is 8888 by default.

Use the IPv6 option parameter indicate the IPv6 sFlow source address. If the destination IPv6 type does not match with the sFlow source IP address then the default behavior will be taken.

The sFlow source UDP for IPv4 is independent of IPv6.

The Null0 option is used to drop the sFlow sample with this source, while maintaining sFlow statistics.

Configuring the sFlow agent interface

The sFlow agent interface is globally defined for all sFlow destinations.

Configuration considerations

- By default, the sFlow agent is not specified, and the sFlow datagram will use the router ID as the agent ID.
- The user has the ability to configure the sFlow agent interface for IPv4 and IPv6.

```
device(config)# sflow agent [ipv6] [[ethernet | loopback | ve | pos <interface-id>]
```

Syntax: [no] sflow agent [**ipv6**] [[**ethernet** | **loopback** | **ve** | **pos** *interface-id*]

The **ipv6** keyword will indicate IP version 6 from the configured interface ID. The optional keyword will be followed by the interface type.

The command**no sflow agent** command with the specific parameters removes the specified agent interface and reassigns the agent-IP to the router-ID as in the default behavior.

Configuring the sFlow management VRF

The **sflow management-vrf-disable** command is used to disable the management VRF for sFlow and using the default VRF instance. By default, the management VRF is enabled on sFlow.

```
device(config)#sflow management-vrf-disable
```

Syntax: [no] sflow [**management-vrf-disable**]

The **no sflow management-vrf-disable** command disables the use of management VRF on sFlow and enables the default VRF instance.

> NOTE
> The output of the **show running-config** command does not show "management-vrf-disable" because it is the default behavior. If the **no sflow management-vrf-disable** command has been used, "management-vrf-disable" will appear in the output to the show running-config command.

sFlow forwarding

sFlow exports data only for the interfaces on which you enable sFlow forwarding. You can enable sFlow forwarding on Ethernet or POS interfaces.

> NOTE
> sFlow forwarding enables sampling of data packets received on sFlow-enabled ports and does not sample data packets that leave sFlow-enabled ports.

To enable sFlow forwarding:

- Globally enable the sFlow feature.
- Enable sFlow forwarding on individual interfaces.

NOTE
Before you enable sFlow, make sure the device has an IP address that sFlow can use as its source address. Refer to Source address on page 182 for the source address requirements.

Enabling sFlow forwarding

To enable sFlow forwarding, enter commands such as the following.

```
device(config)# sflow enable
device(config)# interface ethernet 1/1 to 1/8
device(config-mif-1/1-1/8)# sflow forwarding
```

These commands globally enable sFlow, then enable sFlow forwarding on Ethernet ports 1/1 through 1/8. You must use both the **sflow enable** and **sflow forwarding** commands to enable the feature.

Syntax: [no] sflow enable

Syntax: [no] sflow forwarding

NOTE
Data for POS ports is sampled using Ethernet format. The PPP or HDLC header of the sampled POS packet is replaced with an Ethernet header. PPP or HDLC control packets or IS-IS packets transmitted or received at a POS port are not sampled. Such packets are not included in the number of packets from which each sample is taken.

NOTE
sFlow packets cannot be forwarded from a management interface. You must configure an IP interface on an Interface module to forward sFlow packets.

NOTE
Configuring sFlow with Provider Bridge (PB) or Provider Backbone Bridges (PBB) port-type is not supported on the Extreme NetIron CES Series and Exrteme NetIron CER devices.

ACL-based Inbound sFlow

Multi-Service IronWare software supports using an IPv4 or IPv6 ACL to select sample traffic to be sent to an sFlow collector. The data matching an ACL clause can be collected to observe traffic flow patterns and quantities between a set of switches and routers. To accommodate collecting sFlow through standard procedures and using ACL-filtered traffic, the proprietary Tag Type 1991 encapsulates the sFlow samples obtained through ACL-based sFlow and separates them from the sequence flow of other sFlow samples. Figure 12 shows the format of an sFlow packet, which illustrates the differences between a standard sFlow payload and an ACL-based payload.

Figure 12 shows sFlow in a UDP packet. Within the UDP packet, the sFlow contents are carried in individual samples that are identified by a Tag Type and a Length variable. The standard values for the Tag Types are 1 (sampled packet) and 2 (counter sample). The Length variable describes the length of the sample. Within the sample are other variables including the Sequence number and the Source ID.

Extreme has introduced the proprietary Tag Type 1991 to identify ACL-based sFlow samples. For these samples, standard Tag Type 1 samples collected using ACL-based Inbound sFlow are encapsulated in a Tag Type 1991 sample. The Length variable identifies the entire length of the Tag Type 1991 sample including the encapsulated Tag Type 1 sample.The encapsulated sample has a Length variable of its own that only identifies the length of that sample.

The Tag Type 1991 samples are sequenced separately from the unencapsulated Tag Type 1 samples. For instance, in the packet detail described in "Sequence Flow for sFlow Records" in Figure 12, the top sFlow record with Tag Type 1 begins with the sequence number 1. The next sFlow record is Tag Type 1991, which indicates that the sample contained is from ACL-based sFlow. Encapsulated within this ACL-based sFlow sample is an sFlow sample record of Tag Type 1. The ACL-based sFlow sample (which contains the Tag Type 1 sample) is followed by an unencapsulated Tag Type 1 sFlow sample. That unencapsulated Tag Type 1 sFlow sample follows the sequence numbering of the first unencapsulated Tag Type 1 sFlow sample, which gives it a sequence number of 2.

This is useful in cases where an sFlow collector does not recognize Tag Type 1991. In these situations, the Tag Type 1991 samples can be ignored without disrupting the sFlow sequence numbers. It is also useful for identifying samples obtained using ACL-based sFlow on which other processing might be performed.

FIGURE 12 sFlow packet format

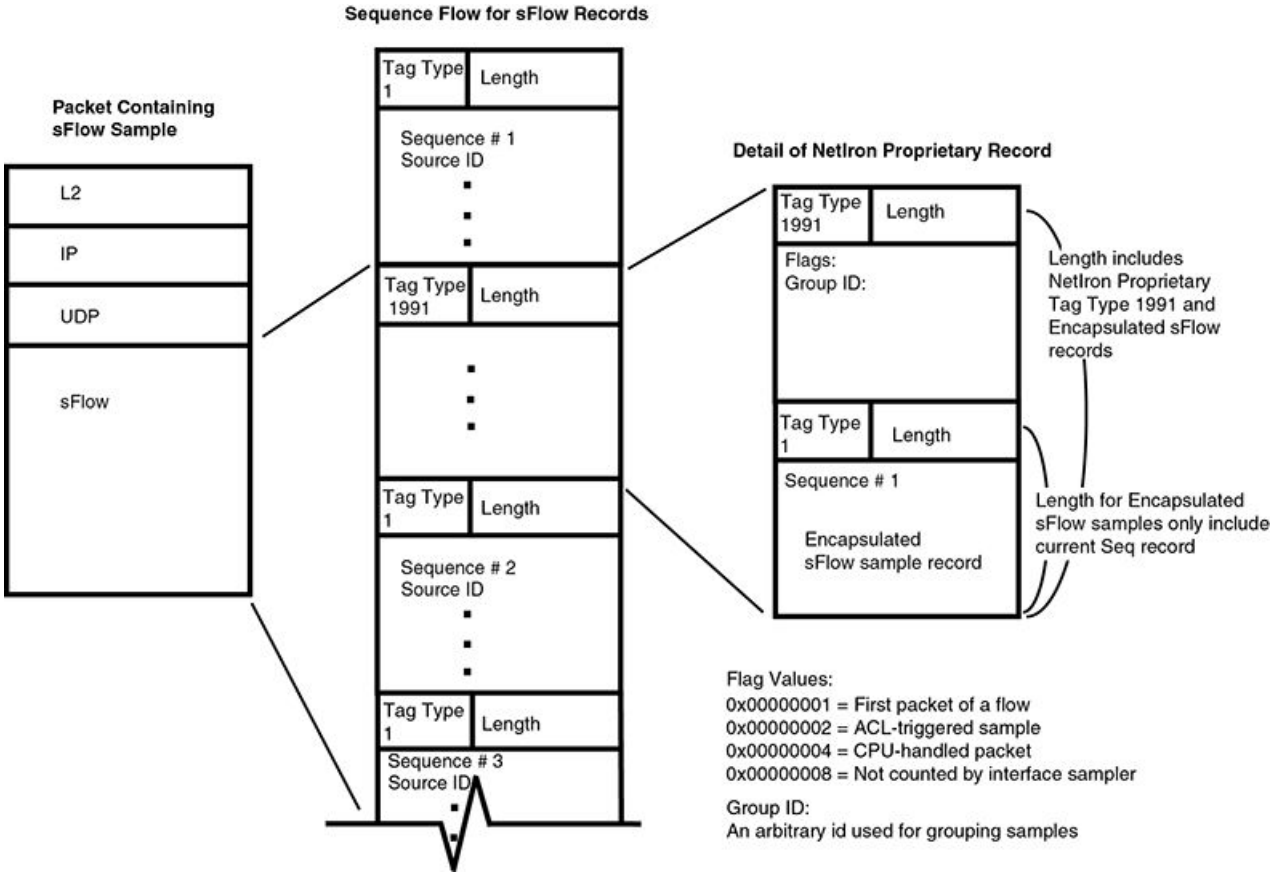

Configuring ACL-based Inbound sFlow

The following sections describe how to configure ACL-based Inbound sFlow:

- Configuration considerations for ACL-based Inbound sFlow on page 192
- Creating an ACL with an sFlow clause on page 192
- Displaying sFlow information on page 192

Configuration considerations for ACL-based Inbound sFlow

The following section describes configuration considerations for ACL-based Inbound sFlow:

- sFlow must be enabled on the router.
- **ACL-based mirroring:** The **mirror** and **copy-sflow** keywords are mutually exclusive on a per-ACL clause basis.
- **Port-based monitoring:** Port-based monitoring and ACL-based sFlow can co-exist on the same interface.
- **Port-based sFlow:** Port-based and ACL-based sFlow can co-exist on the same interface. When both features are configured on an interface, packets that qualify as ACL-based sFlow packets are sent to the collector as ACL sample packets. Also, the user can configure ACL-based sFlow on an interface without configuring port-based sFlow.
- **IP Receive ACLs:** IP Receive ACLs are used for filtering or rate-limiting management traffic. The **copy-sflow** keyword is also supported for IP Receive ACLs.
- **Policy Based Routing:** The **copy-sflow** keyword is applicable for PBR ACLs.
- **IPv4 ACL-based Rate Limiting:** When the **copy-sflow** keyword is used in an IPv4 Rate Limiting ACL, only traffic permitted by the Rate Limiting engine is copied to the CPU for forwarding to the sFlow collector.
- **IPv4 ACLs on VRF endpoints:** You can apply ACL-based sFlow for VRF endpoints; however, such packets are treated as regular sampled sFlow packets and do not carry proprietary encapsulation. This can create a minor skew of statistics projection.
- **Layer 2 ACLs:** The **copy-sflow** keyword is not supported for Layer 2 ACLs.
- If the **copy-sflow** keyword is used for a clause that is applied to the outbound direction, it is ignored.

Creating an ACL with an sFlow clause

The **copy-sflow** keyword has been added for inclusion in IPv4 and IPv6 ACL clauses to direct traffic that meets the criteria in the clause to be sent to the sFlow collector. In the following example, the ACL is used to direct syn-ack packets sent from a server at address 10.10.10.1.

```
access-list 151 permit tcp host 10.10.10.1 any established syn copy-sflow
access-list 151 permit any any
```

The **copy-sflow** keyword directs selected traffic to the sFlow collector. Traffic can only be selected using the **permit** clause.

You must apply the ACL to an interface using the**ip access-group command** as shown in the following example.

```
device(config)# int eth 1/1
device(config-if-e10000-1/1)# ip access-group 151 in
```

Displaying sFlow information

To display sFlow configuration information and statistics, enter the following command at any level of the CLI.

```
device(config)# show sflow
sFlow services are enabled.
sFlow management VRF is enabled.
sFlow management VRF name is blue.
sFlow agent IP address: 10.25.120.1
sFlow agent IPV6 address: unspecified
sFlow source IP address: unspecified, UDP 9999
sFlow source IPv6 address: 22::32, UDP 5544
2 collector destinations configured:
Collector IP 10.25.120.10, UDP 6343
Collector IPV6 10::32, UDP 6343
Polling interval is 20 seconds.
Configured default sampling rate: 1 per 2048 packets.
352 UDP packets exported
1 sFlow samples collected.
0 sFlow management-vrf UDP packets dropped
```

```
0 ACL sFlow samples collected.
sFlow ports Global Sample Rate Port Sample Rate Hardware Sample Rate
1/4 2048 10000 32768
```

Syntax: show sflow

Table 17 shows the output information provided by the **show sflow** command.

TABLE 17 sFlow information

Field	Description
sFlow services	The feature state, which can be one of the following: • disabled • enabled
sFlow management VRF	Indication that sFlow is enabled to use the management VRF. Disabled means that sFlow is using the non-management VRF instance.
sFlow management VRF name	Management VRF name, if the management VRF is enabled on sFlow.
sFlow agent IP address	The IP address that sFlow is using in the agent_address field of packets sent to the collectors. Refer to Source address on page 182.
sFlow agent IPV6 address	The sFlow agent IPv6 address is unspecified by default. If configured, it will correspond to the IP address on the configured interface.
sFlow source IP address	The sFlow source IP address corresponds to the IP address on the configured interface. If an IP address is not configured on the interface, then it will be unspecified. However, if the source interface is null0, then it will be null0 on the interface.
sFlow source IPv6 address	The sFlow source IPv6 address that corresponds to the IP address on the configured interface. If an IP address is not configured, then the interface will be unspecified. However, if the interface is null0, then the configured interface will be null0.
UDP	The sFlow source UDP port default is 8888.
Collector	The collector information. The following information is displayed for each collector: • IP address • UDP port If more than one collector is configured, the line above the collectors indicates how many have been configured.
Polling interval	The port counter polling interval.
Configured default sampling rate	The configured global sampling rate. If you changed the global sampling rate, the value you entered is shown here.
UDP packets exported	The number of sFlow export packets the device has sent. NOTE Each UDP packet can contain multiple samples.
sFlow samples collected	The number of sampled packets that have been sent to the collectors.
sFlow ports	The ports on which you enabled sFlow.
Global Sample Rate	The global sampling rate for the device.
Port Sampling Rate	The sampling rates of a port on which sFlow is enabled.
Hardware Sample Rate	The actual sampling rate. This is the same as the Global Sample Rate

Displaying ACL-based sFlow statistics

Use the **show sflow** command to display the number of sFlow samples collected for ACL-based sFlow. These statistics are shown in bold in the following display.

```
device# show sflow
sFlow services are disabled.
sFlow agent IP address: 10.10.10.254
Collector IP 10.10.10.1, UDP 6343
Polling interval is 30 seconds.
Configured default sampling rate: 1 per 1024 packets.
0 UDP packets exported
0 sFlow samples collected.
5 ACL sFlow samples collected
sFlow ports   Global Sample Rate   Port Sample Rate   Hardware Sample Rate
        4/1                1024               8192   8192
```

Viewing BGP AS path sFlow statistics

The output of the **show sflow** command displays sFlow configuration information, and the elapsed time after the last sampling used for the BGP AS path table, and the interval for cleaning up the AS path table. The following example output shows that the AS path table has not been sampled within the last 51,385 seconds and that 3600 seconds, which is the default value, is the configured clean up interval.

```
device(config)# show sflow

Slot 1 1/1 is disabled for sflow with sample rate = 2048 (actual rate = 2048)

Slot 1 1/2 is disabled for sflow with sample rate = 2048 (actual rate = 2048)

...

Slot 1 1/19 is disabled for sflow with sample rate = 2048 (actual rate = 2048)

Slot 1 1/20 is disabled for sflow with sample rate = 2048 (actual rate = 2048)

sflow destinations :

Total sflow sampling time = 0 (0)

Total sflow UDP time = 0 (0)

Total sflow ppcr tx time = 0 (0)

No sflow sampling on AS Path for 51385 sec

Sflow as path clean up wait interval 3600 sec
```

Clearing sFlow statistics

To clear the UDP packet and sFlow sample counters, use the **clear statistics** command.

```
device(config)# clear statistics
```

Syntax: clear statistics [**sflow**]

The **sflow** option clears the following values:

- UDP packets exported
- sFlow samples collected
- sFlow UDP packets dropped
- ACL sFlow samples collected

System Monitoring

System monitoring overview

System monitoring (Sysmon) is implemented to monitor the overall system's health. Sysmon is a system-wide, modular monitoring service. It monitors different system components of a device to determine if those components are operating correctly.

Sysmon periodically monitors the system for defined event types such as errors on TM and FE links. Sysmon runs as a background process. It has a default policy that controls what is monitored and what actions will be taken if a fault is detected. Sysmon generates the following log outputs for the monitoring information.

- Syslog
- Sysmon internal log

NOTE
Syslog reported Sysmon alarm messages should be reported to Extreme Technical Support.

Internal logs are generated to give more information to Extreme Technical Support when a problem occurs. The existence of internal logs doesn't mean the system is experiencing problems, or that some actions need to be taken. If Sysmon detects a failure, it will report the failure by generating the syslog messages. In some cases the failed device will be shutdown or isolated from the system. In other cases the software may attempt to recover the failed device.

Overall system performance depends on how resources are utilized. Any shortage of resources impacts the overall performance of a system. The system resource histogram feature provides detailed information on how system resources are used. It collects information on task CPU usage, buffer usage and memory usage and stores this information in internal memory.

Runtime diagnostics are a critical component of a networking system to provide maximum uptime by detecting and isolating faults, and then recovering from them. A system runtime diagnostics framework supports execution of diagnostic tests such as the port CRC error monitoring test. It manages this background diagnostic test and provides mechanisms for taking corrective action.

Event monitoring

This section discusses the following topics:

- Event monitoring overview on page 196

Event monitoring overview

Sysmon monitors a number of event types periodically, detecting errors based on polling and interrupt. Polling is the reading of specific hardware registers, while interrupt is an instantaneous event detection. Sysmon continuously monitors management processor (MP) and interface processors (LPs) by the way of polling and interrupt methods. Once a threshold is reached, Sysmon logs the event in the internal Sysmon log and takes one of following actions based on the event type:

- TM_LINK or FE_LINK monitoring:
  - Syslog—Generates a syslog message
  - Shutdown link—Disables the link between the TM and the FE
  - SNMP trap—Generates an SNMP trap
- Port CRC test:
  - Syslog—Generates a syslog message
  - Port down—Disables the port
  - No action—No action is taken
- NP memory error monitoring:
  - No action—Disables monitoring of memory errors on interface modules
  - Syslog—Generates a syslog message
  - Syslog and SNMP trap—Generates a syslog message and an SNMP trap
  - SNMP trap—Generates an SNMP trap

By default, Sysmon is enabled to monitor and detect the defined event types. The following Sysmon event types are defined and implemented:

- TM_LINK—Monitoring TM SerDes links
- FE_LINK —Monitoring FE SerDes links
- NP memory errors—Monitoring memory errors on interface modules
- Port CRC errors—Monitoring for excess packet CRC errors on each port
- LP high CPU usage—Monitoring for high LP CPU usage
- MP high CPU usage—Monitoring for high MP CPU usage

  NOTE
  By default, Sysmon does not monitor LP or MP CPU usage event types. You must enable Sysmon to monitor them.

Event types

TM_LINK

TM link is the link between the line card and the switch fabric module. The event type TM_LINK monitors this link for the errors reported on the link by the TM, such as CRC, misalignment, code group error, and down links. Here is an example from Syslog.

```
Dec 29 15:31:24:W:System: ALARM:LP15/TM3 has 6 links, less than the minimum to maintain line rate
```

FE_LINK

FE link is the link between the line card and the switch fabric module. The event type FE_LINK monitors this link for the errors reported on the link by the FE, such as CRC, misalignment, code group error and down links. Here is an example from Syslog.

```
Dec 29 15:31:24:W:System: ALARM:LP15/TM3 has 6 links, less than the minimum to maintain line rate
```

NP interface memory errors

The NP Memory Error Monitoring event monitors memory errors on interface modules. Monitoring includes parity errors, ECC errors, overflow and underflow errors. Errors are reported as syslog messages or SNMP traps. Here is an example from Syslog.

```
Feb 23 19:27:29:E:PRAM Word 2 Parity Error on port range 3/1 - 3/2
```

LP CPU high-usage

The LP CPU high-usage monitoring event monitors the CPU usage on interface modules. Monitoring is enabled for a default usage value (threshold). The CPU usage is monitored and any excursion above the threshold in a 100 ms window creates a syslog message. If the CPU usage remains above the threshold for 300 ms, a debug file is created with information relevant to identifying the cause. If the CPU usage falls below the set threshold before the 300 ms mark, a syslog message is generated but no debug file is created. The following examples highlight the three cases.

```
SYSLOG: <14>May 15 00:19:51 Eltanin-R3 LP High CPU: LP 1. Status: Threshold EXCEEDED

SYSLOG: <14>May 15 00:19:51 Eltanin-R3 LP High CPU: LP 1. Status: Logs CAPTURED

SYSLOG: <14>May 15 00:19:51 Eltanin-R3 LP High CPU: LP 1. Status: Condition CLEARED
```

For detailed information on LP CPU High-usage monitoring, refer to LP CPU high-usage monitoring on page 236.

MP high CPU usage

The MP high CPU usage event monitors the CPU usage on the active and standby MPs. Monitoring is enabled for default usage and task threshold values. The system monitors the percentage of the CPU uses and the amount of time that a task holds the CPU, and creates a log file when either threshold is exceeded.

For detailed information on MP CPU high-usage monitoring, refer to MP CPU high-usage monitoring on page 237.

Displaying event information

Displaying internal log messages

You can use the following show commands to view the results of the monitoring activity. These show commands display information for all event types in one output.

To display the contents of the internal log, enter the following command.

```
device# show sysmon logs
INFO:May 13 07:29:54: TM Link Error: LP2/TM2/Link2 -- SNM3/FE3/Link43 (disabled)
INFO:May 13 07:29:33: FE Link Error: SNM3/FE3/Link64 -- LP4/TM1/Link2 (disabled)
```

Syntax: show sysmon logs

NOTE

The size of the internal log table is 10,000 logs.

Clearing internal logs

To clear the internal logs, enter the following command.

```
device# clear sysmon logs
```

Syntax: clear sysmon logs

Displaying current SYSMON configuration

Enter the **show sysmon configuration** command to view the current configuration for system monitoring services. Look for output similar to the following:

```
device# show sysmon config
---------------------------+----------------+-----------+----------+------------
EVENT                      | ACTION         |POLL PERIOD|THERESHOLD|LOG BACK-OFF
                           |                |   (SEC)   |#(PER POLL|
                           |                |           | in #POLL)|
---------------------------+----------------+-----------+----------+------------
TM. Link Monitoring        |SHUTDOWN-LINK   |    60     | 5 in 10  |  1800
---------------------------+----------------+-----------+----------+------------
Port CRC Monitoring        |SYSLOG          |    60     | 3 in 5   |  1800
---------------------------+----------------+-----------+----------+------------
FE. Link Monitoring        |SHUTDOWN-LINK   |    60     | 5 in 10  |  1800
---------------------------+----------------+-----------+----------+------------
NP Memory Error Monitoring |SYSLOG-AND-TRAP |    10     |   N/A    |  N/A
---------------------------+----------------+-----------+ ---------+------------
```

Saving system information to Flash overview

System state information can be captured from a Management Processor (MP) as well as the Line card (LP) and stored in the MP storage card (slot 2).

To aid debugging at a customer site, the following system state information can be captured in a memory dump:

- All task-related information
- Registers related to Traffic Manager (TM), Packet Processor (XPP) and PCIe Bus Interface (PBIF)
- Memory pools

The collection of system state information can be initiated from the MP or the LP using CLI commands. When initiated from the MP, only MP state information is captured by default. If configured, memory from only one line card can also be captured along with the MP memory dump. Information can be captured for a specific line card from the corresponding LP.

NOTE
This feature is supported only for MR-2 management modules.

To trigger the memory dumps there are different commands to run on the LP and the MP. When the memory dump is triggered from an LP, the line card resets after capturing the system state information. When the memory dump is triggered from an MP, the management processor reloads after capturing the system state information. The memory dump files can be uploaded to a TFTP server and used for further debugging.

CAUTION
Do not execute a memory dump from the MP and an LP at the same time.

Configuring and triggering a memory dump from a line card

System state information can be captured from a line card to help with debugging.

After remotely connecting to the line card, you can trigger a memory dump.

1. In privileged exec mode, remotely connect to the line card.

```
device# rcon 1
```

2. From the line card prompt, use the **enable** command to enter privileged exec mode.

```
linecard1> enable
```

3. Trigger a memory dump from the line card.

```
linecard1# reset-memdump
```

The following example triggers a memory dump on the line card and resets the line card.

```
device# rcon 1
linecard1> enable
linecard1# reset-memdump
```

Configuring and triggering a memory dump from an MP

System state information can be captured from an MP to help with debugging. After the memory dump, the device reloads in the form of a warm reboot.

Perform this task from a management processor (MP). The first three steps are optional to allow you to configure a memory and register dump from one specified line card.

1. (Optional for line card configuration) Enter global configuration mode.

```
device# configure terminal
```

2. (Optional for line card configuration) Identify a line card from which a memory and register dump can be captured.

```
device(config)# memdump slot 1
```

3. (Optional for line card configuration) Exit to privileged exec mode.

```
device(config)# exit
```

4. Trigger a memory dump from the MP and the line card in slot 1. After the memory dump, the device reloads.

```
device# reload-memdump
```

5. After the reload, display the memory dump files on the MP.

```
device# dir /slot2

Directory of /slot2

09/20/2016 08:34:21      123,691,320  memdump_mp.txt
09/16/2016 10:22:00      116,464,260  memdump_lp.txt
09/16/2016 10:22:00      8,669  memdump_registers.txt
09/16/2016 10:22:00      4,669  memdump_mp_metadata.txt
09/16/2016 10:22:00      2,669  memdump_lp_metadata.txt
   5 File(s)      240,171,587 bytes
```

The following example identifies that a memory dump can be captured from the MP and slot 1.

```
device# configure terminal
device(config)# memdump slot 1
device(config)# exit
device# reload-memdump
```

Histogram information

This section discusses the following topics:

Histogram information overview

The histogram framework feature monitors and records system resource usage information. The main objective of the histogram is to record resource allocation failures and task CPU usage information. The histogram feature keeps track of task execution information, context switch history of tasks, buffer allocation failure and memory allocation failure.

The histogram information is collected and maintained internally, in a cyclical buffer. It can be reviewed to determine if resource allocation failures or task CPU usage may have contributed to an application failure.

NOTE
Histogram information is not maintained accross reboot

Displaying CPU histogram information

The CPU histogram provides information about task CPU usage. The CPU histogram is viewed in the form of buckets i.e., task usage is divided into different interval levels called buckets. For example, the task run time is divided into buckets - bucket 1(0-50ms), bucket2 (50-100ms), bucket3 (100-150ms) etc. The CPU histogram collects the task CPU usage in each bucket. This includes how many times a task run/hold time falls in each bucket, max run time and total run time for each bucket. CPU histogram information is measured for hold-time, wait-time, system timer time, and user interrupt time of the task.

- Hold time - time that the task is holding the CPU without yield.
- Wait time - time that the task is waiting for execution.
- Timer time - time that task is handling the timer routines without yielding the CPU.
- Interrupt time - time that the task is handling the user interrupt routines without yielding the CPU.

Show commands

To display task hold time information, enter the following command:

```
device# show cpu histogram hold
HISTOGRAM CPU HISTOGRAM INFO
----------------------------
 No of Bucket       : 51
 Bucket Granularity : 10 ms
 Last cleared at    : 2012.07.10-07:29:20.704
 No of Task         : 67
```

```
Task Name  Bkt     Bkt      No of Time   HoldTime   HoldTime                  Time
           Num     Time(ms)              Total(s)    Max(ms)
-------------------------------------------------------------------------------
ip_rx       1      000-010          4    .000463     .201  2012.07.10-07:29:20.701
vlan        1      000-010          1    .000025     .025  2012.07.10-07:29:20.700
mac_mgr     1      000-010          1    .000010     .010  2012.07.10-07:29:20.701
mrp         1      000-010          1    .000025     .025  2012.07.10-07:29:20.700
erp         1      000-010          1    .000025     .025  2012.07.10-07:29:20.700
mxrp        1      000-010          1    .000009     .009  2012.07.10-07:29:20.700
rtm         1      000-010          1    .000062     .062  2012.07.10-07:29:20.700
rtm6        1      000-010          1    .000091     .091  2012.07.10-07:29:20.700
ip_tx       1      000-010          1    .000207     .207  2012.07.10-07:29:20.700
l2vpn       1      000-010          1    .000018     .018  2012.07.10-07:29:20.701
ospf        1      000-010          1    .000046     .046  2012.07.10-07:29:20.700
isis        1      000-010          1    .000009     .009  2012.07.10-07:29:20.700
mcast       1      000-010          1    .000017     .017  2012.07.10-07:29:20.700
ospf6       1      000-010          1    .000012     .012  2012.07.10-07:29:20.700
mcast6      1      000-010          1    .000012     .012  2012.07.10-07:29:20.700
web         1      000-010          1    .000029     .029  2012.07.10-07:29:20.700
lacp        1      000-010          1    .000013     .013  2012.07.10-07:29:20.700
loop_detect 1      000-010          1    .000009     .009  2012.07.10-07:29:20.701
cluster_mgr 1      000-010          1    .000011     .011  2012.07.10-07:29:20.701
telnet_0    1      000-010          4       .003        3  2012.07.10-07:29:20.672
-------------------------------------------------------------------------------
```

Syntax: show cpu histogram { **hold** | **wait** | **interrupt** | **timer** } [**taskname***name* | **above** *threshold-value* | **noclear**]

The *hold* parameter displays the task hold time histogram. The *wait* parameter displays the task wait time histogram. The *interrupt* parameter displays the task user-interrupt usage histogram. The *timer* parameter displays the task sys-timer time usage histogram.

When the *taskname name* variable is specified, the histogram information for the specified task only, is displayed. The*above threshold-value* variable specifies the display of histogram information for tasks whose maximum hold time is above the specified threshold level.

By default, task values are cleared on read. The *noclear* parameter displays information without clearing the values.

To display sequence of task execution information, enter the following command:

```
device# show cpu histogram sequence
HISTOGRAM TASK SEQUENCE INFO
----------------------------
 THRESHOLD  : 10 ms
 DURATION   : 30 s
----------------------------
Seq No Task Name    Context  HoldTime   Start Time     End Time       Date
                             Max (ms)
-------------------------------------------------------------------------------
     1 snms         TASK           16 07:33:08.790 07:33:08.806 2012.07.10
     2 snms         TASK           16 07:33:08.772 07:33:08.789 2012.07.10
     3 snms         TASK           17 07:33:08.755 07:33:08.772 2012.07.10
     4 snms         TASK           16 07:23:08.790 07:23:08.806 2012.07.10
     5 snms         TASK           16 07:23:08.772 07:23:08.789 2012.07.10
     6 snms         TASK           17 07:23:08.755 07:23:08.772 2012.07.10
     7 snms         TASK           16 07:13:08.790 07:13:08.806 2012.07.10
     8 snms         TASK           16 07:13:08.772 07:13:08.789 2012.07.10
     9 snms         TASK           17 07:13:08.755 07:13:08.772 2012.07.10
    10 snms         TASK           16 07:03:08.790 07:03:08.806 2012.07.10
    11 snms         TASK           16 07:03:08.772 07:03:08.789 2012.07.10
    12 snms         TASK           17 07:03:08.755 07:03:08.772 2012.07.10
    13 snms         TASK           16 06:53:08.790 06:53:08.806 2012.07.10
    14 telnet_0     TASK           50 09:51:50.091 09:51:50.142 2012.07.05
    15 telnet_0     TASK           50 09:51:35.184 09:51:35.234 2012.07.05
    16 console      TASK           50 09:51:11.451 09:51:11.501 2012.07.05
    17 telnet_0     TASK           50 09:47:01.459 09:47:01.509 2012.07.05
    18 console      TASK           52 09:46:32.443 09:46:32.496 2012.07.05
    19 mpls         TIMER          12 09:46:32.428 09:46:32.441 2012.07.05
    20 telnet_0     TASK           54 09:46:03.018 09:46:03.072 2012.07.05
    21 telnet_0     TASK           52 09:44:31.749 09:44:31.802 2012.07.05
    22 telnet_0     TASK           50 09:44:17.984 09:44:18.034 2012.07.05
    23 telnet_0     TASK           50 09:43:43.638 09:43:43.689 2012.07.05
    34 telnet_0     TASK           12 09:43:43.623 09:43:43.636 2012.07.05
```

```
    35 telnet_0      TASK           54 09:43:20.669 09:43:20.724 2012.07.05
    36 snms          TASK           16 09:43:08.740 09:43:08.756 2012.07.05
    37 snms          TASK           16 09:43:08.723 09:43:08.740 2012.07.05
-------------------------------------------------------------------------------
```

Syntax: show cpu histogram sequence [**taskname** *name* | **above** *threshold-value*| **trace**]

The *sequence* parameter displays sequential task execution information. Sequential execution of task information is recorded when a task's hold time is greater than the specified threshold value. The task sequence is maintained for a specific period of time and stored in a cyclic buffer, so the oldest record is overwritten by a new record.

When the *taskname name* variable is specified, the histogram information for the specified task only, is displayed. The *above threshold-value* variable specifies the display of histogram information for tasks whose maximum hold time is above the specified threshold level.

The *trace* parameter displays high CPU condition task traces.

Clearing task sequence information

To clear CPU histogram sequence information, enter the following command:

```
device(config)# clear cpu histogram sequence
```

Syntax: clear cpu histogram sequence

Displaying buffer histogram information

The main objective of the buffer histogram is to see if there was any buffer exhaustion in the last few seconds (10-60sec). Buffer usage is collected when available buffers in the 2K buffer size pool fall below the reserved limit. The threshold limit is defined in terms of BM allocate request type.

TABLE 18 Threshold values for different buffer allocation request types.

Buffer Pool	BM Allocate Request Type	Buffer allocated if available buffers
2K	OS, SDS, RCON	Above 0
2K	TX, RX Critical	Above 128
2K	IPC High	Above 512
2K	Data High	Above 700
2K	IPC Low	Above 850
2K	RX Low	Above 1024

Show commands

To display buffer histogram information, enter the following command:

```
device# show bm histogram
HISTOGRAM BUFFER SEQUENCE INFO
------------------------------
 DURATION   : 60 s
 SEQ IDX    : 1
 TIME       : 2012.07.10-09:46:59.061
 THRESHHOLD : Below RX limit (1129)
POOL-ID  SIZE(KB)    TOTAL     FREE   IN-USE  APP-OWN
-------------------------------------------------------
       3        2     6144     1024     5120     1248
-------------------------------------------------------
Task Name          App-Owns(buffers)
-----------------------------------------
mac_mgr                  13
```

```
ip_tx                    12
rtm                      14
mcast                   112
console                  11
ip_rx                    16
rtm6                     23
mcast6                   46
mpls                     71
nht                      92
l2vpn                    98
-------------------------------------------
```

To display the buffer allocation stack for the top three tasks (in terms of buffer ownership), enter the following command:

```
device(config)# show bm histogram trace 3
HISTOGRAM BUFFER SEQUENCE INFO
------------------------------
DURATION : 60 s
SEQ IDX : 1
TIME : 2013.02.07-10:39:34.334
THRESHHOLD : Below IPC Critical limit (128)
POOL-ID SIZE(KB) TOTAL FREE IN-USE APP-OWN
-----------------------------------------------------
   3     2        6144 128   6016     58
-----------------------------------------------------
Task Name App-Owns(buffers)
-------------------------------------------
mac_mgr        3
ip_tx          6
rtm            5
mcast         12
console        1
rtm6           4
mcast6         6
mpls           1
nht            2
telnet_34     18
-------------------------------------------
[ Taskname : telnet_34 , AppId : 98 ]
[ Taskname : mcast , AppId : 17 ]
[
00055a38: dev_bm_get_buf_internal
000557c0: dev_bm_get_ipc_buf
00005024: xsyscall
2037791c: ipc_get_buffer
2038fe18: allocate_a_dy_sync_packet
2039120c: init_dy_sync_mgmt
20d08c18: l2mcast_metro_vpls_init_mac_entry_sync_mgmt]
[
00055a38: dev_bm_get_buf_internal
000557c0: dev_bm_get_ipc_buf
00005024: xsyscall
2037791c: ipc_get_buffer
2038fe18: allocate_a_dy_sync_packet
2039120c: init_dy_sync_mgmt
20ced240: l2mcast_init_mdb_sync_mgmt]
```

Syntax: show bm histogram [**priority** *threshold-value* | **trace**]

The *priority threshold-value* variable displays histogram information for the specified buffer priority level only. The valid range is 0-5 (0-Critical, 1-Hi Tx, 2-Hi IPC Rx, 3-Hi Data Rx, 4-Low IPC Rx, 5-Low Data Rx).

The *trace* parameter displays the buffer allocation stack of the top three tasks (in terms of buffer ownership).

Clearing buffer historgram data

To clear the buffer histogram data, enter the following command:

```
device(config)# clear bm histogram
```

Syntax: clear bm histogram

Low buffer syslogs

Syslog messages are generated when when available buffers fall below the 20, 10 and 5 percent buffer thresholds.

```
SYSLOG: <14>Feb 7 10:39:58 Ni-MLX-Sys-6 System: Low buffer, Available buffer goes Below 20%, Available
Buffer (2243) on MP
SYSLOG: <12>Feb 7 10:40:40 Ni-MLX-Sys-6 System: Low buffer, Available buffer goes Below 10%, Available
Buffer (1633) on MP
SYSLOG: <9>Feb 7 10:41:11 Ni-MLX-Sys-6 System: Low buffer, Available buffer goes Below 5%, Available Buffer
(1328) on MP
```

SYSLOG: <10>Feb 7 10:47:34 Ni-MLX-Sys-6 System: Out of buffer, Below IPC Critical limit (128) on MP

Displaying memory histogram information

System memory is divided into five memory pools: OS, Shared, Global, User Private and DMA. The memory histogram keeps track of each memory allocation/deallocation request from an application. It helps to identify memory leak and memory usage accross the task. It also monitors the under usage condition and reports to the system. The memory histogram is recorded when available memory goes below the threshold limit on each memory pool. The threshold limit is defined in terms of percentage of available memory (20%, 10% or 5%).

To display memory histogram information, enter the following command:

```
device# show memory histogram
HISTOGRAM MEMORY SEQUENCE INFO
------------------------------
 DURATION   : 60 s
 SEQ IDX    : 1
 TIME       : 2012.07.10-11:14:08.539
 AVAIL MEM : below 5 %
-----------------------------------------------------------
POOL      Total Memory      Used Memory Available Memory
               (bytes)          (bytes)           (bytes)
-----------------------------------------------------------
Global      2855272448       2843262976         12009472
-----------------------------------------------------------
Task Name            Alloc-Number    Alloc-Size(bytes)
-----------------------------------------------------------
main                         1355             28486529
itc                             4                  645
tmr                            63                10173
ip_rx                         425               396453
scp                           748             17995881
lpagent                        63                31309
console                       101              3515673
vlan                           44              5814177
mac_mgr                        40              2305485
mrp                            26                 8541
vsrp                           28                 8557
erp                            28                 8557
mxrp                           26                 7527
snms                          192               188337
rtm                            98             33724605
rtm6                          109              1918717
ip_tx                         151              1274437
rip                            70               323733
ospf_msg_task                  17                 7453
telnet_0                       28                 7689
telnet_1                       29                 7817
-----------------------------------------------------------
```

Syntax: show memory histogram [**pool** *pool-id* | **below** *threshold-value*]

The *pool pool-id* variable specifies the display of memory histogram information for a specific memory pool. The valid range for the *pool pool-id* variable is 0-3, where 0 = OS, 1 = Shared, 2 = Global and 3 = User Private. The *below threshold-value* variable specifies the display of memory histogram information when available memory falls below the specified percentage (5, 10 or 20 percent).

Low memory syslogs

Syslog messages are generated when available memory falls below the 20, 10, and 5 percent thresholds.

```
SYSLOG: <14>Feb 7 10:50:11 Ni-MLX-Sys-6 System: Low physical memory, Pool(2-Global) below 20%, available
pool memory (225480704), physical memory (225480704) on MP
SYSLOG: <9>Feb 7 10:50:11 Ni-MLX-Sys-6 System: Low pool memory, Pool(2-Global) below 5%, available pool
memory (171204608), physical memory (171204608) on MP
SYSLOG: <12>Feb 7 10:50:12 Ni-MLX-Sys-6 System: Low physical memory, Pool(2-Global) below 10%, available
pool memory (118108160), physical memory (118108160) on MP
SYSLOG: <9>Feb 7 10:50:12 Ni-MLX-Sys-6 System: Low physical memory, Pool(2-Global) below 5%, available pool
memory (64421888), physical memory (64421888) on MP
SYSLOG: <10>Feb 7 10:50:12 Ni-MLX-Sys-6 System: Low pool memory, Pool(2-Global) below 1%, available pool
memory (28532736), physical memory (28532736) on MP
SYSLOG: <10>Feb 7 10:50:12 Ni-MLX-Sys-6 System: Low physical memory, Pool(2-Global) below 1%, available
pool memory (10731520), physical memory (10731520) on MP
```

Clearing memory historgram data

To clear the memory histogram data, enter the following command:

```
device(config)# clear memory histogram
```

Syntax: clear memory histogram

NP memory error monitoring

This sections discusses the following topics:

NP memory error monitoring overview

It can be useful to know when memory errors occur on interface modules. NP memory error monitoring periodically monitors for external and internal memory errors and reports these errors as syslog messages or generates SNMP traps.

For details of specific errors that may occur on interface cards that support NP memory error monitoring, refer to the NP memory errors section.

NP memory error monitoring: basic configuration

By default:

- NP memory error monitoring is enabled.
- Errors generate both a syslog message and a SNMP trap.
- The polling period time is 60 seconds.

Configuring NP memory error monitoring

You can configure:

- The polling frequency.
- How the errors are reported.

To set the polling frequency for NP memory errors at 10 second intervals, enter the following command:

```
device(config)# sysmon np memory-errors polling-period 10
```

To configure NP memory error monitoring to generate syslog messages, use the following command:

```
device(config)# sysmon np memory-errors action syslog
```

You may want to disable error reporting if, for example, a hardware fault exists and is generating a lot of errors. To disable reporting of NP memory errors, use the following command:

```
device(config)# sysmon np memory-errors action none
```

The following example disables monitoring of memory errors on interface modules.

```
device(config)# no sysmon np memory-errors
```

The **no** form of the command specifying a *poll-interval* value restores the default polling period. For example, the following command restores the polling period to 60 seconds.

```
device(config)# no sysmon np memory-errors polling-period 1000
```

The **no** form of the command specifying the **action** as **syslog-and-trap** , **syslog** , or **trap** will remove that action. The following command removes the **syslog** action.

```
device(config)# no sysmon np memory-errors action syslog
```

The **no** form of the command specifying the **action** as **none** will restore the default action (**syslog-and-trap)** . To restore the NP memory error action to **syslog-and-trap** , enter the following command:

```
device(config)# no sysmon np memory-errors action none
```

Syntax: [no] sysmon np memory-errors { **polling-period** *secs* | **action** { **syslog-and-trap** | **syslog** | **trap** | **none** } }

The *polling-period secs* variable specifies the frequency of polling for NP memory errors. The range is from 1 through 65535. The default value is 60 seconds.

The *action* parameter specifies the action taken when NP memory errors are detected. If the *action* parameter is set to *none* , NP memory errors are not reported. Setting the *action* parameter to *syslog* specifies the generation of a syslog message. Setting the *action* parameter to *trap* specifies the generation of a SNMP trap. If *action* is configured as *syslog* followed by configuration as *trap*, then the *action* will become *syslog-and-trap*. The default *action* is *syslog-and-trap.*

The **no** form of this command restores the default action.

NOTE
The *polling-period* parameter determines the interval between checks for NP memory errors. Reporting may not happen within the polling interval. It may be delayed by factors such as a high CPU load on the interface module or the management module, by low memory etc.

NOTE
The *action* parameter controls the generation of syslog messages or SNMP traps: they cannot be controlled by the **no snmp-server enable traps** command or the **no logging enable** command.

NOTE
Memory errors are detected on the interface module. Errors may not be reported if there is a communication problem between the management module and the interface module.

NP memory errors

The Sysmon NP memory error monitoring event monitors memory errors on interface modules. The following table lists the interface cards that support NP memory error monitoring and details the NP memory errors that are supported on each interface card.

The following interface cards support NP memory error monitoring:

- BR-MLX-10GX4-IPSEC
- BR-MLX-10GX20
- BR-MLX-10GX20-X2
- BR-MLX-10GX24
- BR-MLX-100GX2-CFP2
- BR-MLX-100GX2-CFP2-X2
- BR-MLX-100GX2-X(100G)
- BR-MLX-40GX4-X
- Gen-1
  - NI-MLX-10GX4
  - NI-XMR-10GX4
- Gen-1.1
  - BR-MLX-10GX4-X
  - BR-MLX-1GCX24-X
  - BR-MLX-1GFX24-X
- Gen-2
  - BR-MLX-10GX8-X
  - NI-MLX-10GX8-D
  - NI-MLX-10GX8-M

TABLE 19 NP memory errors supported on BR-MLX-100GX2-CFP2, BR-MLX-10GX20, and BR-MLX-10GX4-IPSEC interface cards

Error	Description
External memory errors	
1	CAM1 Packet Error 0
2	CAM1 Packet Error 1
3	CAM3 Packet Error 0
4	CAM3 Packet Error 1
5	PRAM Word 0 Parity Error
6	PRAM Word 1 Parity Error
7	PRAM Word 2 Parity Error
8	PRAM Word 3 Parity Error
9	PRAM Word 4 Parity Error
10	PRAM Word 5 Parity Error

TABLE 19 NP memory errors supported on BR-MLX-100GX2-CFP2, BR-MLX-10GX20, and BR-MLX-10GX4-IPSEC interface cards (continued)

Error	Description
11	PRAM Word 6 Parity Error
12	PRAM Word 7 Parity Error
13	CAM2PRAM Word 0 Single Bit Parity Error
14	CAM2PRAM Word 1 Single Bit Parity Error
15	CAM2PRAM Word 2 Single Bit Parity Error
16	CAM2PRAM Word 3 Single Bit Parity Error
17	CAM2PRAM Word 0 Double Bit Parity Error
18	CAM2PRAM Word 1 Double Bit Parity Error
19	CAM2PRAM Word 2 Double Bit Parity Error
20	CAM2PRAM Word 3 Double Bit Parity Error
21	LBLRAM Word 0 Parity Error
22	LBLRAM Word 1 Parity Error
23	CAM1 Word 0 Parity Error
24	CAM1 Word 1 Parity Error
25	CAM1 GIO Parity Error
26	CAM1 PEO Parity Error
27	CAM1 Operation Error
28	CAM1 Result Bus Parity Error
29	CAM2 Word 0 Parity Error
30	CAM2 Word 1 Parity Error
31	CAM2 GIO Parity Error
32	CAM2 PEO Parity Error
33	CAM2 Operation Error
34	CAM2 Result Bus Parity Error
35	CAM3 Word 0 Parity Error
36	CAM3 Word 1 Parity Error
37	CAM3 GIO Parity Error
38	CAM3 PEO Parity Error
39	CAM3 Operation Error
40	CAM3 Result Bus Parity Error
Internal memory errors	
1	Interlaken CRC32 Error on lane 0
2	Interlaken CRC32 Error on lane 1
3	Interlaken CRC32 Error on lane 2
4	Interlaken CRC32 Error on lane 3
5	Interlaken CRC32 Error on lane 4
6	Interlaken CRC32 Error on lane 5
7	Interlaken CRC32 Error on lane 6
8	Interlaken CRC32 Error on lane 7
9	Interlaken CRC32 Error on lane 8

TABLE 19 NP memory errors supported on BR-MLX-100GX2-CFP2, BR-MLX-10GX20, and BR-MLX-10GX4-IPSEC interface cards (continued)

Error	Description
10	Interlaken CRC32 Error on lane 9
11	Interlaken CRC32 Error on lane 10
12	Interlaken CRC32 Error on lane 11
13	Interlaken CRC32 Error on lane 12
14	Interlaken CRC32 Error on lane 13
15	Interlaken CRC32 Error on lane 14
16	Interlaken CRC32 Error on lane 15
17	Interlaken CRC32 Error on lane 16
18	Interlaken CRC32 Error on lane 17
19	Interlaken CRC32 Error on lane 18
20	Interlaken CRC32 Error on lane 19
21	Interlaken CRC32 Error on lane 20
22	Interlaken CRC32 Error on lane 21
23	Interlaken CRC32 Error on lane 22
24	Interlaken CRC32 Error on lane 23
25	Interlaken CRC24 Error
26	Interlaken RG Overflow
27	Interlaken Core RDC Overflow
28	Interlaken Core Control FIFO Overflow
29	Interlaken Core Tx Underflow
30	Interlaken Core Tx Overflow
31	Interlaken Sync FIFO Rx Overflow
32	Interlaken Flow Control DIP Error
33	Tx Deframer MVLAN Flag FIFO Parity Error
34	Tx Deframer MVLAN control Packet FIFO Parity Error
35	Tx Deframer MVLAN replacement table Parity Error
36	Tx Deframer MVLAN start offset FIFO Parity Error
37	Tx Deframer MVLAN sop FIFO Parity Error
38	Tx Deframer MVLAN payload Data FIFO Parity Error
39	Tx Packet Edit Data FIFO Parity Error
40	Tx Packet Edit Next Hop Table Parity Error
41	ACL Data FIFO Parity Error
42	ACL Control FIFO Parity Error
43	ACL QoS Done FIFO Parity Error
44	ACL Port Number FIFO Parity Error
45	ACL Priority Encode Table Parity Error
46	ACL Tx VLAN Table Parity Error
47	Tx Priority Encode Table Lookup Result Parity Error
48	MAC2 Frame LSTD Parity Error
49	MAC2 Frame Data Parity Error

TABLE 19 NP memory errors supported on BR-MLX-100GX2-CFP2, BR-MLX-10GX20, and BR-MLX-10GX4-IPSEC interface cards (continued)

Error	Description
50	MAC2 Frame Control Parity Error
51	MAC1 Frame LSTD Parity Error
52	MAC1 Frame Data Parity Error
53	MAC1 Frame Control Parity Error
54	MAC0 Frame LSTD Parity Error
55	MAC0 Frame Data Parity Error
56	MAC0 Frame Control Parity Error
57	Tx Packet Edit Data FIFO Parity Error
58	Tx Packet Edit Control FIFO Parity Error
59	Tx Packet Edit Nhlk FIFO Parity Error
60	Tx Packet Edit Pipe LBLe FIFO Parity Error
61	Start Offset Table CPU Read Parity Error
62	Replacement Table CPU Read Parity Error
63	Next Hop Table CPU Read Parity Error
64	Tx VLAN Table CPU Read Parity Error
65	Priority Encode Table CPU Read Parity Error
66	MAC9 Frame LSTD Parity Error
67	MAC9 Frame Data Parity Error
68	MAC9 Frame Control Parity Error
69	MAC8 Frame LSTD Parity Error
70	MAC8 Frame Data Parity Error
71	MAC8 Frame Control Parity Error
72	MAC7 Frame LSTD Parity Error
73	MAC7 Frame Data Parity Error
74	MAC7 Frame Control Parity Error
75	MAC6 Frame LSTD Parity Error
76	MAC6 Frame Data Parity Error
77	MAC6 Frame Control Parity Error
78	MAC5 Frame LSTD Parity Error
79	MAC5 Frame Data Parity Error
80	MAC5 Frame Control Parity Error
81	MAC4 Frame LSTD Parity Error
82	MAC4 Frame Data Parity Error
83	MAC4 Frame Control Parity Error
84	MAC3 Frame LSTD Parity Error
85	MAC3 Frame Data Parity Error
86	MAC3 Frame Control Parity Error
87	Rx MAC0 Data FIFO Read Parity Error
88	Rx MAC0 Flag FIFO Read Parity Error
89	Rx MAC1 Data FIFO Read Parity Error

TABLE 19 NP memory errors supported on BR-MLX-100GX2-CFP2, BR-MLX-10GX20, and BR-MLX-10GX4-IPSEC interface cards (continued)

Error	Description
90	Rx MAC1 Flag FIFO Read Parity Error
91	CAM Result Scheduler FIFO Overflow
92	Rx Port Pipeline HQoS Data Parity Error
93	Rx Port Pipeline Rx Data-in Parity Error
94	Rx Port Pipeline Rxctrl FIFO Read Data Parity Error
95	Rx Port Pipeline Read Rx QoS Id FIFO Parity Error
96	Rx Port Pipeline Rx portnum FIFO Parity Error
97	Rx Port Pipeline Rx QoS Done FIFO Parity Error
98	Rx Port Pipeline Rx Flag FIFO Parity Error
99	Rx Port Pipeline Rx Header FIFO Parity Error
100	Rx Port Pipeline PRAM Packet Id Mismatch
101	Rx Port Pipeline Data Path Packet Id Mismatch
102	EXM IP Address FIFO Overflow
103	CAM Result FIFO Parity Error
104	CAM Result FIFO Underflow
105	ECMP FIFO Underflow
106	LBL FIFO Underflow
107	Invalid CAM Result
108	CAM2PRAM CAM Intf Data FIFO Overflow
109	CAM2PRAM CAM Intf Count FIFO Overflow
110	LBLRAM SRVT0 Lookup FIFO Overflow
111	LBLRAM SRVT1 Lookup FIFO Overflow
112	Rx Service PRAM Result FIFO Parity Error
113	Rx Packet Header FIFO MISC RAM Parity Error
114	Rx Packet Header FIFO RAM3 Parity Error
115	Rx Packet Header FIFO RAM2 Parity Error
116	Rx Packet Header FIFO RAM1 Parity Error
117	CAM1 Lookup FIFO Valid Word FIFO Overflow
118	LBL Lookup FIFO Overflow
119	Packet Decode FIFO Parity Error
120	ECMP FIFO Overflow
121	LBL2 ECMP FIFO Overflow
122	Rx port Pipeline Rx QoS Id FIFO Overflow
123	Rx port Pipeline Rx Flag FIFO Overflow
124	Rx port Pipeline Rx Header FIFO Overflow
125	Rx CAM Result FIFO Parity Error
126	Rx CAM Result FIFO Underflow
127	Rx CAM Result FIFO Overflow
128	Rx Packet Decode FIFO Overflow
129	Rx Packet Decode FIFO Underflow

TABLE 19 NP memory errors supported on BR-MLX-100GX2-CFP2, BR-MLX-10GX20, and BR-MLX-10GX4-IPSEC interface cards (continued)

Error	Description
130	Rx topotos FIFO Parity Error
131	Rx topotos FIFO Underflow
132	Rx topotos FIFO Overflow
133	CAM1 Lookup Misc FIFO Overflow
134	CAM1 Lookup Data FIFO Overflow
135	CAM SCI FIFO Parity Error
136	CAM1 Response Packet FIFO1 Underflow
137	CAM1 Response Packet FIFO1 Overflow
138	CAM1 Response Packet FIFO2 Underflow
139	CAM1 Response Packet FIFO2 Overflow
140	CAM1 Response Packet FIFO3 Underflow
141	CAM1 Response Packet FIFO3 Overflow
142	CAM1 Response Packet FIFO4 Underflow
143	CAM1 Response Packet FIFO4 Overflow
144	CAM SCI FIFO Underflow
145	CAM SCI FIFO Overflow
146	CAM1 Lookup Misc FIFO Underflow
147	CAM1 Lookup Data FIFO Underflow
148	CAM1 Lookup Misc FIFO Parity Error
149	CAM1 Lookup Data FIFO Parity Error
150	CAM ILA Core FIFO Underflow
151	CAM ILA Core FIFO Overflow
152	EXM CPU2HashBucket rData Parity Error
153	EXM HW Search Hash Bucket rData Parity Error
154	EXM Hash Index Table Parity Error
155	EXM IP Address FIFO Underflow
156	EXM IP Address FIFO Parity Error
157	EXM VPN Id FIFO Underflow
158	EXM VPN Id FIFO Parity Error
159	CAM1 Lookup Misc FIFO Underflow
160	CAM1 Lookup Data FIFO Underflow
161	CAM2PRAM QDR Interface FIFO3 Parity Error
162	CAM2PRAM QDR Interface FIFO2 Parity Error
163	CAM2PRAM QDR Interface FIFO1 Parity Error
164	CAM2PRAM QDR Interface FIFO0 Parity Error
165	CAM2PRAM QDR Interface FIFO3 Underflow
166	CAM2PRAM QDR Interface FIFO2 Underflow
167	CAM2PRAM QDR Interface FIFO1 Underflow
168	CAM2PRAM QDR Interface FIFO0 Underflow
169	CAM2PRAM QDR Interface FIFO3 Overflow

TABLE 19 NP memory errors supported on BR-MLX-100GX2-CFP2, BR-MLX-10GX20, and BR-MLX-10GX4-IPSEC interface cards (continued)

Error	Description
170	CAM2PRAM QDR Interface FIFO2 Overflow
171	CAM2PRAM QDR Interface FIFO1 Overflow
172	CAM2PRAM QDR Interface FIFO0 Overflow
173	CAM2PRAM QDR Interface Read Request FIFO3 Parity Error
174	CAM2PRAM QDR Interface Read Request FIFO2 Parity Error
175	CAM2PRAM QDR Interface Read Request FIFO1 Parity Error
176	CAM2PRAM QDR Interface Read Request FIFO0 Parity Error
177	CAM2PRAM QDR Interface Read Request FIFO3 Underflow
178	CAM2PRAM QDR Interface Read Request FIFO2 Underflow
179	CAM2PRAM QDR Interface Read Request FIFO1 Underflow
180	CAM2PRAM QDR Interface Read Request FIFO0 Underflow
181	CAM2PRAM QDR Interface Read Request FIFO3 Overflow
182	CAM2PRAM QDR Interface Read Request FIFO2 Overflow
183	CAM2PRAM QDR Interface Read Request FIFO1 Overflow
184	CAM2PRAM QDR Interface Read Request FIFO0 Overflow
185	CAM2PRAM cpu FIFO Parity Error
186	CAM2PRAM CAM Interface Data FIFO Underflow
187	CAM2PRAM CAM Interface Count FIFO Underflow
188	CAM2PRAM Result FIFO0 Parity Error
189	CAM2PRAM Result FIFO1 Parity Error
190	CAM2PRAM Result FIFO2 Parity Error
191	CAM2PRAM Result FIFO3 Parity Error
192	CAM2PRAM Result FIFO4 Parity Error
193	CAM2PRAM Result FIFO5 Parity Error
194	CAM2PRAM Result FIFO0 Underflow
195	CAM2PRAM Result FIFO1 Underflow
196	CAM2PRAM Result FIFO2 Underflow
197	CAM2PRAM Result FIFO3 Underflow
198	CAM2PRAM Result FIFO4 Underflow
199	CAM2PRAM Result FIFO5 Underflow
200	CAM2PRAM Result FIFO0 Overflow
201	CAM2PRAM Result FIFO1 Overflow
202	CAM2PRAM Result FIFO2 Overflow
203	CAM2PRAM Result FIFO3 Overflow
204	CAM2PRAM Result FIFO4 Overflow
205	CAM2PRAM Result FIFO5 Overflow
206	CAM2PRAM Result Scheduler Underflow
207	CAM2PRAM Result Scheduler Overflow
208	PRAM CAM Interface Data FIFO0 Underflow
209	PRAM CAM Interface Data FIFO0 Overflow

TABLE 19 NP memory errors supported on BR-MLX-100GX2-CFP2, BR-MLX-10GX20, and BR-MLX-10GX4-IPSEC interface cards (continued)

Error	Description
210	PRAM CAM Interface Data FIFO1 Underflow
211	PRAM CAM Interface Data FIFO1 Overflow
212	PRAM CAM Interface Data FIFO2 Underflow
213	PRAM CAM Interface Data FIFO2 Overflow
214	PRAM CAM Interface Data FIFO3 Underflow
215	PRAM CAM Interface Data FIFO3 Overflow
216	PRAM Channel0 Data0 Multi bit Parity Error
217	PRAM Channel0 Data1 Multi bit Parity Error
218	PRAM Channel0 Data2 Multi bit Parity Error
219	PRAM Channel0 Data3 Multi bit Parity Error
220	PRAM Channel0 Data0 Single bit Parity Error
221	PRAM Channel0 Data1 Single bit Parity Error
222	PRAM Channel0 Data2 Single bit Parity Error
223	PRAM Channel0 Data3 Single bit Parity Error
224	PRAM Channel1 Data0 Multi bit Parity Error
225	PRAM Channel1 Data1 Multi bit Parity Error
226	PRAM Channel1 Data2 Multi bit Parity Error
227	PRAM Channel1 Data3 Multi bit Parity Error
228	PRAM Channel1 Data0 Single bit Parity Error
229	PRAM Channel1 Data1 Single bit Parity Error
230	PRAM Channel1 Data2 Single bit Parity Error
231	PRAM Channel1 Data3 Single bit Parity Error
232	PRAM Channel2 Data0 Multi bit Parity Error
233	PRAM Channel2 Data1 Multi bit Parity Error
234	PRAM Channel2 Data2 Multi bit Parity Error
235	PRAM Channel2 Data3 Multi bit Parity Error
236	PRAM Channel2 Data0 Single bit Parity Error
237	PRAM Channel2 Data1 Single bit Parity Error
238	PRAM Channel2 Data2 Single bit Parity Error
239	PRAM Channel2 Data3 Single bit Parity Error
240	PRAM Channel3 Data0 Multi bit Parity Error
241	PRAM Channel3 Data1 Multi bit Parity Error
242	PRAM Channel3 Data2 Multi bit Parity Error
243	PRAM Channel3 Data3 Multi bit Parity Error
244	PRAM Channel3 Data0 Single bit Parity Error
245	PRAM Channel3 Data1 Single bit Parity Error
246	PRAM Channel3 Data2 Single bit Parity Error
247	PRAM Channel3 Data3 Single bit Parity Error
248	CAM2Age L2 FIFO0 Underflow
249	CAM2Age L2 FIFO1 Underflow

TABLE 19 NP memory errors supported on BR-MLX-100GX2-CFP2, BR-MLX-10GX20, and BR-MLX-10GX4-IPSEC interface cards (continued)

Error	Description
250	CAM2Age ACL FIFO0 Underflow
251	CAM2Age ACL FIFO1 Underflow
252	CAM2Age L3 Underflow
253	L2 Aged FIFO Underflow
254	L2 Aged FIFO Overflow
255	L2 Aged FIFO Parity Error
256	L2 Aged mem Parity Error
257	L3 Aged FIFO Underflow
258	L3 Aged FIFO Overflow
259	L3 Aged FIFO Parity Error
260	L3 Aged Mem Parity Error
261	ACL Aged FIFO Underflow
262	ACL Aged FIFO Overflow
263	ACL Aged FIFO Parity Error
264	ACL Aged mem Parity Error
265	Rx QoS Id FIFO Underflow
266	Rx QoS Id FIFO Overflow
267	Rx Flag FIFO Underflow
268	Rx QoS Done FIFO Underflow
269	Rx QoS Done FIFO Overflow
270	Rx FID FIFO Underflow
271	CAM Lookup Misc FIFO Underflow
272	CAM Lookup Data FIFO Underflow
273	PRAM QDR Interface Read Request FIFO0 Parity Error
274	PRAM QDR Interface Read Request FIFO1 Parity Error
275	PRAM QDR Interface Read Request FIFO2 Parity Error
276	PRAM QDR Interface Read Request FIFO3 Parity Error
277	PRAM QDR Interface Read Request FIFO0 Underflow
278	PRAM QDR Interface Read Request FIFO1 Underflow
279	PRAM QDR Interface Read Request FIFO2 Underflow
280	PRAM QDR Interface Read Request FIFO3 Underflow
281	PRAM QDR Interface Read Request FIFO0 Overflow
282	PRAM QDR Interface Read Request FIFO1 Overflow
283	PRAM QDR Interface Read Request FIFO2 Overflow
284	PRAM QDR Interface Read Request FIFO3 Overflow
285	PRAM QDR Interface CPU Read FIFO Parity Error
286	PRAM Result FIFO0 Parity Error
287	PRAM Result FIFO1 Parity Error
288	PRAM Result FIFO2 Parity Error
289	PRAM Result FIFO3 Parity Error

TABLE 19 NP memory errors supported on BR-MLX-100GX2-CFP2, BR-MLX-10GX20, and BR-MLX-10GX4-IPSEC interface cards (continued)

Error	Description
290	PRAM Result FIFO4 Parity Error
291	PRAM Result FIFO5 Parity Error
292	PRAM Result FIFO0 Overflow
293	PRAM Result FIFO1 Overflow
294	PRAM Result FIFO2 Overflow
295	PRAM Result FIFO3 Overflow
296	PRAM Result FIFO4 Overflow
297	PRAM Result FIFO5 Overflow
298	CAM2Age L3 Overflow
299	CAM2Age L2 FIFO0 Overflow
300	CAM2Age L2 FIFO1 Overflow
301	CAM2Age ACL FIFO0 Overflow
302	CAM2Age ACL FIFO1 Overflow
303	CAM2PRAM Data FIFO Overflow
304	CAM2PRAM Count FIFO Overflow
305	CAM3 SCI FIFO Parity Error
306	CAM3 Response Packet FIFO1 Underflow
307	CAM3 Response Packet FIFO1 Overflow
308	CAM3 Response Packet FIFO2 Underflow
309	CAM3 Response Packet FIFO2 Overflow
310	CAM3 Response Packet FIFO3 Underflow
311	EXM VPN ID FIFO Overflow
312	CAM3 Response Packet FIFO3 Overflow
313	CAM3 Response Packet FIFO4 Underflow
314	CAM3 Response Packet FIFO4 Overflow
315	CAM3 SCI FIFO Underflow
316	CAM3 SCI FIFO Overflow
317	CAM3 Lookup Misc FIFO Underflow
318	CAM3 Lookup Data FIFO Underflow
319	CAM3 Lookup Misc FIFO Parity Error
320	CAM3 Lookup Data FIFO Parity Error
321	CAM3 ILA Core FIFO Underflow
322	CAM3 ILA Core FIFO Overflow
323	CAM3 ILA Core FIFO Parity Error
324	PCIe Cmd Read Data Parity Error
325	PCIe Rx DRAM FIFO Underflow
326	PCIe Rx Frame FIFO Underflow
327	CAM3 Lookup Misc FIFO Overflow
328	CAM3 Lookup Data FIFO Overflow
329	Service CAM Lookup Control FIFO Overflow

TABLE 19 NP memory errors supported on BR-MLX-100GX2-CFP2, BR-MLX-10GX20, and BR-MLX-10GX4-IPSEC interface cards (continued)

Error	Description
330	Service CAM Lookup Control FIFO Underflow
331	Service CAM Lookup FIFO Overflow
332	Service CAM Lookup FIFO Underflow
333	eACL CAM Lookup FIFO0 Underflow
334	eACL CAM Lookup Control FIFO0 Underflow
335	eACL CAM Lookup FIFO1 Underflow
336	eACL CAM Lookup Control FIFO1 Underflow
337	CAM3 Result FIFO Underflow
338	CAM3 Result Data FIFO Underflow
339	CAM3 Result FIFO Parity Error
340	CAM3 Result Data FIFO Parity Error
341	CAM Result FIFO Overflow
342	CAM Result FIFO Underflow

TABLE 20 NP memory errors supported on BR-MLX-40Gx4-X interface cards

Error	Description
External memory errors	
1	PRAM Word 0 Parity Error
2	PRAM Word 1 Parity Error
3	PRAM Word 2 Parity Error
4	PRAM Word 3 Parity Error
5	PRAM Word 4 Parity Error
6	PRAM Word 5 Parity Error
7	PRAM Word 6 Parity Error
8	PRAM Word 7 Parity Error
9	CAM2PRAM Word 0 Single Bit Parity Error
10	CAM2PRAM Word 1 Single Bit Parity Error
11	CAM2PRAM Word 2 Single Bit Parity Error
12	CAM2PRAM Word 3 Single Bit Parity Error
13	CAM2PRAM Word 0 Double Bit Parity Error
14	CAM2PRAM Word 1 Double Bit Parity Error
15	CAM2PRAM Word 2 Double Bit Parity Error
16	CAM2PRAM Word 3 Double Bit Parity Error
17	LBLRAM Word 0 Parity Error
18	LBLRAM Word 1 Parity Error
19	CAM1 Word 0 Parity Error
20	CAM1 Word 1 Parity Error
21	CAM1 GIO Parity Error
22	CAM1 PEO Parity Error
23	CAM1 Operation Error

TABLE 20 NP memory errors supported on BR-MLX-40Gx4-X interface cards (continued)

Error	Description
24	CAM1 Result Bus Parity Error
25	CAM2 Word 0 Parity Error
26	CAM2 Word 1 Parity Error
27	CAM2 GIO Parity Error
28	CAM2 PEO Parity Error
29	CAM2 Operation Error
30	CAM2 Result Bus Parity Error
31	CAM3 Word 0 Parity Error
32	CAM3 Word 1 Parity Error
33	CAM3 GIO Parity Error
34	CAM3 PEO Parity Error
35	CAM3 Operation Error
36	CAM3 Result Bus Parity Error
Internal memory errors	
1	Tx Deframer MVLAN Flag FIFO Parity Error
2	Tx Deframer MVLAN control Packet FIFO Parity Error
3	Tx Deframer MVLAN replication table Parity Error
4	Tx Deframer MVLAN start offset FIFO Parity Error
5	Tx Deframer MVLAN sop FIFO Parity Error
6	Tx Deframer MVLAN payload Data FIFO Parity Error
7	Tx Packet Edit Data FIFO Parity Error
8	Tx Packet Edit Next Hop Table Parity Error
9	ACL PRAM Results FIFO Parity Error
10	ACL Data FIFO Parity Error
11	ACL Control FIFO Parity Error
12	ACL QoS Done FIFO Parity Error
13	ACL Port Number FIFO Parity Error
14	ACL Priority Encode Table Parity Error
15	ACL Tx VLAN Table Parity Error
16	Tx Priority Encode Table Lookup Result Parity Error
17	MAC0 Frame LSTD Parity Error
18	MAC0 Frame Data Parity Error
19	MAC0 Frame Control Parity Error
20	MAC1 Frame LSTD Parity Error
21	MAC1 Frame Data Parity Error
22	MAC1 Frame Control Parity Error
23	Tx Packet Edit Data FIFO Parity Error
24	Tx Packet Edit Control FIFO Parity Error
25	Tx Packet Edit nhlk FIFO Parity Error
26	Tx Packet Edit pipe LBLe FIFO Parity Error
27	Start Offset Table CPU Read Parity Error

TABLE 20 NP memory errors supported on BR-MLX-40Gx4-X interface cards (continued)

Error	Description
28	Replacement Table CPU Read Parity Error
29	Next Hop Table CPU Read Parity Error
30	Tx VLAN Table CPU Read Parity Error
31	Priority Encode Table CPU Read Parity Error
32	Rx MAC Data FIFO Read Parity Error
33	Rx MAC Flag FIFO Read Parity Error
34	Rx MAC Data FIFO Read Parity Error
35	Rx MAC Flag FIFO Read Parity Error
36	CAM Result Scheduler FIFO Overflow
37	CAM1 Lookup FIFO3 Overflow
38	CAM1 Lookup FIFO2 Overflow
39	CAM1 Lookup FIFO1 Overflow
40	CAM2 Lookup FIFO3 Overflow
41	CAM2 Lookup FIFO2 Overflow
42	CAM2 Lookup FIFO1 Overflow
43	Rx Port Pipeline HQoS Data Parity Error
44	Rx Port Pipeline Rx Data-in Parity Error
45	Rx Port Pipeline Rxctrl FIFO Read Data Parity Error
46	Rx Port Pipeline Read Rx QoS Id FIFO Parity Error
47	Rx Port Pipeline Rx portnum FIFO Parity Error
48	Rx Port Pipeline Rx QoS Done FIFO Parity Error
49	Rx Port Pipeline Rx Flag FIFO Parity Error
50	Rx Port Pipeline Rx Header FIFO Parity Error
51	Rx Port Pipeline PRAM Packet Id Mismatch
52	Rx Port Pipeline Data Path Packet Id Mismatch
53	EXM IP Address FIFO Overflow
54	Rx Service PRAM Result FIFO Parity Error
55	Rx Packet Header FIFO MISC RAM Parity Error
56	Rx Packet Header FIFO RAM3 Parity Error
57	Rx Packet Header FIFO RAM2 Parity Error
58	Rx Packet Header FIFO RAM1 Parity Error
59	Rx Packet Header FIFO Parity Error
60	LBL Lookup FIFO Overflow
61	Packet Decode FIFO Parity Error
62	ECMP FIFO Overflow
63	LBL2 ECMP FIFO Overflow
64	Rx port Pipeline Rx QoS Id FIFO Overflow
65	Rx port Pipeline Rx Flag FIFO Overflow
66	Rx port Pipeline Rx Header FIFO Overflow
67	Rx CAM Result FIFO Parity Error
68	Rx CAM Result FIFO Underflow

TABLE 20 NP memory errors supported on BR-MLX-40Gx4-X interface cards (continued)

Error	Description
69	Rx CAM Result FIFO Overflow
70	Rx Packet Decode FIFO Overflow
71	Rx Packet Decode FIFO Underflow
72	Rx topotos FIFO Parity Error
73	Rx topotos FIFO Underflow
74	Rx topotos FIFO Overflow
75	CAM1 Lookup FIFO 1 Parity Error
76	CAM1 Lookup FIFO 2 Parity Error
77	CAM1 Lookup FIFO 3 Parity Error
78	CAM1 Asc FIFO Parity Error
79	CAM2 Lookup FIFO 1 Parity Error
80	CAM2 Lookup FIFO 2 Parity Error
81	CAM2 Lookup FIFO 3 Parity Error
82	CAM2 Asc FIFO Parity Error
83	LBLRAM srvt Lookup FIFO Underflow
84	LBLRAM extd Service rd Data Parity Error
85	LBLRAM LBL Lookup FIFO Underflow
86	LBLRAM srvp Lookup FIFO Overflow
87	LBLRAM srvp Lookup FIFO Underflow
88	LBLRAM Read Request FIFO Overflow
89	LBLRAM Read Request FIFO Underflow
90	LBLRAM extd Service Read FIFO Overflow
91	LBLRAM extd Service Read FIFO Underflow
92	Service Table Lookup FIFO rd Data Parity Error
93	CAM Result Scheduler FIFO Underflow
94	CAM1 Result FIFO Overflow
95	CAM1 Result FIFO Underflow
96	CAM2 Result FIFO Overflow
97	CAM2 Result FIFO Underflow
98	Label Result FIFO Overflow
99	Label Result FIFO Underflow
100	LBL hold FIFO Underflow
101	LBL hold FIFO Overflow
102	ECMP FIFO Underflow
103	LBL Result sync FIFO Underflow
104	LBL Result sync FIFO Overflow
105	LBL2ECMP FIFO Underflow
106	EXM CPU2HashBucket rData Parity Error
107	EXM HW Search Hash Bucket rData Parity Error
108	EXM Hash Index Table Parity Error
109	EXM IP Address FIFO Underflow

TABLE 20 NP memory errors supported on BR-MLX-40Gx4-X interface cards (continued)

Error	Description
110	EXM IP Address FIFO Parity Error
111	EXM VPN Id FIFO Underflow
112	EXM VPN Id FIFO Parity Error
113	CAM2PRAM QDR Interface FIFO3 Parity Error
114	CAM2PRAM QDR Interface FIFO2 Parity Error
115	CAM2PRAM QDR Interface FIFO1 Parity Error
116	CAM2PRAM QDR Interface FIFO0 Parity Error
117	CAM2PRAM QDR Interface FIFO3 Underflow
118	CAM2PRAM QDR Interface FIFO2 Underflow
119	CAM2PRAM QDR Interface FIFO1 Underflow
120	CAM2PRAM QDR Interface FIFO0 Underflow
121	CAM2PRAM QDR Interface FIFO3 Overflow
122	CAM2PRAM QDR Interface FIFO2 Overflow
123	CAM2PRAM QDR Interface FIFO1 Overflow
124	CAM2PRAM QDR Interface FIFO0 Overflow
125	CAM2PRAM QDR Interface Read Request FIFO3 Parity Error
126	CAM2PRAM QDR Interface Read Request FIFO2 Parity Error
127	CAM2PRAM QDR Interface Read Request FIFO1 Parity Error
128	CAM2PRAM QDR Interface Read Request FIFO0 Parity Error
129	CAM2PRAM QDR Interface Read Request FIFO3 Underflow
130	CAM2PRAM QDR Interface Read Request FIFO2 Underflow
131	CAM2PRAM QDR Interface Read Request FIFO1 Underflow
132	CAM2PRAM QDR Interface Read Request FIFO0 Underflow
133	CAM2PRAM QDR Interface Read Request FIFO3 Overflow
134	CAM2PRAM QDR Interface Read Request FIFO2 Overflow
135	CAM2PRAM QDR Interface Read Request FIFO1 Overflow
136	CAM2PRAM QDR Interface Read Request FIFO0 Overflow
137	CAM2PRAM cpu FIFO Parity Error
138	CAM2PRAM CAM Interface Data FIFO Underflow
139	CAM2PRAM CAM Interface Count FIFO Underflow
140	CAM2PRAM Result FIFO0 Parity Error
141	CAM2PRAM Result FIFO1 Parity Error
142	CAM2PRAM Result FIFO2 Parity Error
143	CAM2PRAM Result FIFO3 Parity Error
144	CAM2PRAM Result FIFO4 Parity Error
145	CAM2PRAM Result FIFO5 Parity Error
146	CAM2PRAM Result FIFO0 Underflow
147	CAM2PRAM Result FIFO1 Underflow
148	CAM2PRAM Result FIFO2 Underflow
149	CAM2PRAM Result FIFO3 Underflow
150	CAM2PRAM Result FIFO4 Underflow

TABLE 20 NP memory errors supported on BR-MLX-40Gx4-X interface cards (continued)

Error	Description
151	CAM2PRAM Result FIFO5 Underflow
152	CAM2PRAM Result FIFO0 Overflow
153	CAM2PRAM Result FIFO1 Overflow
154	CAM2PRAM Result FIFO2 Overflow
155	CAM2PRAM Result FIFO3 Overflow
156	CAM2PRAM Result FIFO4 Overflow
157	CAM2PRAM Result FIFO5 Overflow
158	CAM2PRAM Result Scheduler Underflow
159	CAM2PRAM Result Scheduler Overflow
160	PRAM CAM Interface Data FIFO0 Underflow
161	PRAM CAM Interface Data FIFO0 Overflow
162	PRAM CAM Interface Data FIFO1 Underflow
163	PRAM CAM Interface Data FIFO1 Overflow
164	PRAM CAM Interface Data FIFO2 Underflow
165	PRAM CAM Interface Data FIFO2 Overflow
166	PRAM CAM Interface Data FIFO3 Underflow
167	PRAM CAM Interface Data FIFO3 Overflow
168	PRAM Channel0 Data0 Multi bit Parity Error
169	PRAM Channel0 Data1 Multi bit Parity Error
170	PRAM Channel0 Data2 Multi bit Parity Error
171	PRAM Channel0 Data3 Multi bit Parity Error
172	PRAM Channel0 Data0 Single bit Parity Error
173	PRAM Channel0 Data1 Single bit Parity Error
174	PRAM Channel0 Data2 Single bit Parity Error
175	PRAM Channel0 Data3 Single bit Parity Error
176	PRAM Channel1 Data0 Multi bit Parity Error
177	PRAM Channel1 Data1 Multi bit Parity Error
178	PRAM Channel1 Data2 Multi bit Parity Error
179	PRAM Channel1 Data3 Multi bit Parity Error
180	PRAM Channel1 Data0 Single bit Parity Error
181	PRAM Channel1 Data1 Single bit Parity Error
182	PRAM Channel1 Data2 Single bit Parity Error
183	PRAM Channel1 Data3 Single bit Parity Error
184	PRAM Channel2 Data0 Multi bit Parity Error
185	PRAM Channel2 Data1 Multi bit Parity Error
186	PRAM Channel2 Data2 Multi bit Parity Error
187	PRAM Channel2 Data3 Multi bit Parity Error
188	PRAM Channel2 Data0 Single bit Parity Error
189	PRAM Channel2 Data1 Single bit Parity Error
190	PRAM Channel2 Data2 Single bit Parity Error
191	PRAM Channel2 Data3 Single bit Parity Error

TABLE 20 NP memory errors supported on BR-MLX-40Gx4-X interface cards (continued)

Error	Description
192	PRAM Channel3 Data0 Multi bit Parity Error
193	PRAM Channel3 Data1 Multi bit Parity Error
194	PRAM Channel3 Data2 Multi bit Parity Error
195	PRAM Channel3 Data3 Multi bit Parity Error
196	PRAM Channel3 Data0 Single bit Parity Error
197	PRAM Channel3 Data1 Single bit Parity Error
198	PRAM Channel3 Data2 Single bit Parity Error
199	PRAM Channel3 Data3 Single bit Parity Error
200	CAM2Age L2 FIFO0 Underflow
201	CAM2Age L2 FIFO1 Underflow
202	CAM2Age ACL FIFO0 Underflow
203	CAM2Age ACL FIFO1 Underflow
204	CAM2Age L3 Underflow
205	L2 Aged FIFO Underflow
206	L2 Aged FIFO Overflow
207	L2 Aged FIFO Parity Error
208	L2 Aged mem Parity Error
209	L3 Aged FIFO Underflow
210	L3 Aged FIFO Overflow
211	L3 Aged FIFO Parity Error
212	L3 Aged Mem Parity Error
213	ACL Aged FIFO Underflow
214	ACL Aged FIFO Overflow
215	ACL Aged FIFO Parity Error
216	ACL Aged mem Parity Error
217	Rx QoS Id FIFO Underflow
218	Rx QoS Id FIFO Overflow
219	Rx Flag FIFO Underflow
220	Rx QoS Done FIFO Underflow
221	Rx QoS Done FIFO Overflow
222	PRAM QDR Interface Read Request FIFO0 Parity Error
223	PRAM QDR Interface Read Request FIFO1 Parity Error
224	PRAM QDR Interface Read Request FIFO2 Parity Error
225	PRAM QDR Interface Read Request FIFO3 Parity Error
226	PRAM QDR Interface Read Request FIFO0 Underflow
227	PRAM QDR Interface Read Request FIFO1 Underflow
228	PRAM QDR Interface Read Request FIFO2 Underflow
229	PRAM QDR Interface Read Request FIFO3 Underflow
230	PRAM QDR Interface Read Request FIFO0 Overflow
231	PRAM QDR Interface Read Request FIFO1 Overflow
232	PRAM QDR Interface Read Request FIFO2 Overflow

TABLE 20 NP memory errors supported on BR-MLX-40Gx4-X interface cards (continued)

Error	Description
233	PRAM QDR Interface Read Request FIFO3 Overflow
234	PRAM QDR Interface CPU Read FIFO Parity Error
235	PRAM Result FIFO0 Parity Error
236	PRAM Result FIFO1 Parity Error
237	PRAM Result FIFO2 Parity Error
238	PRAM Result FIFO3 Parity Error
239	PRAM Result FIFO4 Parity Error
240	PRAM Result FIFO5 Parity Error
241	PRAM Result FIFO0 Overflow
242	PRAM Result FIFO1 Overflow
243	PRAM Result FIFO2 Overflow
244	PRAM Result FIFO3 Overflow
245	PRAM Result FIFO4 Overflow
246	PRAM Result FIFO5 Overflow
247	CAM2Age L3 Overflow
248	CAM2Age L2 FIFO0 Overflow
249	CAM2Age L2 FIFO1 Overflow
250	CAM2Age ACL FIFO0 Overflow
251	CAM2Age ACL FIFO1 Overflow
252	CAM2PRAM Data FIFO Overflow
253	CAM2PRAM Count FIFO Overflow
254	CAM1 Lookup FIFO1 Underflow
255	CAM1 Lookup FIFO2 Underflow
256	CAM1 Lookup FIFO3 Underflow
257	CAM2 Lookup FIFO1 Underflow
258	CAM2 Lookup FIFO2 Underflow
259	CAM2 Lookup FIFO3 Underflow
260	CAM1 Asc FIFO Underflow
261	CAM1 Asc FIFO Overflow
262	CAM2 Asc FIFO Underflow
263	CAM2 Asc FIFO Overflow
264	EXM VPN Id FIFO Overflow
265	CAM3 Lookup FIFO1 Underflow
266	CAM3 Lookup FIFO2 Underflow
267	CAM3 Lookup FIFO3 Underflow
268	CAM3 Asc FIFO Underflow
269	CAM3 Asc FIFO Overflow
270	CAM3 Lookup FIFO1 Parity Error
271	CAM3 Lookup FIFO2 Parity Error
272	CAM3 Lookup FIFO3 Parity Error
273	CAM3 Asc FIFO Parity Error

TABLE 20 NP memory errors supported on BR-MLX-40Gx4-X interface cards (continued)

Error	Description
274	PCIe Cmd Read Data Parity Error
275	PCIe Rx DRAM FIFO Underflow
276	PCIe Rx Frame FIFO Underflow
277	CAM3 Lookup FIFO1 Overflow
278	CAM3 Lookup FIFO2 Overflow
279	CAM3 Lookup FIFO3 Overflow
280	Service CAM Block Mux FIFO Overflow
281	Service CAM Block Mux FIFO Underflow
282	eACL CAM Block Mux FIFO1 Overflow
283	eACL CAM Block Mux FIFO1 Underflow
284	eACL CAM Block Mux FIFO2 Overflow
285	eACL CAM Block Mux FIFO2 Underflow
286	Service CAM Result FIFO Overflow
287	Service CAM Result FIFO Underflow
288	eACL CAM Result FIFO Overflow
289	eACL CAM Result FIFO Underflow
290	eACL CAM Block Mux FIFO2 Parity Error
291	eACL CAM Block Mux FIFO1 Parity Error
292	Service CAM Block Mux FIFO Parity Error
293	CAM3 Service Result FIFO Parity Error
294	CAM3 eACL Result FIFO Parity Error

TABLE 21 NP memory errors supported on BR-MLX-10Gx24 interface cards

Error	Description
External memory errors	
1	ISL TCAM Parity Error
2	ISL PRAM ECC Single Error
3	ISL BKT Memory ECC Single Error
4	ISL Mask Memory ECC Single Error
5	ISL Index Memory ECC Single Error
6	ISL PRAM ECC Double Error
7	ISL BKT Memory ECC Double Error
8	ISL Mask Memory ECC Double Error
9	ISL Index Memory ECC Double Error
Internal memory errors	
	IPT
1	IPT Topology Table Memory Un-Correctable ECC Error
2	IPT Topology Table Memory Correctable ECC Error
3	IPT DSCP Table Memory Un-Correctable ECC Error
4	IPT DSCP Table Memory Correctable ECC Error
5	IPT PCP Table Memory Un-Correctable ECC Error

TABLE 21 NP memory errors supported on BR-MLX-10Gx24 interface cards (continued)

Error	Description
6	IPT PCP Table Memory Correctable ECC Error
7	IPT EXP Table Memory Un-Correctable ECC Error
8	IPT EXP Table Memory Correctable ECC Error
9	IPT Byte count Table Memory Parity Error
10	IPT Framee count Table Memory Parity Error
	EFE
11	EFE Frame Control Parity Error
12	EFE Frame Data Parity Error
13	EFE HW NextHop Table Lookup 1bit Parity Error
14	EFE HW NextHop Table Lookup 2bit Parity Error
	PIB
15	PIB Tx Channel 0 FIFO DMA Parity Error
16	PIB Tx Channel 1 FIFO DMA Parity Error
17	PIB Tx Channel 2 FIFO DMA Parity Error
18	PIB Tx Channel 3 FIFO DMA Parity Error
19	PIB Tx Keep Alive FIFO DMA Parity Error
20	PIB Tx Keep Alive Sequence Id Parity Error
21	PIB Tx Keep Alive Scheduler Parity Error
22	PIB Tx Keep Alive Descriptor Parity Error
23	PIB RA IO Read Parity Error
24	PIB RA IO Write Parity Error
25	PIB Tx Channel 0 PCIe Read Parity Error
26	PIB Tx Channel 1 PCIe Read Parity Error
27	PIB Tx Channel 2 PCIe Read Parity Error
28	PIB Tx Channel 3 PCIe Read Parity Error
29	PIB Keep Alive Tx Channel PCIe Read Parity Error
30	PIB lOR PCIe Read Parity Error
31	PIB lOW PCIe Read Parity Error
32	PIB lO PCIe Write Parity Error
	ICC
33	ICC ACL CAM Request FIFO Overflow
34	ICC L2 CAM Request FIFO Overflow
35	ICC L3 CAM Request FIFO Overflow
36	ICC ACL CAM result FIFO Overflow
37	ICC L2 CAM result FIFO Overflow
38	ICC L3 CAM result FIFO Overflow
39	ICC ACL CAM Request FIFO Parity Error
40	ICC L2 CAM Request FIFO Parity Error
41	ICC L3 CAM Request FIFO Parity Error
42	ICC ACL CAM result FIFO Parity Error
43	ICC L2 CAM result FIFO Parity Error

TABLE 21 NP memory errors supported on BR-MLX-10Gx24 interface cards (continued)

Error	Description
44	ICC L3 CAM result FIFO Parity Error
45	ICC L2 CAM result interface Parity Error
46	ICC L3 CAM result interface Parity Error
47	ICC L2 CAM result Pipeline fatal Errors
48	ICC L3 CAM result Pipeline fatal Errors
49	ICC passthrough FIFO (hash,labels) Overflow
50	ICC passthrough FIFO read Parity Error
51	ICC ACL CAM read Parity Error
	ERA
52	ERA RPC Credit update err
53	ERA Tx PRAM HW Read 1bit err
54	ERA Tx PRAM HW Read 2bit err
55	ERA Tx PRAM refresh err
56	ERA Tx VLAN Table HW Read 1bit err
57	ERA Tx VLAN Table HW Read 2bit err
58	ERA Tx VLAN Table refresh err
59	ERA RPC RMRAM ECC serr
60	ERA RPC RMRAM ECC merr
61	ERA RPC CNTRAM ECC serr
62	ERA RPC CNTRAM ECC merr
63	ERA RPC IRRAM ECC serr
64	ERA RPC IRRAM ECC merr
65	ERA RPC BSRAM ECC serr
66	ERA RPC BSRAM ECC merr
67	ERA RPC ACRAM ECC serr
68	ERA RPC ACRAM ECC merr
69	ERA Priet 1bit err
70	ERA Priet 2bit err
71	ERA TCAM Scrub Parity Error
72	ERA Tx Packet Data FIFO Parity Error
73	ERA Tx Packet Control Info Parity Error
74	ERA RPC FIFO Overflow
	RPP
75	RPF FIFO Overflow
76	RHF NullFIFO Overflow
77	LEP FIFO Overflow
78	RCF Status FIFO Overflow
79	CCM Status FIFO Overflow
80	CCM Status FIFO undrflow
81	Rx Header FIFO Parity Error
82	CCM Status FIFO Parity Error

TABLE 21 NP memory errors supported on BR-MLX-10Gx24 interface cards (continued)

Error	Description
83	CCM Session Table Read Uncorrectible ECC Error
84	CCM Session Table Read Correctible ECC Error
85	CCM Hash Table Read Uncorrectible ECC Error
86	CCM Hash Table Read Correctible ECC Error
87	CCM Hash Bucket Read Uncorrectible ECC Error
88	CCM Hash Bucket Read Correctible ECC Error
89	BFD Session Table Read Uncorrectible ECC Error
90	BFD Session Table Read Correctible ECC Error
91	BFD Session Rx Count Statistics Read Parity Error
	TPP
92	Transmit Packet FIFO Read Parity Error
93	TPP Statistics Parity Error
94	Write Command Reply Packet Parity Error
95	Read Command Reply Packet Parity Error
	IFE
96	RPC ACRAM Table Memory Un-Correctable ECC Error
97	RPC accumulator Table Memory Correctable ECC Error
98	RPC Max Burst size Table Memory Un-Correctable ECC Error
99	RPC Max Burst size Table Memory Correctable ECC Error
100	RPC Credit Increment Table Memory Un-Correctable ECC Error
101	RPC Credit Increment Table Memory Correctable ECC Error
102	RPC Remap Table Memory Un-Correctable ECC Error
103	RPC Remap Table Memory Correctable ECC Error
104	RPC Count Table Memory Un-Correctable ECC Error
105	RPC Count Table Memory Correctable ECC Error
106	Rx Port Pipeline QOS done FIFO Parity Error
107	Rx Port Pipeline Flag FIFO Parity Error
108	Rx Port Pipeline Header FIFO Parity Error
109	Port Number FIFO UnCorrectable ECC Error
110	Port Number FIFO Correctable ECC Error
111	Rx Control iNPt FIFO UnCorrectable ECC Error
112	Rx Control iNPt FIFO Correctable ECC Error
113	Rx Port Pipeline FIFO Error
114	Rx Port Pipeline iNPt QOS ID FIFO Parity Error
115	Per port priority indexed counter Memory Parity Error

TABLE 22 NP memory errors supported on BR-MLX-100Gx2-X(100G) interface cards

Error	Description
External memory errors	
1	LBLRAM Parity Errors
2	Age RAM 1 Parity Errors

TABLE 22 NP memory errors supported on BR-MLX-100Gx2-X(100G) interface cards (continued)

Error	Description
3	Age RAM 2 Parity Errors
4	CAM1 Interface Parity Errors
5	CAM2 Interface Parity Errors
6	CAM3 Interface Parity Errors
7	TXCAM Interface Parity Error
Internal memory errors	
1	Multicast VLAN flag FIFO Parity
2	Multicast VLAN cPacket FIFO Parity
3	Multicast VLAN sfsTable Parity
4	Multicast VLAN repTable Parity
5	Multicast VLAN sfs FIFO Parity
6	Multicast VLAN sop FIFO Parity
7	Multicast VLAN pld FIFO Parity
8	Packet Edit Data FIFO Parity
9	Packet Edit sop FIFO Parity
10	Packet Edit merge FIFO Parity
11	Nexthoptable lkup Data Parity
12	ACL PRAM Results FIFO Parity
13	ACL feed FIFO Parity
14	ACL Data FIFO Parity
15	ACL ctrl FIFO Parity
16	ACL qosdone Parity
17	ACL portnum Parity
18	ACL priet Parity
19	Tx vlan Result Parity
20	Framer ctrl FIFO Parity
21	Framer Data FIFO Parity
22	Packet Edit Data FIFO rdData Parity
23	Packet Edit ctrl FIFO rdData Parity
24	Packet Edit nhlk FIFO rdData Parity
25	Packet Edit lble FIFO rdData Parity
26	CPU2startofs Read Data Parity
27	CPU2replace Read Data Parity
28	CPU2gentable nhtable Read Data Parity
29	CPU2gentable Txvlan Read Data Parity
30	CPU priet Read Data Parity
31	Rx MAC Data FIFO Parity Status
32	Rx MAC Flag FIFO Parity Status
33	CAM1 Result Data FIFO Parity Status
34	CAM2 Result Data FIFO Parity Status
35	CAM3 Result Data FIFO Parity Status

TABLE 22 NP memory errors supported on BR-MLX-100Gx2-X(100G) interface cards (continued)

Error	Description
36	CAM Result Scheduler FIFO Parity Status
37	CAM1 PacketID Mismatch
38	CAM2 PacketID Mismatch
39	CAM3 PacketID Mismatch
40	CAM Result FIFO Parity Status
41	Label PRAM Result Scheduler FIFO Parity Status
42	Service PRAM Result FIFO Parity Status
43	Rx Packethdr Result FIFO Parity Status
44	Rx Packet ID Mismatch
45	Rx Control FIFO Parity Status
46	Rx Data FIFO Parity Status
47	Ageram1 FIFO 1 Parity Status
48	Ageram1 FIFO 2 Parity Status
49	Ageram2 FIFO 1 Parity Status
50	Ageram1 Aging Entry FIFO Parity Status
51	Ageram2 Aging Entry FIFO Parity Status
52	Rx Data FIFO Mismatch
53	Service PRAM Result FIFO Parity Status
54	Packet Header Misc Parity Status
55	Packet Header 2 Parity Status
56	Packet Header 1 Parity Status
57	Packet Header 0 Parity Status
58	Service CAM Lookup FIFO Parity Status
59	CAM1 ASC FIFO Parity Status
60	CAM1 Lookup FIFO-1 Parity Status
61	CAM1 Lookup FIFO-2 Parity Status
62	CAM2 ASC FIFO Parity Status
63	CAM2 Lookup FIFO-1 Parity Status
64	CAM2 Lookup FIFO-2 Parity Status
65	CAM3 ASC FIFO Parity Status
66	CAM3 Lookup FIFO-1 Parity Status
67	CAM3 Lookup FIFO-2 Parity Status
68	PRAM Result FIFO Parity
69	Trunk adjusted header Parity
70	Packet Tablerd Parity
71	Service PRAM Result FIFO Parity Status
72	Packet Header Misc Parity Status
73	Packet Header 2 Parity Status
74	Packet Header 1 Parity Status
75	Packet Header 0 Parity Status
76	Rx Packetdecode FIFO Parity

TABLE 22 NP memory errors supported on BR-MLX-100Gx2-X(100G) interface cards (continued)

Error	Description
77	Rx topotos FIFO Parity
78	Eval0 trunk group Table Parity
79	Eval1 trunk group Table Parity
80	Eval2 trunk group Table Parity
81	Eval3 trunk group Table Parity
82	Eval4 trunk group Table Parity
83	Eval5 trunk group Table Parity
84	Merged CAM Result FIFO0 Parity
85	Merged CAM Result FIFO1 Parity
86	Merged PRAM Result FIFO0 Parity
87	Merged PRAM Result FIFO1 Parity
88	Ored 7 ram m p Result Parity
89	Rx Data in Parity
90	Rxctrl FIFO Read Data Parity
91	Read Rx qosid FIFO Parity
92	Rx portnum FIFO Parity
93	Rx qosdone FIFO Parity
94	Rx flag FIFO Parity
95	Rx header FIFO Parity
96	HQoS Table Parity
97	PRAM1 ECMP FIFO rd Parity
98	PRAM1 Read Request FIFO Parity
99	PRAM1 CPU rd FIFO Parity
100	PRAM2 Read Request FIFO Parity
101	PRAM2 CPU rd FIFO Parity
102	PRAM3 Read Request FIFO Parity
103	PRAM3 CPU rd FIFO Parity
104	CAM2PRAM1 Read Request FIFO Parity
105	CAM2PRAM1 ECMP FIFO rdData Parity
106	CAM2PRAM1 CPU FIFO rdData Parity
107	CAM2PRAM1 mw FIFO rd Parity
108	CAM2PRAM2 Read Request FIFO Parity
109	CAM2PRAM2 ECMP FIFO rdData Parity
110	CAM2PRAM2 CPU FIFO rdData Parity
111	CAM2PRAM2 mw FIFO rd Parity
112	CAM2PRAM3 Read Request FIFO Parity
113	CAM2PRAM3 ECMP FIFO rdData Parity
114	CAM2PRAM3 CPU FIFO rdData Parity
115	CAM2PRAM3 mw FIFO rd Parity
116	CAM Result FIFO Parity
117	Rx MAC Data FIFO Parity Status

TABLE 22 NP memory errors supported on BR-MLX-100Gx2-X(100G) interface cards (continued)

Error	Description
118	Rx MAC Flag FIFO Parity Status

TABLE 23 NP memory errors supported on Gen-1 and Gen-1.1 interface cards

Internal memory errors	
1	Rx Data FIFO Pointer Mismatch
2	Rx Control FIFO Pointer Mismatch
3	CAM1 ASC FIFO Mismatch
4	CAM2 ASC FIFO Mismatch
5	CAM3 ASC FIFO Mismatch

TABLE 24 NP memory errors supported on Gen-2 interface cards

Error	Description
External memory errors	
1	PRAM Word 0 Parity Error
2	PRAM Word 1 Parity Error
3	PRAM Word 2 Parity Error
4	PRAM Word 3 Parity Error
5	PRAM Word 4 Parity Error
6	PRAM Word 5 Parity Error
7	PRAM Word 6 Parity Error
8	PRAM Word 7 Parity Error
9	CAM2PRAM Word 0 Parity Error
10	CAM2PRAM Word 1 Parity Error
11	CAM2PRAM Word 2 Parity Error
12	CAM2PRAM Word 3 Parity Error
13	LBLRAM Word 0 Parity Error
14	LBLRAM Word 1 Parity Error
15	LBLRAM Word 2 Parity Error
16	LBLRAM Word 3 Parity Error
17	CAM1 Word 0 Parity Error
18	CAM1 Word 1 Parity Error
19	CAM1 GIO Parity Error
20	CAM1 PEO Parity Error
21	CAM1 Operation Error
22	CAM1 Dbase Parity Error
23	CAM2 Word 0 Parity Error
24	CAM2 Word 1 Parity Error
25	CAM2 GIO Parity Error
26	CAM2 PEO Parity Error
27	CAM2 Operation Error
28	CAM2 Dbase Parity Error

TABLE 24 NP memory errors supported on Gen-2 interface cards (continued)

Error	Description
29	CAM3 Word 0 Parity Error
30	CAM3 Word 1 Parity Error
31	CAM3 GIO Parity Error
32	CAM3 PEO Parity Error
33	CAM3 Operation Error
34	CAM3 Dbase Parity Error
Internal memory errors	
1	Tx ACL PRAM Results FIFO Parity Error
2	Tx VLAN Result Parity Error
3	Tx Frame Control Parity Error
4	Tx Frame Data Parity Error
5	Tx NextHop Table lookup Data Parity Error
6	Stats Data Parity Error
7	Spix Multicast VLAN Replace rData Parity Error
8	Rx Dispatch QoS Done FIFO Parity Error
9	Rx Dispatch Flag FIFO Parity Error
10	Rx Dispatch Header FIFO Parity Error
11	Rx Dispatch Port Number FIFO Parity Error
12	Rx Dispatch Dpath Cpath Packet Id Tag Parity Error
13	Rx Dispatch PRAM Decode Packet Id Tag Parity Error
14	Rx Dispatch Control Path Parity Error
15	Rx MAC0 Data FIFO Parity Error
16	Rx MAC1 Data FIFO Parity Error
17	Rx MAC2 Data FIFO Parity Error
18	Rx MAC3 Data FIFO Parity Error
19	Rx MAC0 ctrl FIFO Parity Error
20	Rx MAC1 ctrl FIFO Parity Error
21	Rx MAC2 ctrl FIFO Parity Error
22	Rx MAC3 ctrl FIFO Parity Error
23	SPI0 Multicast VLAN SOP FIFO Parity Error
24	SPI0 Multicast VLAN SFS FIFO Parity Error
25	SPI0 Multicast VLAN Replication Table Parity Error
26	SPI0 Multicast VLAN Cpkt FIFO Parity Error
27	SPI0 Multicast VLAN Flag FIFO Parity Error
28	SPI1 Multicast VLAN SOP FIFO Parity Error
29	SPI1 Multicast VLAN SFS FIFO Parity Error
30	SPI1 Multicast VLAN Replication Table Parity Error
31	SPI1 Multicast VLAN Cpkt FIFO Parity Error
32	SPI1 Multicast VLAN Flag FIFO Parity Error
33	SPI2 Multicast VLAN SOP FIFO Parity Error
34	SPI2 Multicast VLAN SFS FIFO Parity Error

TABLE 24 NP memory errors supported on Gen-2 interface cards (continued)

Error	Description
35	SPI2 Multicast VLAN Replication Table Parity Error
36	SPI2 Multicast VLAN Cpkt FIFO Parity Error
37	SPI2 Multicast VLAN Flag FIFO Parity Error
38	SPI3 Multicast VLAN SOP FIFO Parity Error
39	SPI3 Multicast VLAN SFS FIFO Parity Error
40	SPI3 Multicast VLAN Replication Table Parity Error
41	SPI3 Multicast VLAN Cpkt FIFO Parity Error
42	SPI3 Multicast VLAN Flag FIFO Parity Error
43	Agezero Read Data Parity Error
44	CAM3 Async FIFO rbus Parity Error
45	CAM3 SyncFIFO rdData Lo Parity Error
46	CAM3 SyncFIFO rdData Hi Parity Error
47	CAM3 Lookup FIFO Parity Error
48	CAM2 Async FIFO rbus Parity Error
49	CAM2 SyncFIFO rdData Lo Parity Error
50	CAM2 SyncFIFO rdData Hi Parity Error
51	CAM2 Lookup FIFO Parity Error
52	CAM1 Async FIFO rbus Parity Error
53	CAM1 SyncFIFO rdData lo Parity Error
54	CAM1 SyncFIFO rdData hi Parity Error
55	CAM1 Lookup FIFO Parity Error
56	Eval0 Trunk group Table Parity Error
57	Eval1 Trunk group Table Parity Error
58	LBLPRAM Result Scheduler FIFO Parity Error
59	PRAM Result Scheduler FIFO Parity Error
60	Rx topotos FIFO Parity Error
61	Aged FIFO Read Data Parity Error
62	Cpu2replace rdData Parity Error
63	Rxctrl FIFO Read Data Parity Error
64	Config Read Data Parity Error
65	Cmd Read Data Parity Error
66	Cmpl Data Parity Error
67	Rx pktdecode FIFO Parity Error
68	Topo Table Read Packet Parity Error
69	Exp Table Read Packet Parity Error
70	PCP Table Read Packet Parity Error
71	DSCP Table Read Packet Parity Error
72	PRAM Result FIFO0 Parity Error
73	PRAM Result FIFO1 Parity Error
74	PRAM Result FIFO2 Parity Error
75	PRAM Result FIFO3 Parity Error

TABLE 24 NP memory errors supported on Gen-2 interface cards (continued)

Error	Description
76	Rx Packet Header Service PRAM FIFO Parity Error
77	Rx Packet Header Misc FIFO Parity Error
78	Rx Packet Header2 FIFO Parity Error
79	Rx Packet Header1 FIFO Parity Error
80	Rx Packet Header0 FIFO Parity Error
81	Packet Table Read Parity Error
82	NextHop Table Read Data Parity Error
83	TxVLAN Read Data Parity Error
84	Label Lookup FIFO Overflow
85	LBLRAM Read Request FIFO Parity Error
86	LBLRAM Txp lkupFIFO Parity Error
87	LBLRAM LBL lkupFIFO Parity Error
88	LBLRAM cpu FIFO rdData Parity Error
89	PRAM ecmp FIFO rd Parity Error
90	PRAM Read Request FIFO Parity Error
91	PRAM cpu rdFIFO Parity Error
92	Rx CAMResult FIFO Parity Error
93	CAM2PRAM mwFIFO Parity Error
94	CAM2PRAM Read Request FIFO Parity Error
95	CAM2PRAM ecmp FIFO rdData Parity Error
96	CAM2PRAM cpu FIFO rdData Parity Error
97	sCAM Result ReadData Parity Error
98	mCAM Result ReadData Parity Error
99	LBLlkup Lookup FIFO Underflow
100	Txplkup Lookup FIFO Underflow
101	Txplkup Lookup FIFO Overflow
102	LBLRAM Read Request FIFO Underflow
103	LBLRAM Read Request FIFO Overflow
104	MAC0 Frame ctrl Parity Error
105	MAC0 Frame Data Parity Error
106	MAC1 Frame ctrl Parity Error
107	MAC1 Frame Data Parity Error
108	MAC2 Frame ctrl Parity Error
109	MAC2 Frame Data Parity Error
110	MAC3 Frame ctrl Parity Error
111	MAC3 Frame Data Parity Error
112	Sp0 Tx Frame ctrl Parity Error
113	Sp0 Tx Frame Data Parity Error
114	Sp1 Tx Frame ctrl Parity Error
115	Sp1 Tx Frame Data Parity Error
116	Sp2 Tx Frame ctrl Parity Error

TABLE 24 NP memory errors supported on Gen-2 interface cards (continued)

Error	Description
117	Sp2 Tx Frame Data Parity Error
118	Sp3 Tx Frame ctrl Parity Error
119	Sp3 Tx Frame Data Parity Error

LP CPU high-usage monitoring

This sections discusses the following topics:

LP CPU high-usage monitoring overview

When the CPU usage on the interface card goes high, it may lead to protocol flaps, timeouts, network convergence issues, etc. While in this state, it is possible to collect some data about the state and analyze it to find the root cause leading to the high CPU usage. The system can monitor itself and collect data to a file which is exportable and easy to analyze. Use the **show sysmon logs** command to see if information has been collected.

LP CPU high-usage monitoring: basic configuration

By default:

- LP CPU high-usage monitoring is disabled.
- When monitoring is enabled, LP CPU utilization is checked every 100 ms using a timer. Only when the LP CPU utilization remains above the specified threshold for three consecutive readings, a syslog message and a debug file are generated. You can set the threshold from 50% to 100%, the default is 80%.
- If LP CPU utilization is below the specified threshold at any of the three consecutive sampling points, the tracking logic is reset and no debug file is generated.
- If the LP CPU utilization cannot be sampled for three consecutive periods (300 ms), the system logs it as a LP CPU high-usage condition and captures data to help determine the task that caused the condition even though the CPU utilization may have been lower than the specified threshold at the sampling points.

Configuring LP CPU high usage monitoring

You can configure:

- LP CPU usage monitoring on all or individual LPs.
- The threshold for one or more LPs.

To enable LP CPU usage monitoring on all LPs, enter the following command:

```
device(config)# sysmon lp-high-cpu enable all
```

To enable LP CPU usage monitoring on the LP for interface slot 3, enter the following command:

```
device(config)# sysmon lp-high-cpu enable 3
```

To disable LP CPU usage monitoring on all LPs, enter the following command:

```
device(config)# no sysmon lp-high-cpu enable all
```

To set the LP CPU usage threshold to 90% on all LPs for which monitoring is enabled, enter the following command:

```
device(config)# sysmon lp-high-cpu threshold 90
```

To reset the LP CPU usage threshold to the default value (80%) on all LPs for which monitoring is enabled, enter the following command:

```
device(config)# no sysmon lp-high-cpu threshold
```

MP CPU high-usage monitoring

MP CPU high-usage monitoring allows the automatic detection of high CPU conditions and the logging of them in a text file. These conditions occur when the percentage of CPU usage or the amount of time that the task holds the CPU exceed their defined thresholds.

When the CPU usage on the MP becomes high or any task holds the CPU for a long time, it may lead to timeouts, network convergence issues, the IPC transmit queue and ITC queues becoming full, and heartbeat loss. If you are available to observe these states on the system, you can collect and analyze some data about the states to find the root cause leading to the high CPU usage. However, if you are unavailable when these issues occur, you can configure the system to monitor itself and to collect the data in log files automatically.

MP CPU high-usage monitoring and data collection

By default, CPU monitoring is disabled on the active and standby MP CPUs. When enabled, the system monitors the CPUs when the MPs are in the up state for the following events:

- Overall high CPU usage that exceeds 90 percent by default for three seconds or more, based on average CPU utilization
- A task holding the CPU that exceeds 400 milliseconds (ms), the default value
- The CPU cannot be checked and the detection timer cannot run for one second

When any of these events occur, the system logs it in a text file. The system can have a maximum of 10 high-CPU log files. After 10 files, the system stores the most recent files and replaces the older files starting from file number 1. The file name format is $$$ $MHC*count-date-timestamp*. The *count* is the file number from 1 to 10.

NOTE
MP high CPU usage data capture is part of the **show sysmon logs** command output and does not appear on the MP console.

The collected data in the log includes the output from the following **show** commands:

- **show emac stat detail** (three times)
- **show cpu** (two times)
- **show cpu histogram sequence trace 3**
- **show cpu histogram hold** (two times)

Configuring MP CPU high-usage monitoring

Configure MP CPU high-usage monitoring to allow the system to automatically detect high CPU usage on the active MP CPU and standby MP CPU. When the percentage of the CPU used or the amount of time that the task holds the CPU exceeds the configured threshold, the system triggers data collection into a log file. By default, MP CPU high-usage monitoring is disabled.

To configure MP CPU high-usage monitoring, perform the following steps.

1. In privileged EXEC mode, enter global configuration mode.

```
device # configure terminal
```

2. Enable MP CPU high-usage monitoring.

```
device(config)# sysmon mp-high-cpu enable
```

   By default, MP CPU high-usage monitoring is disabled.

3. Configure the MP CPU high-usage threshold to trigger data collection.

```
device(config)# sysmon mp-high-cpu cpu-threshold 70
```

   In this example, the threshold to log the event is 70 percent. The default setting is 90 percent.

4. Configure the threshold for the amount of time that the task holds the MP CPU to trigger data collection.

```
device(config)# sysmon mp-high-cpu task-threshold 300
```

   In this example, the task threshold to log the event is 300 milliseconds (ms). The default setting is 400 ms.

The following configuration is an example of the previous steps.

```
device # configure terminal
device(config)# sysmon mp-high-cpu enable
device(config)# sysmon mp-high-cpu cpu-threshold 70
device(config)# sysmon mp-high-cpu task-threshold 300
```

LP and MP IPC reliable TX queue monitoring

Interprocess Communication (IPC) reliable transmission (TX) queue monitoring allows the system to detect when the IPC Reliable TX queues on the LP and MP become stuck and the system generates syslog messages. When the queue recovers, the system detects it and generates a related syslog message. By default, the system does not monitor the LP and MP IPC reliable TX queues.

In highly scaled systems, IPC packets may be dropped and acknowledgments (ACKs) may be lost or delayed. Meanwhile, the sender waits for an ACK before sending the next packet. As a result, the IPC queue becomes stuck.

The system maintains separate IPC reliable TX queues on the source side for all the destination slots. The MP has separate queues for all LPs. An LP has separate queues for the active MP and the standby MP.

On the MP, the queue may become stuck due to the following occurrences:

- The buffer becomes corrupt and the packets may be dropped but are not removed from the queue. The MP retransmits but does not receive an ACK.
- A buffer overflow occurs. The MP runs out of buffer space and cannot receive an ACK.

On the LP, the queue may become stuck due to the following occurrences:

- When IPC packets are received, they are verified for checksum. If corruption with their checksum occurs and verification fails, the LP drops these packets. The ACK is not sent.
- A delay in processing and in sending an ACK occurs.

Enabling LP and MP IPC reliable TX queue monitoring

Enable LP and MP IPC reliable TX queue monitoring for the system to automatically detect when the IPC TX queue is stuck and to generate a syslog message. When the queue recovers, the system detects it and generates a related syslog message.

By default, IPC reliable TX monitoring is disabled. To enable monitoring, perform the following steps.

1. In privileged EXEC mode, enter global configuration mode.

```
device # configure terminal
```

2. Enable IPC reliable TX queue monitoring.

```
device(config)# sysmon ipc rel-q-mon enable
```

   By default, IPC reliable TX queue monitoring is disabled.

The following configuration is an example of the previous steps.

```
device # configure terminal
device(config)# sysmon ipc rel-q-mon enable
```

Port CRC error monitoring test

This section discusses the following topics:

Port CRC error monitoring overview

The port CRC error monitoring test is a background diagnostic test which monitors each port and checks if the number of packets with CRC errors (MAC CRC error eounter) exceeds a pre-configured limit. This limit or threshold is configured as the number of CRC errors occurring over the polling interval of the diagnostic test. If the test fails on a port for more than a configured threshold, a diagnostic action, if enabled, will be triggered. The diagnostic action can be configured to disable the port where the CRC errors exceed the configured threshold.

The threshold for diagnostic action, is configured as the ratio of the number of test failures to the number of diagnostics tests run. For example, if the threshold is set to three failures out of five diagnostic test runs, then the diagnostic action, when enabled, will be triggered if the test fails three times in five consecutive diagnostic tests.

A syslog is generated every time a port CRC error monitoring test fails. A syslog message is also generated after a port is disabled in a port CRC error diagnostic action.

NOTE

Optionally, syslogs can be disabled , before they are logged again, for a specific number of events (refer to Configuring 'log-backoff ' for the port CRC error monitoring test on page 241). This applies to the syslog which is sent after the port CRC error monitoring test fails, but not to the syslog sent after a port is disabled. When a port is disabled in the port CRC error diagnostic action, a syslog will be logged to notify the user of the port state change irrespective of this command.

Port CRC error monitoring: basic configuration

By default the:

- Port CRC error monitoring is enabled

- Port CRC error monitoring test diagnostic action is set to **syslog** i.e. a syslog message is generated when port CRC errors exceed the configured threshold.

Configuring the port CRC error monitoring test

1. Configure the port CRC error counter limit.
2. Configure the polling period for the test.
3. Configure the threshold to trigger diagnostic action

   To configure the port CRC error counter limit to 20, enter the following command:

   ```
   device(config)# sysmon port port-crc-test counter port-crc-counter less-than 20
   ```

   Syntax: sysmon port port-crc-test counter port-crc-counter less-than *crc-count*

   The variable *crc-count* specifies the port CRC error count limit for the configured polling period. The range of values is 0 through 65535. The default value is 20.

   To configure the port CRC error monitoring test to run every 60 seconds, enter the following command:

   ```
   device(config)# sysmon port port-crc-test polling-period 60
   ```

   Syntax: sysmon port port-crc-test polling-period *secs*

   The variable *secs* specifies the polling period in seconds. The range of values is 0 through 65535. The default value is 60 seconds.

   To configure the threshold to trigger the diagnostic action, if the test fails more than three times during five continuous polls, enter the following command:

   ```
   device(config)# sysmon port port-crc-test threshold 3 5
   ```

   Syntax: sysmon port port-crc-test threshold *num-failuresnum-polls*

   The *num-failures* variable specifes the number of failed test runs. The range of values is 1 through 31.

   The*num-polls* variable specifies the number of polls (tests). The range of values is 2 through 31.

   The default theshold is 3 failed test runs out of 5 polls.

Disabling the port CRC error monitoring test

The port CRC error monitoring test is enabled by default.

To disable the port CRC error monitoring test, enter the following command:

```
device(config)# no sysmon port port-crc-test
```

To enable the test again, enter the following command:

```
device(config)# sysmon port port-crc-test
```

Syntax: [no] sysmon port port-crc-test

Configuring the port CRC error monitoring test diagnostic action

The port CRC error monitoring test diagnostic action can be configured as:

- *none* - no action is taken.
- *port-disable* - disable the port.

- *syslog* - generate a syslog message.

The default port CRC error monitoring test diagnostic action is *syslog*.

NOTE
When the diagnostic action is configured as **port-disable,** a syslog message will also be generated after a port is disabled.

Table 25 lists the commands to transition between port CRC error monitoring test diagnostic action states.

TABLE 25 Port CRC error monitoring test: diagnostic action states

Action State	none	syslog	port-disable
none		no sysmon port port-crc action none	sysmon port port-crc action port-disable
syslog	sysmon port port-crc action none		sysmon port port-crc action port-disable
port-disable	sysmon port port-crc action none	no sysmon port port-crc action port-disable	

To disable the port CRC error monitoring test diagnostic action, enter the following command:

```
device(config)# sysmon port port-crc-test action none
```

To set the diagnostic action to disable a port when the port CRC error limit crosses the configured threshold, enter the following command:

```
device(config)# sysmon port port-crc-test action port-disable
```

Syntax: sysmon port port-crc-test action { **none** | **syslog** | **port-disable** }

Configuring 'log-backoff ' for the port CRC error monitoring test

Syslog messages sent after a port CRC diagnostic test fails, can be disabled for a certain number of events. Syslog action will resume after the specified number of events.

To disable syslog for 1,000 events:

```
device(config)# sysmon port port-crc-test log-backoff 1000
```

Syntax: sysmon port port-crc-test log-backoff *num*

The variable *num* specifies the number of events to skip before logging syslog messages again. The range of values is 1 through 14,400.

CRC check on Hi-Gig header in Rx path

You can enable Hi-Gig CRC check on Rx path using the **higig-crc-check-rx** command.

By default, Hi-Gig CRC check on Rx path is disabled.

NOTE
This command is valid for 48x1G-T card only.

The **higig-crc-check-rx** command provides the following options.

```
device# higig-crc-check-rx ?
disable Disable Higig CRC check on Rx path
enable Enable Higig CRC check on Rx path
status Show status of higig CRC check on Rx path
```

TM DRAM CRC error monitoring

TM DRAM CRC error monitoring overview

The TM DRAM CRC error monitoring feature monitors CRC errors. A total of 30 ingress dram CRC errors in a minute is considered as one event. If the number of events are more than three, then the action will be taken depending on the user configuration. Threshold and number of events needed to take action are fixed and cannot be configured.

TM DRAM CRC error monitoring: basic configuration

The basic configuration of TM DRAM CRC error monitoring is as follows.

Syntax: sysmon tm ingress-dram-crc action disable-ports | none | reset-linecard | syslog

The default configuration is disable-ports.

The **disable-ports** keyword disables ports for DRAM CRC errors.

The **none** keyword specifies no action.

The **reset-linecard** keyword resets line cards for DRAM CRC errors.

The **syslog** keyword adds system log messages for the DRAM CRC errors.

When **disable-ports** is configured, then all ports belonging to the affected TM are disabled.

```
Feb 18 11:46:09:A:System: LP15/TM0: all ports down due to dram crc errors
Feb 18 11:46:09:I:System: Interface ethernet 15/7, state down - ingress dram crc
Feb 18 11:46:09:I:System: Interface ethernet 15/2, state down - ingress dram crc
```

When you configure **none**, then there will be no action taken even after 3 events.

When **reset-linecard** is configured, the affected LP will be reset.

```
Feb 18 11:47:22:D:System: Module reset in slot 15, TM errors detected
```

When **syslog** is configured, only the system log message is generated at the 4th event and no other action is taken.

```
May 18 12:05:47:A:System: LP15/TM0: dram crc errors are detected
```

Scheduled System Monitor

The system monitoring Runtime Diagnostic (RTD) framework supports scheduling in future, on-demand testing, and interface module (LP) specific testing/monitoring. These features are only for those tests that adopt the new RTD test execution framework.

Scheduled system monitor consists of the following tests:

- Future scheduling on page 243

- On-demand testing on page 243
- Slot specific monitoring and testing on page 243

Future scheduling

System monitoring RTD framework runs tests for diagnostics periodically when the system starts or when a line card comes up. As part of the RTD framework, future scheduling supports the ability to schedule a test at some point of time in the future. This test is required in any monitoring system as the user may want to schedule a test based on the condition of the system or the available resources. Users can view the results asynchronously using CLI commands when the test is completed. Scheduling feature does not work if the test is already running in the continuous polling mode.

NOTE
This test supports a single scheduled configuration and does not support multiple scheduling configurations.

On-demand testing

On-demand testing is a specific test that a user runs while monitoring the system depending upon the symptoms as and when required. On-demand testing feature does not work if the test is already running in the continuous polling mode. Users can specify the number of times the test is to be run.

NOTE
When a on-demand test is scheduled, then the user can not cancel the test before its completion.

Slot specific monitoring and testing

This feature enables running a test on one or more specific LPs that have the same configuration. Users can specify LPs on which this test has to be run. If the user does not specify a slot, the test runs on all LPs by default. This feature helps in reducing the number of other unwanted execution of tests.

NOTE
This test cannot be run on individual LPs with different configuration for a specific test. For example, for a port CRC test, the test cannot be scheduled at different time on different LPs, or it cannot be scheduled on one LP and run in continuous polling on other LPs.

Longest Prefix Match Next Hop Walk monitoring

The Longest Prefix Match (LPM) Next Hop (NH) Walk monitoring detects inconsistencies between the LPM next hop programming in the software and hardware, and it can generate a syslog warning or take corrective action to clear the affected routes. This feature is used only on the CES and CER series of devices.

You can configure the feature to operate automatically at scheduled intervals (polling-period intervals), or you can execute the recovery actions manually. You have different choices for recovery actions: specifying (1) to take no action, (2) to recovery only affect routes in any VRF, (3) to recover all routes in all VRFs, or (4) to issue a syslog for any LPM errors detected. You can also set the threshold (number of errors) at which automatic recovery action will initiate.

When a recovery action is initiated (either automatically or manually), IPv4 HW forwarding entries are deleted and re-installed for the affected routes in VRFs or for all routes in all VRFs (depending upon the recovery action that was specified).

For additional information, refer to the **sysmon lpm nh-walk** command in the *NetIron Command Reference Guide*.

Using Syslog

This appendix describes how to display Syslog messages and how to configure the Syslog facility, and lists the Syslog messages that a Extreme device can display during standard operation.

NOTE

This appendix does not list Syslog messages that can be displayed when a debug option is enabled.

A device's software can write syslog messages to provide information at the following severity levels:

- Emergencies
- Alerts
- Critical

- Errors
- Warnings
- Notifications
- Informational
- Debugging

The device writes the messages to a local buffer, which can hold up to 5000 entries.

You also can specify the IP address or host name of up to six Syslog servers. When you specify a Syslog server, the device writes the messages both to the system log and to the Syslog server.

Using a Syslog server ensures that the messages remain available even after a system reload. The device's local Syslog buffer is cleared during a system reload or reboot, but the Syslog messages sent to the Syslog server remain on the server.

The Syslog service on a Syslog server receives logging messages from applications on the local host or from devices such as a MLX Series. Syslog adds a time stamp to each received message and directs messages to a log file. Most Unix workstations come with Syslog configured. Some third party vendor products also provide Syslog running on NT.

Syslog uses UDP port 514 and each Syslog message thus is sent with destination port 514. Each Syslog message is one line with Syslog message format. The message is embedded in the text portion of the Syslog format. There are several subfields in the format. Keywords are used to identify each subfield, and commas are delimiters. The subfield order is insensitive except that the text subfield should be the last field in the message. All the subfields are optional.

Displaying Syslog messages

To display the Syslog messages in the device's local buffer, enter the following command at any level of the CLI.

```
device# show logging
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
    Buffer logging: level ACDMEINW, 3 messages logged
    level code: A=alert C=critical D=debugging M=emergency E=error
                I=informational N=notification W=warning
Static Log Buffer:
Dec 15 19:04:14:A:Fan 1, fan on right connector, failed
Dynamic Log Buffer (50 entries):
Dec 15 18:46:17:I:Interface ethernet 1/4, state up
Dec 15 18:45:21:I:Bridge topology change, vlan 4095, interface 4, changed
state to forwarding
Dec 15 18:45:15:I:Warm start
```

For information about the Syslog configuration information, time stamps, and dynamic and static buffers, refer to Displaying the Syslog configuration on page 247.

Enabling real-time display of Syslog messages

By default, to view Syslog messages generated by a device, you need to display the Syslog buffer or the log on a Syslog server used by the device.

You can enable real-time display of Syslog messages on the management console. When you enable this feature, the software displays a Syslog message on the management console when the message is generated.

When you enable the feature, the software displays Syslog messages on the serial console when they occur. However, to enable display of real-time Syslog messages in Telnet or SSH sessions, you also must enable display within the individual sessions.

To enable real-time display of Syslog messages, enter the following command at the global CONFIG level of the CLI.

```
device(config)# logging console
```

Syntax: [no] logging console

This command enables the real-time display of Syslog messages on the serial console. You can enter this command from the serial console or a Telnet or SSH session.

To also enable the real-time display for a Telnet or SSH session, enter the following command from the Privileged EXEC level of the session.

```
telnet@device# terminal monitor
Syslog trace was turned ON
```

Syntax: [no] terminal monitor

Notice that the CLI displays a message to indicate the status change for the feature. To disable the feature in the management session, enter the **terminal monitor** command again. The command toggles the feature on and off.

```
telnet@device# terminal monitor
Syslog trace was turned OFF
```

Here is an example of how the Syslog messages are displayed.

```
telnet@device# terminal monitor
Syslog trace was turned ON
SYSLOG: <9>device, Power supply 2, power supply on left connector, failed
SYSLOG: <14>device, Interface ethernet 1/6, state down
SYSLOG: <14>device, Interface ethernet 1/2, state up
```

Configuring the Syslog service

The procedures in this section describe how to perform the following Syslog configuration tasks:

- Specify a Syslog server. You can configure the device to use up to six Syslog servers. (Use of a Syslog server is optional. The system can hold up to 5000 Syslog messages in an internal buffer.)
- Change the level of messages the system logs.
- Change the number of messages the local Syslog buffer can hold.
- Display the Syslog configuration.
- Clear the local Syslog buffer.

Logging is enabled by default, with the following settings:

- Messages of all severity levels (Emergencies - Debugging) are logged.
- By default, up to 50 messages are retained in the local Syslog buffer. This can be changed.
- No Syslog server is specified.

Displaying the Syslog configuration

To display the Syslog parameters currently in effect on a device, enter the following command from any level of the CLI.

```
device> show logging
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
    Buffer logging: level ACDMEINW, 3 messages logged
    level code: A=alert C=critical D=debugging M=emergency E=error
                I=informational N=notification W=warning
Static Log Buffer:
```

```
Dec 15 19:04:14:A:Fan 1, fan on right connector, failed
Dynamic Log Buffer (50 entries):
Dec 15 18:46:17:I:Interface ethernet 1/4, state up
Dec 15 18:45:21:I:Bridge topology change, vlan 4095, interface 4, changed
state to forwarding
Dec 15 18:45:15:I:Warm start
```

Syntax: show logging

The Syslog display shows the following configuration information, in the rows above the log entries themselves.

TABLE 26 CLI Display of Syslog buffer configuration

This field...	Displays...
Syslog logging	The state (enabled or disabled) of the Syslog buffer.
messages dropped	The number of Syslog messages dropped due to user-configured filters. By default, the software logs messages for all Syslog levels. You can disable individual Syslog levels, in which case the software filters out messages at those levels. Refer to Disabling logging of a message level on page 253. Each time the software filters out a Syslog message, this counter is incremented.
flushes	The number of times the Syslog buffer has been cleared by the **clear logging** command. refer to Clearing the Syslog messages from the local buffer on page 256.
overruns	The number of times the dynamic log buffer has filled up and been cleared to hold new entries. For example, if the buffer is set for 100 entries, the 101st entry causes an overrun. After that, the 201st entry causes a second overrun.
level	The message levels that are enabled. Each letter represents a message type and is identified by the key (level code) below the value. If you disable logging of a message level, the code for that level is not listed.
messages logged	The total number of messages that have been logged since the software was loaded.
level code	The message levels represented by the one-letter codes.

Static and dynamic buffers

The software provides two separate buffers:

- **Static** - logs power supply failures, fan failures, and temperature warning or shutdown messages
- **Dynamic** - logs all other message types. In previous releases, power supply messages were displayed in static logs only, with only the last event logged in. The power supply messages are now displayed in both static and dynamic logs.

In the static log, new messages replace older ones, so only the most recent message is displayed. For example, only the most recent temperature warning message will be present in the log. If multiple temperature warning messages are sent to the log, the latest one replaces the previous one. The static buffer is not configurable.

The message types that appear in the static buffer do not appear in the dynamic buffer. The dynamic buffer contains up to the maximum number of messages configured for the buffer (50 by default), then begins removing the oldest messages (at the bottom of the log) to make room for new ones.

The static and dynamic buffers are both displayed when you display the log.

```
device(config)# show logging
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
    Buffer logging: level ACDMEINW, 3 messages logged
    level code: A=alert C=critical D=debugging M=emergency E=error
          I=informational N=notification W=warning
Static Log Buffer:
Dec 15 19:04:14:A:Fan 1, fan on right connector, failed
Dec 15 19:00:14:A:Fan 2, fan on left connector, failed
Dynamic Log Buffer (50 entries):
Dec 15 18:46:17:I:Interface ethernet 1/4, state up
Dec 15 18:45:21:I:Bridge topology change, vlan 4095, interface 4, changed
```

```
state to forwarding
Dec 15 18:45:15:I:Warm start
```

Notice that the static buffer contains two separate messages for fan failures. Each message of each type has its own buffer. Thus, if you replace fan 1 but for some reason that fan also fails, the software replaces the first message about the failure of fan 1 with the newer message. The software does not overwrite the message for fan 2, unless the software sends a newer message for fan 2.

When you clear log entries, you can selectively clear the static or dynamic buffer, or you can clear both. For example, to clear only the dynamic buffer, enter the following command at the Privileged EXEC level.

```
device# clear logging dynamic-buffer
```

Syntax: clear logging [**dynamic-buffer** | **static-buffer**]

You can specify **dynamic-buffer** to clear the dynamic buffer or **static-buffer** to clear the static buffer. If you do not specify a buffer, both buffers are cleared.

Time stamps

The contents of the time stamp differ depending on whether you have set the time and date on the onboard system clock:

- If you have set the time and date on the onboard system clock, the date and time are shown in the following format:*mm dd hh:mm:ss* where:
  - *mm* - abbreviation for the name of the month
  - *dd* - day
  - *hh* - hours
  - *mm* - minutes
  - *ss* - seconds

For example, "Oct 15 17:38:03" means October 15 at 5:38 PM and 3 seconds.

- If you have not set the time and date on the onboard system clock, the time stamp shows the amount of time that has passed since the device was booted, in the following format:*numdnumhnummnums*where:
  - *numd* - day
  - *numh* - hours
  - *numm* - minutes
  - *nums* - seconds

For example, "188d1h01m00s" means the device had been running for 188 days, 11 hours, one minute, and zero seconds when the Syslog entry with this time stamp was generated.

Example of Syslog messages on a device whose onboard clock is set

The example shows the format of messages on a device whose onboard system clock has been set. Each time stamp shows the month, the day, and the time of the system clock when the message was generated. For example, the system time when the most recent message (the one at the top) was generated was October 15 at 5:38 PM and 3 seconds.

```
device(config)# show log
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
    Buffer logging: level ACDMEINW, 38 messages logged
    level code: A=alert C=critical D=debugging M=emergency E=error
                I=informational N=notification W=warning
Static Log Buffer:
Dec 15 19:04:14:A:Fan 1, fan on right connector, failed
Dec 15 19:00:14:A:Fan 2, fan on left connector, failed
Dynamic Log Buffer (50 entries):
Oct 15 17:38:03:warning:list 101 denied tcp 10.157.22.191(0)(Ethernet 4/18
0000.001f.77ed) -> 10.99.4.69(http), 1 event(s)
```

```
Oct 15 07:03:30:warning:list 101 denied tcp 10.157.22.26(0)(Ethernet 4/18
0000.001f.77ed) -> 10.99.4.69(http), 1 event(s)
Oct 15 06:58:30:warning:list 101 denied tcp 10.157.22.198(0)(Ethernet 4/18
0000.001f.77ed) -> 10.99.4.69(http), 1 event(s)
```

Example of Syslog messages on a device whose onboard clock is not set

The example shows the format of messages on a device whose onboard system clock is not set. Each time stamp shows the amount of time the device had been running when the message was generated. For example, the most recent message, at the top of the list of messages, was generated when the device had been running for 21 days, seven hours, two minutes, and 40 seconds.

```
device(config)# show log
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
    Buffer logging: level ACDMEINW, 38 messages logged
    level code: A=alert C=critical D=debugging M=emergency E=error
                I=informational N=notification W=warning
Static Log Buffer:
Dynamic Log Buffer (50 entries):
21d07h02m40s:warning:list 101 denied tcp 10.157.22.191(0)(Ethernet 4/18
0000.001f.77ed) -> 10.99.4.69(http), 1 event(s)
19d07h03m30s:warning:list 101 denied tcp 10.157.22.26(0)(Ethernet 4/18
0000.001f.77ed) -> 10.99.4.69(http), 1 event(s)
17d06h58m30s:warning:list 101 denied tcp 10.157.22.198(0)(Ethernet 4/18
0000.001f.77ed) -> 10.99.4.69(http), 1 event(s)
```

Configuring an encrypted syslog server

You can configure up to six encrypted syslog servers, but only one is active at any time, with the other servers acting as standby. When you add an encrypted syslog server, if there is no active syslog server, a session is established with the configured server. If a new connection is added when an active session exists, a new session with another encrypted syslog server is not attempted.

A new syslog server session is attempted in the following scenarios:

- Current active encrypted syslog server configuration is removed or the SSL connection to the active syslog server is closed
- During a device reload
- During switch over of the management module
- No active syslog server is found when the device sends syslog messages

Attempts to connect to a new syslog server starts with the first configured syslog server. The device attempts to establish an SSL connection with a server until a successful SSL connection is established. During this interval, the trap hold down timer is started and all the syslog messages are queued. When the timer expires, the device sends queued log messages to the connected syslog server.

Configuring encrypted syslog servers requires two steps:

- Installing the SSL Client certificate from a remote machine
- Adding encrypted syslog servers

Installing the SSL client certificate

Before you can configure an encrypted syslog server for the device, you must install the SSL client certificate. Do one of the following to install the SSL client certificate.

Using TFTP:

Use TFTP to copy the SSL Client Certificate and private key from the remote machine if TFTP is enabled on the device. Enter the following commands in sequence in any order:

```
device# copy tftp flash 10.25.101.121 cert.p12 client-certificate
device# copy tftp flash 10.25.101.121 privkeyfile client-private-key
```

Syntax: copy tftp flash *remote_ip cert_file* **client-certificate**

and

Syntax: copy tftp flash *remote_ippriv_key_file* **client-private-key**

The *remote_ip* keyword specifies the IP address of the remote host where the SSL Client certificate and private key are present. The *cert_file* keyword specifies the filename of the SSL Client Certificate, and the *priv_key_file* keyword specifies the filename of the private key.

Using SCP

Use SCP to copy the SSL Client Certificate and private key from the remote machine. Enter the following commands in sequence in any order at the remote host where the SSL Client Certificate and private key are present:

```
Host# scp cert.p12 user@10.25.105.121:sslclientcert
Host# scp privkeyfile user@10.25.105.121:sslclientprivkey
```

Syntax: scp *cert_file* **user@***remote_ip* **:sslclientcert**

and

Syntax: scp *priv_key_file* **user@** *remote_ip* **:sslclientprivkey**

The *remote_ip* keyword specifies the IP address of the device. The *cert_file* keyword specifies the filename of the SSL Client Certificate, and the *priv_key_file* keyword specifies the filename of the private key.

Adding an encrypted syslog server

To configure an encrypted server connection, enter the following command:

```
device(config)# logging host 10.25.105.201 ssl-port 60514
```

Syntax: logging host [**ipv6**] *ip_address* | *ipv6_address* **ssl-port** *port*

The *ip-address* variable specifies the syslog server. The **port** variable specifies the SSL port that will be used to connect to the specified syslog server.

NOTE

You can configure an encrypted syslog server connection only after the device has been placed in the Common Criteria mode. While you can configure these when the device is in the Administrative mode, the configuration takes effect only after the device is put in the Common Criteria Operational mode.

Displaying the configured server connections

You can display the active encrypted syslog server connection with the **show ip ssl** command:

```
device# show ip ssl

Session Source IP    Source Port     Remote IP        Remote Port
0       10.25.105.80 633                                10.25.105.201        60514
```

In addition, you can use the show logging command to display the active SSL-encrypted syslog server along with the logging level information.

```
device# show logging
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
    Buffer logging: level ACDMEINW, 27 messages logged
    level code: A=alert C=critical D=debugging M=emergency E=error
                I=informational N=notification W=warning

Current active SSL syslog server: 10.25.105.201:60514
```

Ascending or descending option for show log command

A new option was added to the **show log** command that allows you to display the log in either ascending or descending order based on time. The command will still work without the option selected and will display the log in default descending chronological order. The command is executed as shown

```
device# show log ascending
```

Syntax: show log [**ascending** | **descending**]

The **ascending** option displays the oldest log entry first.

The **descending** option displays the most recent log entry first. This is the default condition and consistent with previous versions of the Multi-Service IronWare.

Disabling or re-enabling Syslog

Syslog is enabled by default. To disable it, enter the following command at the global CONFIG level.

```
device(config)# no logging on
```

Syntax: [no] logging on [*udp-port*]

The *udp-port* parameter specifies the application port used for the Syslog facility. The default is 514.

To re-enable logging, enter the following command.

```
device(config)# logging on
```

This command enables local Syslog logging with the following defaults:

- Messages of all severity levels (Emergencies - Debugging) are logged.
- Up to 50 messages are retained in the local Syslog buffer.
- No Syslog server is specified.

Disabling Syslog of an event

Enter the **no logging enable** command to disable syslogs of a particluar event. In the following example, the **nologging enable** command disables the syslog for SNMP authentication failure.

```
device(config) # no logging enable snmp-auth-failure
```

Syntax: [no] logging enable [bfd | cfm | config-changed | fan-speed-change | fan-state-change | link-state-change | mgmt-mod-redun-state-change | module-hotswap | mpls | mvrp-vlan |ntp | ospf| snmp-auth-failure |temp-error |user-login |vrrp-if-state-change]

The **bfd** option defines the log of changes in the status of the BFD session.

The **cfm** option defines the log of changes in the CFM operations.

The **config-changed** option defines the log of changes in the configuration data.

The **fan-speed-change** option defines the log of changes in the speed of the fan.

The **fan-state-change** option defines the log of changes in the state of the fan.

The **link-state-change** option defines the log of changes in the state of the link.

The **mgmt-mod-redun-state-change** option defines the log of changes in the redundant state of the management module.

The **module-hotswap** option defines the log of insertion and removal of modules.

The **mpls** option defines the log of changes in the state of MPLS VPLS and MPLS VLL.

The **mvrp-vlan** defines the log of changes in the state of MVRP VLAN.

The **ntp** option defines the log of changes in the state of the NTP response.

The **ospf** option defines the log of changes in the state of OSPF.

The **snmp-auth-failure** option defines the log of SNMP authentication failure events.

The **temp-error** option defines the log of temporary errors.

The **user-login** option defines the log of user names for login.

The **vrrp-if-state-change** option defines the log of changes in the state of VRRP interface.

Specifying a Syslog server

To specify a Syslog server, enter a command such as the following

```
device(config)# logging host 10.0.0.99
```

For backward compatibility, the software reads the old command syntax from the startup configuration, and converts it to the new command syntax in the running configuration.

Syntax: [no] logging host *ip-address* | *server-name*

Specifying an additional Syslog server

To specify an additional Syslog server, enter the **logging host***ip-addr* command again, as in the following example. You can specify up to six Syslog servers.

Enter a command such as the following

```
device(config)# logging host 10.0.0.99
```

For backward compatibility, the software reads the old command syntax from the startup configuration, and converts it to the new command syntax in the running configuration.

Syntax: [no] logging host *ip-address* |*server-name*

Disabling logging of a message level

If you want to disable the logging of a message level, you must disable each message level individually.

For example, to disable logging of debugging and informational messages, enter the following commands

```
device(config)# no logging buffered debugging
device(config)# no logging buffered informational
```

Syntax: [no] logging buffered *level|num-entries*

The *level* parameter can have one of the following values:

- *alerts*
- *critical*
- *debugging*
- *emergencies*
- *errors*
- *informational*
- *notifications*
- *warnings*

The commands in the example above change the log level to notification messages or higher. The software will not log informational or debugging messages. The changed message level also applies to the Syslog servers.

On a NetIron XMR and NetIron MLX, enter 1 - 5000 for *num-entries*.

On a NetIron CES and NetIron CER 2000, enter 1 - 5000 for *num-entries*.

Changing the number of entries for the local buffer

You also can use the **logging buffered** command to change the number of entries the local Syslog buffer can store.

```
device(config)# logging buffered 100
```

Syntax: [no] logging buffered *level|num-entries*

On a NetIron XMR and NetIron MLX, enter 1 - 5000 for *num-entries*.

The default number of messages is 50. The change takes effect immediately and does not require you to reload the software.

Changing the log facility

The Syslog daemon on the Syslog server uses a facility to determine where to log the messages from the device. The default facility for messages the device sends to the Syslog server is "user". You can change the facility using the following command.

> NOTE
> You can specify only one facility. If you configure the device to use two Syslog servers, the device uses the same facility on both servers.

```
device(config)# logging facility local0
```

Syntax: [no] logging facility *facility-name*

The *facility-name* can be one of the following:

- kern - kernel messages
- user - random user-level messages
- mail - mail system

- daemon - system daemons
- auth - security or authorization messages
- syslog - messages generated internally by Syslog
- lpr - line printer subsystem
- news - netnews subsystem
- uucp - uucp subsystem
- sys9 - cron or at subsystem
- sys10 - reserved for system use
- sys11 - reserved for system use
- sys12 - reserved for system use
- sys13 - reserved for system use
- sys14 - reserved for system use
- cron - cron or at subsystem
- local0 - reserved for local use
- local1 - reserved for local use
- local2 - reserved for local use
- local3 - reserved for local use
- local4 - reserved for local use
- local5 - reserved for local use
- local6 - reserved for local use
- local7 - reserved for local use

Displaying the interface name in Syslog messages

By default, an interface's slot number (if applicable) and port number are displayed when you display Syslog messages. If you want to display the name of the interface instead of its number, enter the following command.

```
device(config)# ip show-portname
```

This command is applied globally to all interfaces on the device.

Syntax: [no] ip show-portname

When you display the messages in the Syslog, you refer to the interface name under the Dynamic Log Buffer section. The actual interface number is appended to the interface name. For example, if the interface name is "lab" and its port number is "2", you refer to "lab2" displayed as in the example below.

```
device# show logging
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
    Buffer logging: level ACDMEINW, 3 messages logged
    level code: A=alert C=critical D=debugging M=emergency E=error
          I=informational N=notification W=warning
Static Log Buffer:
Dec 15 19:04:14:A:Fan 1, fan on right connector, failed
Dynamic Log Buffer (50 entries):
Dec 15 18:46:17:I:Interface ethernet Lab2
, state up
Dec 15 18:45:15:I:Warm start
```

Clearing the Syslog messages from the local buffer

To clear the Syslog messages stored in the device's local buffer, use the following command.

```
device# clear logging
```

Syntax: clear logging

Logging all CLI commands to Syslog

This feature allows you to log all valid CLI command from each user session into the system log.

To enable CLI command logging, enter the following command.

```
device(config)# logging cli-command
```

Syntax: [no] logging cli-command

Example of CLI command logging

In the following example, two CLI sessions are run. In the first example, a telnet session enables CLI command logging and configures **router bgp** and the BGP **no neighbor** command as shown.

```
telnet@ device(config)# logging cli-command
telnet@ device(config)# router bgp
telnet@ device(config-bgp)# no nei 10.1.1.8 remote 10
```

In the next example, a console session configures **router bgp** and the **BGP neighbor** command as shown.

```
device(config)# router bgp
device(config-bgp)# nei 10.1.1.8 remote 10
```

Using the **show log** command, you would refer to a series of log records as shown in the following.

```
device(config-bgp)# show log
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
    Buffer logging: level ACDMEINW, 24 messages logged
    level code: A=alert C=critical D=debugging M=emergency E=error
                I=informational N=notification W=warning
Dynamic Log Buffer (50 lines):
Sep 9 18:38:23:I:CLI CMD: "nei 10.1.1.8 remote 10" from console
Sep 9 18:38:21:I:CLI CMD: "router bgp" from console
Sep 9 18:38:07:I:CLI CMD: "no nei 10.1.1.8 remote 10" from telnet client 10.1.1.1
Sep 9 18:38:05:I:CLI CMD: "router bgp" from telnet client 10.1.1.1
```

Syslog messages

The tables that follow list all of the Syslog messages. The messages are listed by message level, in the following order:

- Emergencies (none)
- Alerts
- Critical
- Errors
- Warnings
- Notifications
- Informational

- Debugging

Syslog messages system

Message `CAM partition partition name warning total total-count , free current free-count, slot slot-number , ppcr ppcr-id`

Explanation Indicates that the CAM partition specified by the *partition name* has exceeded a threshold (configurable with a default value of within 5% of the capacity of the partition) and may soon overflow the threshold. The *free-count* specifies the amount of free space still available in the partition. The *slot-number* and ppcr *ppcr-id* indicate where the overflow is occurring. The *partition-name* includes the sub-partition ID if applicable.

Message Level Warning

Message `Error Failed to shutdown Power Supply PS-Num . Write Failed (offset 0x2, value 44, size 2).( Extreme NetIron XMR and Extreme MLX only).`

Explanation A power supply failed to shutdown because of its failure to access its registers.

Message Level Error

Message `Error Module down in slot 3, reason CARD_DOWN_REASON_BOOT_FAILED.Error Code (1).`

Explanation

- The error message displayed on the Management Module console when the Interface Module fails to boot up. The message will display the error code reason.
- When the Interface Module is in DOWN state, the error code is included in the dynamic buffer.

The error code is 0 when there is no error code reported from the Interface Module.

Message Level Error

Message `ISIS Memory Limit Exceeded`

Explanation IS-IS is requesting more memory than is available.

Message Level Alert

Message `System Cold start`

Explanation The device has been powered on.

Message Level Informational

Message `System Enough power available to power on module in slot num .`

Explanation There is enough power available in the chassis to power on the module in the specific slot number. The slot *num* refers to the slot number in the chassis.

Message Level Notification

Message `System Fan num , location , failed`

Explanation A fan has failed. The *num* is the power supply number. The *location* describes where the failed power supply is in the chassis. The location can be one of the following

Message Level Alert

Message `System: Health Monitoring: FE access failure detected on SFM num /FE num (NetIron XMR and NetIron MLX only)`

Explanation The management processor is unable to access the specified fabric element. This syslog message will be generated a maximum of once per ten minute period. The SFM and FE *num* parameters indicate the number of the switch fabric module and fabric element that could not be accessed

Message Level Alert

Message `System: Health Monitoring: TM Egress data errors detected on LP num /TM num`

Explanation The system has detected egress data errors on the specified line processor and traffic manager. The LP and TM *num* parameters indicate the number of the line processor and traffic manager on which the errors were detected.

Message Level Alert

Message `System IfIndex assignment was changed.)`
Explanation The maximum number of ifIndex per module has been changed.
Message Level Informational

Message `System Interface portnum is down (remote fault)`
Explanation The interface is down due to Remote Fault. This is indicated as "(remote fault)". The *portnum* is the port number of the interface.
Message Level Informational

Message `System Interface portnum , line protocol down`
Explanation The line protocol on a port has gone down. The *portnum* is the port number.
Message Level Informational

Message `System Interface portnum , line protocol up`
Explanation The line protocol on a port has come up. The *portnum* is the port number.
Message Level Informational

Message `System Interface portnum , state down`
Explanation A port has gone down.
The *portnum* is the port number.
Message Level Informational

Message `System Interface portnum , state up`
Explanation A port has come up.
The *portnum* is the port number.
Message Level Informational

Message `System: LP's IPC Reliable TX Queue: slot slot-number:recovered`
Explanation The syslog message is generated when the LP IPC reliable transmission (TX) queue recovers from being stuck.
Message Level Informational

Message `System: LP's IPC Reliable TX Queue:slot slot-number:stuck`
Explanation The syslog message is generated when the LP IPC reliable TX queue is stuck.
Message Level Warning

Message `System Management module at slot slot-num state changed from module-state to module-state due to reason .`
Explanation Indicates a state change in a management module. The *slot-num* indicates the chassis slot containing the module. The *module-state* can be one of the following:

- active
- standby
- crashed
- coming-up
- unknown

A **due to** clause has been added to this message. The *reason* variable can be either or the following:

- MP upgrade to ver *version number* where *version number* is the version number of the Multi-Service IronWare software that the management module was upgraded to.
- Active Reboot

Message Level Alert

Message `System: Mbridge FPGA mismatch between Active and standby Module`

Explanation There is a mismatch in the field-programmable gate array (FPGA) versions between the active and standby management module.

Message Level Warning

Message `System Module n CPU m crashed`

Explanation

Message Level Informational

Message `System Module down in slot n reason`

Explanation Indicates that the module in the slot specified by the *n* variable is down for one of the following reasons as specified by the *reason* variable:

- CARD_DOWN_REASON_NONE
- CARD_DOWN_REASON_ADMIN_DOWN
- CARD_DOWN_REASON_CONFIG_MISMATCH
- CARD_DOWN_REASON_LOSS_HEARTBEAT
- CARD_DOWN_REASON_BOOT_FAILED
- CARD_DOWN_REASON_TIMEOUT
- CARD_DOWN_REASON_STRIPE_SYNC_FAILED
- CARD_DOWN_REASON_REBOOTED
- CARD_DOWN_REASON_OVER_HEAT
- CARD_DOWN_REASON_POWERED_OFF_BY_USER
- CARD_DOWN_REASON_LINK_DOWN

Message Level Notification

Message `System Module n powered on`

Explanation

Message Level Notification

Message `System Module n powered off`

Explanation

Message Level Notification

Message `System Module up in slot n`

Explanation

Message Level Notification

Message `System Module was inserted to slot slot-num`

Explanation Indicates that a module was inserted into a chassis slot. The *slot-num* is the number of the chassis slot into which the module was inserted.

Message Level Notification

Message `System Module was removed from slot slot-num`

Explanation Indicates that a module was removed from a chassis slot. The *slot-num* is the number of the chassis slot from which the module was removed.

Message Level Notification

Message `System: MP's IPC Reliable TX Queue: slot slot-number:recovered`

Explanation The syslog message is generated when the MP IPC reliable TX queue recovers from being stuck.

Message Level Informational

Message `System: MP's IPC Reliable TX Queue:slot slot-number:stuck`

Explanation The syslog message is generated when the MP IPC reliable TX queue is stuck.
Message Level Warning

Message `System Not enough power to power on module in slot num`
Explanation There is not enough power available in the chassis to power on the module in the specific slot number.
The slot *num* refers to the slot number in the chassis.
Message Level Warning

Message `System num-modules modules and 1 power supply, need more power supply!`
Explanation Indicates that the chassis needs more power supplies to run the modules in the chassis. The *num-modules* parameter indicates the number of modules in the chassis.
Message Level Alert

Message `System portnum is down (local fault)`
Explanation The port is down due to Local Fault. This is indicated as "(local fault)".
The *portnum* is the port number of the interface.
Message Level Informational

Message `System Power supply num, location , failed`
Explanation A power supply has failed. The *num* is the power supply number. The *location* describes where the failed power supply is in the chassis.
Message Level Alert

Message `System Power Supply PS-Num is shutdown due to flapping.( Extreme XMR and Extreme MLX only).`
Explanation A power supply is shut down because of flapping. The *PS-Num* is the power supply number.
Message Level Informational

Message `System Power Supply PS-Num will be shutdown due to flapping next time it becomes available. ( Extreme NetIron XMR and Extreme MLX only).`
Explanation A power supply will shutdown because of flapping the next time it is available. The *PS-Num* is the power supply number.
Message Level Informational

Message `System power type Power Supply num, location , state`
Explanation The *power type* refers to the AC or DC power supply. The *num* is the power supply number as positioned in the chassis. The *location* describes where the power supply is in the chassis in relation to its state. The *state* refers to how the power supply is functioning in the chassis. The *state* can be one of the following:

- Installed (OK): The power supply is installed and operating normally.
- Installed (Failed or Disconnected): The power supply has failed, or the power cord is disconnected.
- Not Installed (FAILED): The power supply is physically removed from the chassis.

Message Level Alert

Message `System Set fan speed to speed percentage`
Explanation Indicates that the fan speed has been changed to the value described in the *speed* variable and that the fan is now operating at the *percentage* of capacity described. The possible *speedpercentage* values are:

- LOW (50%)
- MEDIUM (75%)
- MEDIUM-HIGH (90%)
- HIGH (100%)

Message Level Notification

Message	`System SSH	telnet server enabled	disabled from console	telnet	ssh	web	snmp session [by user username]`
Explanation	A user enabled or disabled an SSH or Telnet session, or changed the SSH enable or disable configuration through the Web, SNMP, console, SSH, or Telnet session.
Message Level	Informational
Message	`System Switch fabric n powered off`
Explanation	
Message Level	Notification
Message	`System Switch fabric n powered on`
Explanation	
Message Level	Notification
Message	`System Syslog server IP-address deleted	added	modified from console	telnet	ssh	web	snmp OR Syslog operation enabled	disabled from console	telnet	ssh	web	snmp`
Explanation	A user made Syslog configuration changes to the specified Syslog server address, or enabled or disabled a Syslog operation through the Web, SNMP, console, SSH, or Telnet session.
Message Level	Informational
Message	`System Temperature degrees C degrees, warning level warn-degrees C degrees, shutdown level shutdown-degrees C degrees`
Explanation	Indicates an overtemperature condition on the active module. The *degrees* value indicates the temperature of the module. The *warn-degrees* value is the warning threshold temperature configured for the module. The *shutdown-degrees* value is the shutdown temperature configured for the module.
Message Level	Alert
Message	`System Warm start`
Explanation	The system software (flash code) has been reloaded.
Message Level	Informational

Syslog messages security

Message	`Security Port security violation at interface portnum , address mac , vlan id`
Explanation	
Message Level	Warning
Message	`Security Interface portnum was shut down due to port security violation`
Explanation	
Message Level	Warning
Message	`Security console login {by user I null } to USER EXEC mode Security {telnet I ssh} login {by user I null } from src {IP ip I IPv6 ipv6-addr } to USER EXEC mode`
Explanation	A user has logged into the USER EXEC mode of the CLI. The *user* is the user name.
Message Level	Informational
Message	`Security console logout {by user I null } from USER EXEC mode Security {telne I ssh} logout {by user I null } from src {IP ip I IPv6 ipv6-addr } from USER EXEC mode`
Explanation	A user has logged out of the USER EXEC mode of the CLI.

The *user* is the user name.

Message Level Informational

Message `Security console login {by user I null } to Privileged EXEC mode Security {telnet I ssh} login {by user I null } from src {IP ip I IPv6 ipv6-addr } to Privileged EXEC mode`

Explanation A user has logged into the Privileged EXEC mode of the CLI.

The *user* is the user name.

Message Level Informational

Message `Security console logout {by user I null } from Privileged EXEC mode Security {telnet I ssh} logout {by user I null } from src {IP ip I IPv6 ipv6-addr } from Privileged EXEC mode`

Explanation A user has logged out of Privileged EXEC mode of the CLI.

The *user* is the user name.

Message Level Informational

Message `Security outbound telnet session number login to server IP ip from SSH session session number`

Explanation A user has initiated an outbound Telnet session from an inbound SSH session.

The first *session number* is the number of the outbound Telnet session.

The *ip* is the IP address to which the Telnet session is connected.

The second *sessions number* is the number of the inbound SSH session.

Message Level Informational

Message `Security outbound telnet session number logout from server IP ip from SSH session session number`

Explanation A user has terminated an outbound Telnet session initiated from an inbound SSH session.

The first *session number* is the number of the outbound Telnet session.

The *ip* is the IP address from which the Telnet session has disconnected.

The second *sessions number* is the number of the inbound SSH session.

Message Level Informational

Message `Security startup-config was changed {by user } from {web management I snmp management I ssh client ip I telnet client ip }`

Explanation A configuration change was saved to the startup configuration file.

The *user* is the user's ID, if they entered a user ID to log in.

Message Level Informational

Message `Security running-config was changed {by user } from {web management I snmp management I ssh client ip I telnet client ip }`

Explanation A configuration change was saved to the running configuration file.

The *user* is the user's ID, if they entered a user ID to log in.

Message Level Informational

Message `Security telnet | SSH | web access [by username ] from src IP source ip address , src MAC source MAC address rejected, n attempts`

Explanation There were failed web, SSH, or Telnet login access attempts from the specified source IP and MAC address.

- [by *user username*] does not appear if telnet or SSH clients are specified.
- *n* is the number of times this SNMP trap occurred in the last five minutes, or other configured number of minutes.

Message Level Informational

Message `Security user username added | deleted | modified from console | telnet | ssh | web | snmp`

Explanation A user created, modified, or deleted a local user account through the Web, SNMP, console, SSH, or Telnet session.

Message Level Informational

Message `Security Enable super | port-config | read-only password deleted | added | modified from console | telnet | ssh | web | snmp OR Line password deleted | added | modified from console | telnet | ssh | web | snmp`

Explanation A user created, re-configured, or deleted an Enable or Line password through the Web, SNMP, console, SSH, or Telnet session.

Message Level Informational

Message `Apr 2 11:00:39:I:Security: telnet access from src IP ip address rejected, 1 attempt(s).`

Explanation A user sees this message when attempting to login with Telnet when the standby MP is in SYNC_SW State.

Message Level Informational

Message `Apr 2 11:00:39:I:Security:SSH access from src IP ip address rejected, 1 attempt(s).`

Explanation A user sees this message when attempting to login with SSH when the standby MP is in SYNC_SW State.

Message Level Informational

Syslog messages VLAN

Message `VLAN Id vlan-id added | deleted | modified from console | telnet | ssh | web | snmp session`

Explanation A user created, modified, or deleted a VLAN through the Web, SNMP, console, SSH, or Telnet session.

Message Level Informational

Syslog messages STP

Message `STP VLAN id - New RootBridge string RootPort portnum (reason )`

Explanation A Spanning Tree Protocol (STP) topology change has occurred.

The *id* is the ID of the VLAN in which the STP topology change occurred.

The *portnum* is the number of the port connected to the new root bridge.

Message Level Informational

Message `STP VLAN id - Bridge is RootBridge string ( reason )`

Explanation A Spanning Tree Protocol (STP) topology change has occurred, resulting in the device becoming the root bridge.

The *id* is the ID of the VLAN in which the STP topology change occurred.

Message Level Informational

Message `STP VLAN id Port portnum - Bridge TC Event ( reason )`

Explanation	A Spanning Tree Protocol (STP) topology change has occurred on a port. The *id* is the ID of the VLAN in which the STP topology change occurred. The *portnum* is the port number.
Message Level	Informational
Message	`STP VLAN vlanid Port portnum - State state ( reason )`
Explanation	
Message Level	Informational
Message	`STP Root Guard Port portnum , VLAN vlan-id inconsistent (Received superior BPDU)`
Explanation	The specified port was blocked because i t has Root Guard enabled and received a superior BPDU.
Message Level	Informational
Message	`STP Root Guard Port portnum , VLAN vlan-id consistent (Timeout)`
Explanation	The specified block Root Guard-protected port was unblocked.
Message Level	Informational
Message	`STP BPDU Guard port portnum disable System Interface ethernet portnum , state down - disabled`
Explanation	The **spanning-tree protect do-disable** command is configured on the specified port and the port became disabled due to a receipt of a BPDU packet.
Message Level	Informational
Message	`STP BPDU Guard re-enabled on ports ethe portnum System Interface ethernet portnum , state up`
Explanation	The **spanning-tree protect re-enable** was issued to re-enable the specified port
Message Level	Informational

Syslog messages RSTP

Message	`RSTP VLAN id Port portnum - Bridge TC Event (reason)`
Explanation	802.1W recognized a topology change event in the bridge. The topology change event is the forwarding action that started on a non-edge Designated port or Root port.
Message Level	Informational
Message	`RSTP VLAN id Port portnum - STP State state (reason)`
Explanation	802.1W changed the state of a port to a new state forwarding, learning, blocking. If the port changes to blocking, the bridge port is in discarding state.
Message Level	Informational
Message	`RSTP VLAN id - New RootPort portnum (reason)`
Explanation	802.1W changed the port's role to Root port, using the root selection computation.
Message Level	Informational
Message	`RSTP VLAN id - New RootBridge string RootPort portnum (reason)`
Explanation	802.1W selected a new root bridge as a result of the BPDUs received on a bridge port.
Message Level	Informational
Message	`RSTP VLAN id - Bridge is RootBridge string (reason)`
Explanation	802.1W changed the current bridge to be the root bridge of the given topology due to administrative change in bridge priority.
Message Level	Informational

Message `vlan vlan-id Bridge is RootBridge mac-address (MsgAgeExpiry)`
Explanation The message age expired on the Root port so 802.1W changed the current bridge to be the root bridge of the topology.
Message Level Informational

Syslog messages LAG

Message `LAG group ( ports ) created by 802.3ad link-aggregation module.`
Explanation 802.3ad link aggregation is configured on the device, and the feature has dynamically created a LAG group (aggregate link).

The *ports* is a list of the ports that were aggregated to make the LAG group.
Message Level Informational

Syslog messages MRP

Message `MRP interface ethernet portnum vlan vlan-master , changing to state-string`
Explanation
Message Level Informational

Message `MRP metro ring ring-id cannot be enabled. No free CAM entries`
Explanation
Message Level Informational

Syslog messages UDLD

Message `UDLD Logical link on interface ethernet portnum is up`
Explanation
Message Level Informational

Message `UDLD Logical link on interface ethernet portnum is down`
Explanation
Message Level Informational

Syslog messages VSRP

Message `VSRP VLAN` *`vlanid`* `VRID` *`id`* `- transition to` *`state-string`*
Explanation
Message Level Informational

Message `VSRP VLAN` *`vlanid`* `VRID` *`id`* `- aware change` *`old-portnum`* `->` *`new-portnum`*`\n`
Explanation
Message Level Informational

Message `VSRP VLAN vlanid VRID id - aware learn portnum`
Explanation
Message Level Informational

Syslog messages VRRP

Message `VRRP intf state changed, intf portnum , vrid virtual-router-id ,state vrrp-state`

Explanation A state change has occurred in a Virtual Router Redundancy Protocol (VRRP) interface.

The *portnum* is the port.

The *virtual-router-id* is the virtual router ID (VRID) configured on the interface.

The *vrrp-state* can be one of the following:

- init
- master
- backup
- unknown

Message Level Notification

Syslog messages IP

Message `IP Dup IP ip-addr detected, sent from MAC mac-addr interface portnum`

Explanation Indicates that the device received a packet from another device on the network with an IP address that is also configured on the device.

The *ip-addr* is the duplicate IP address.

The *mac-addr* is the MAC address of the device with the duplicate IP address.

The *portnum* is the port that received the packet with the duplicate IP address. The address is the packet's source IP address.

Message Level Warning

Syslog messages ICMP

Message `ICMP Local ICMP exceeds burst-max burst packets, stopping for lockup seconds!`

Explanation The number of ICMP packets exceeds the *burst-max* threshold set by the **ip icmp burst** command. The device may be the victim of a Denial of Service (DoS) attack.

All ICMP packets will be dropped for the number of seconds specified by the *lockup* value. When the lockup period expires, the packet counter is reset and measurement is restarted.

Message Level Notification

Message `ICMP Transit ICMP in interface portnum exceeds num burst packets, stopping for num seconds!`

Explanation Threshold parameters for ICMP transit (through) traffic have been configured on an interface, and the maximum burst size for ICMP packets on the interface has been exceeded.

The *portnum* is the port number.

The first *num* is the maximum burst size (maximum number of packets allowed).

The second *num* is the number of seconds during which additional ICMP packets will be blocked on the interface.

NOTE
This message can occur in response to an attempted Smurf attack.

Message Level Notification

Syslog messages ACL

Message `ACL list acl-num denied ip-proto src-ip-addr (src-tcp/udp-port) (Ethernet portnummac-addr) -> dst-ip-addr (dst-tcp/udp-port), 1 events`

Explanation Indicates that an Access Control List (ACL) denied (dropped) packets.

The *acl-num* indicates the ACL number. Numbers 1 - 99 indicate standard ACLs. Numbers 100 - 199 indicate extended ACLs.

The *ip-proto* indicates the IP protocol of the denied packets.

The *src-ip-addr* is the source IP address of the denied packets.

The *src-tcp/udp-port* is the source TCP or UDP port, if applicable, of the denied packets.

The *portnum* indicates the port number on which the packet was denied.

The *mac-addr* indicates the source MAC address of the denied packets.

The *dst-ip-addr* indicates the destination IP address of the denied packets.

The *dst-tcp/udp-port* indicates the destination TCP or UDP port number, if applicable, of the denied packets.

Message Level Warning

Message `ACL:rip filter list list-num direction V1 | V2 denied ip-addr , num packets`

Explanation Indicates that a RIP route filter denied (dropped) packets.

The *list-num* is the ID of the filter list.

The *direction* indicates whether the filter was applied to incoming packets or outgoing packets. The value can be one of the following:

- in
- out

The V1 or V2 value specifies the RIP version (RIPv1 or RIPv2).

The *ip-addr* indicates the network number in the denied updates.

The *num* indicates how many packets matching the values above were dropped during the five-minute interval represented by the log entry.

Message Level Warning

Message `ACL insufficient L4 session resource, using flow based ACL instead`

Explanation The device does not have enough Layer 4 session entries.

To correct this condition, allocate more memory for sessions. To allocate more memory, enter the following command at the global CONFIG level of the CLI interface

system-max session-limit *num*

Message Level Notification

Message `ACL system fragment packet inspect rate rate exceeded`

Explanation The fragment rate allowed on the device has been exceeded.

	The *rate* indicates the maximum rate allowed. This message can occur if fragment throttling is enabled.
Message Level	Notification
Message	`AC port fragment packet inspect rate rate exceeded on port portnum`
Explanation	The fragment rate allowed on an individual interface has been exceeded. The *rate* indicates the maximum rate allowed. The *portnum* indicates the port. This message can occur if fragment throttling is enabled.
Message Level	Notification
Message	`ACL Port portnum , exceed configured L4 rule-based CAM size, larger L4 partition size required`
Explanation	
Message Level	Notification
Message	`ACL Port portnum , exceed configured L2 ACL rule-based CAM size, larger partition size is required`
Explanation	
Message Level	Notification
Message	`ACL Port portnum , exceed configured outbound L4 rule-based CAM size, larger outbound L4 partition size required`
Explanation	
Message Level	Notification
Message	`ACL Port portnum , exceed configured IPv6 L4 rule-based CAM size, larger IPv6 L4 partition size required`
Explanation	
Message Level	Notification
Message	`ACL Port portnum , exceed configured IPv6 outbound L4 rule-based CAM size, larger IPv6 outbound L4 partition size required`
Explanation	
Message Level	Notification
Message	`ACL Port portnum, error in allocating inbound L4 rule-based ACL CAM entry`
Explanation	
Message Level	Notification
Message	`ACL Port portnum , error in allocating outbound L4 rule-based ACL CAM entry`
Explanation	
Message Level	Notification
Message	`ACL Port portnum , inbound ACL CAM programming incomplete`
Explanation	
Message Level	Notification
Message	`ACL Port portnum , outbound ACL CAM programming incomplete`
Explanation	
Message Level	Notification

Message `ACL aclid added | deleted | modified from console | telnet | ssh | web | snmp session`
Explanation A user created, modified, deleted, or applied an ACL through the Web, SNMP, console, SSH, or Telnet session.
Message Level Informational

Syslog messages RACL

Message `RACL Port portnum , IP Receive ACL exceed configured CAM size, larger partition size required`
Explanation
Message Level Notification

Message `RACL Port portnum , IP Receive ACL exceed configured RL class limit`
Explanation
Message Level Notification

Message `RACL Port portnum , IP Receive ACL CAM malloc error`
Explanation
Message Level Notification

Syslog messages OSPF

Message `OSPF Memory Overflow`
Explanation OSPF has run out of memory.
Message Level Alert

Message `OSPF LSA Overflow, LSA Type = lsa-type`
Explanation Indicates an LSA database overflow.

The *lsa-type* parameter indicates the type of LSA that experienced the overflow condition. The LSA type is one of the following:

- 1 - Router
- 2 - Network
- 3 - Summary
- 4 - Summary
- 5 - External

Message Level Alert

Message `OSPF interface state changed,rid router-id , intf addr ip-addr , state ospf-state`
Explanation Indicates that the state of an OSPF interface has changed.

The *router-id* is the router ID of the device.

The *ip-addr* is the interface's IP address.

The *ospf-state* indicates the state to which the interface has changed and can be one of the following:

- down
- loopback
- waiting

- point-to-point
- designated router
- backup designated router
- other designated router
- unknown

Message Level Notification

Message `OSPF virtual intf state changed, rid router-id , area area-id , nbr ip-addr , state ospf-state`

Explanation Indicates that the state of an OSPF virtual routing interface has changed.

The *router-id* is the router ID of the router the interface is on.

The *area-id* is the area the interface is in.

The *ip-addr* is the IP address of the OSPF neighbor.

The *ospf-state* indicates the state to which the interface has changed and can be one of the following:

- down
- loopback
- waiting
- point-to-point
- designated router
- backup designated router
- other designated router
- unknown

Message Level Notification

Message `OSPF nbr state changed, rid router-id , nbr addr ip-addr , nbr rid nbr-router-Id , state ospf-state`

Explanation Indicates that the state of an OSPF neighbor has changed.

The *router-id* is the router ID of the device.

The *ip-addr* is the IP address of the neighbor.

The *nbr-router-id* is the router ID of the neighbor.

The *ospf-state* indicates the state to which the interface has changed and can be one of the following:

- down
- attempt
- initializing
- 2-way
- exchange start
- exchange
- loading
- full
- unknown

Message Level Notification

Message `OSPF virtual nbr state changed, rid router-id , nbr addr ip-addr , nbr rid nbr-router-id , state ospf-state`

Explanation Indicates that the state of an OSPF virtual neighbor has changed.

The *router-id* is the router ID of the device.

The *ip-addr* is the IP address of the neighbor.

The *nbr-router-id* is the router ID of the neighbor.

The *ospf-state* indicates the state to which the interface has changed and can be one of the following:

- down
- attempt
- initializing
- 2-way
- exchange start
- exchange
- loading
- full
- unknown

Message Level Notification

Message `OSPF intf config error, rid router-id , intf addr ip-addr , pkt src addr src-ip-addr , error type error-type , pkt type pkt-type`

Explanation Indicates that an OSPF interface configuration error has occurred.

The *router-id* is the router ID of the device.

The *ip-addr* is the IP address of the interface on the device.

The *src-ip-addr* is the IP address of the interface from which the device received the error packet.

The *error-type* can be one of the following:

- bad version
- area mismatch
- unknown NBMA neighbor
- unknown virtual neighbor
- authentication type mismatch
- authentication failure
- network mask mismatch
- hello interval mismatch
- dead interval mismatch
- option mismatch
- unknown

The *packet-type* can be one of the following:

- hello
- database description

- link state request
- link state update
- link state ack
- unknown

Message Level Notification

Message `OSPF virtual intf config error, rid router-id , intf addr ip-addr , pkt src addr src-ip-addr , error type error-type , pkt type pkt-type`

Explanation Indicates that an OSPF virtual routing interface configuration error has occurred.

The *router-id* is the router ID of the device.

The *ip-addr* is the IP address of the interface on the device.

The *src-ip-addr* is the IP address of the interface from which the device received the error packet.

The *error-type* can be one of the following:

- bad version
- area mismatch
- unknown NBMA neighbor
- unknown virtual neighbor
- authentication type mismatch
- authentication failure
- network mask mismatch
- hello interval mismatch
- dead interval mismatch
- option mismatch
- unknown

The *packet-type* can be one of the following:

- hello
- database description
- link state request
- link state update
- link state ack
- unknown

Message Level Notification

Message `OSPF intf authen failure, rid router-id , intf addr ip-addr , pkt src addr src-ip-addr , error type error-type , pkt type pkt-type`

Explanation Indicates that an OSPF interface authentication failure has occurred.

The *router-id* is the router ID of the device.

The *ip-addr* is the IP address of the interface on the device.

The *src-ip-addr* is the IP address of the interface from which the device received the authentication failure.

The *error-type* can be one of the following:

- bad version
- area mismatch
- unknown NBMA neighbor
- unknown virtual neighbor
- authentication type mismatch
- authentication failure
- network mask mismatch
- hello interval mismatch
- dead interval mismatch
- option mismatch
- unknown

The *packet-type* can be one of the following:

- hello
- database description
- link state request
- link state update
- link state ack
- unknown

Message Level Notification

Message `OSPF virtual intf authen failure,rid router-id , intf addr ip-addr , pkt src addr src-ip-addr , error type error-type , pkt type pkt-type`

Explanation Indicates that an OSPF virtual routing interface authentication failure has occurred.

The *router-id* is the router ID of the device.

The *ip-addr* is the IP address of the interface on the device.

The *src-ip-addr* is the IP address of the interface from which the device received the authentication failure.

The *error-type* can be one of the following:

- bad version
- area mismatch
- unknown NBMA neighbor
- unknown virtual neighbor
- authentication type mismatch
- authentication failure
- network mask mismatch
- hello interval mismatch
- dead interval mismatch
- option mismatch
- unknown

The *packet-type* can be one of the following:

- hello
- database description
- link state request
- link state update
- link state ack
- unknown

Message Level Notification

Message `OSPF intf rcvd bad pkt, rid router-id , intf addr ip-addr , pkt src addr src-ip-addr , pkt type pkt-type`

Explanation Indicates that an OSPF interface received a bad packet.

The *router-id* is the router ID of the device.

The *ip-addr* is the IP address of the interface on the device.

The *src-ip-addr* is the IP address of the interface from which the device received the authentication failure.

The *packet-type* can be one of the following:

- hello
- database description
- link state request
- link state update
- link state ack
- unknown

NOTE
This message is typically generated during BFD or OSPF reconverge within the following scenarios:

- The router is undergoing hitless upgrade
- Management module switchover,
- Interface module CPU utilization is at 95% or more,
- The **clear ip ospf neighbor all** command is issued.

During these processes, OSPF adj is deleted due to BFD time out while the router can still receive OSPF packets destined to a previous session from its neighbor because the neighbor has an inconsistent OSPF state due to timing. This message will go away shortly when BFD or OSPF re-establishes neighbor.

Message Level Notification

Message `OSPF virtual intf rcvd bad pkt, rid router-id , intf addr ip-addr , pkt src addr src-ip-addr , pkt type pkt-type`

Explanation Indicates that an OSPF interface received a bad packet.

The *router-id* is the router ID of the device.

The *ip-addr* is the IP address of the interface on the device.

The *src-ip-addr* is the IP address of the interface from which the device received the authentication failure.

The *packet-type* can be one of the following:

- hello
- database description
- link state request
- link state update
- link state ack
- unknown

Message Level Notification

Message `OSPF intf retransmit, rid router-id , intf addr ip-addr , nbr rid nbr-router-id ,pkt type is pkt-type , LSA type lsa-type ,LSA id lsa-id , LSA rid lsa-router-id`

Explanation An OSPF interface on the device has retransmitted a Link State Advertisement (LSA).

The *router-id* is the router ID of the device.

The *ip-addr* is the IP address of the interface on the device.

The *nbr-router-id* is the router ID of the neighbor router.

The *packet-type* can be one of the following:

- hello
- database description
- link state request
- link state update
- link state ack
- unknown

The *lsa-type* is the type of LSA.

The *lsa-id* is the LSA ID.

The *lsa-router-id* is the LSA router ID.

Message Level Notification

Message `OSPF virtual intf retransmit, rid router-id , intf addr ip-addr , nbr rid nbr-router-id , pkt type is pkt-type , LSA type lsa-type ,LSA id lsa-id , LSA rid lsa-router-id`

Explanation An OSPF interface on the device has retransmitted a Link State Advertisement (LSA).

The *router-id* is the router ID of the device.

The *ip-addr* is the IP address of the interface on the device.

The *nbr-router-id* is the router ID of the neighbor router.

The *packet-type* can be one of the following:

- hello
- database description
- link state request
- link state update
- link state ack

- unknown

The *lsa-type* is the type of LSA.

The *lsa-id* is the LSA ID.

The *lsa-router-id* is the LSA router ID.

Message Level Notification

Message `OSPF originate LSA, rid router-id , area area-id , LSA type lsa-type , LSA id lsa-id , LSA router id lsa-router-id`

Explanation An OSPF interface has originated an LSA.

The *router-id* is the router ID of the device.

The *area-id* is the OSPF area.

The *lsa-type* is the type of LSA.

The *lsa-id* is the LSA ID.

The *lsa-router-id* is the LSA router ID.

Message Level Notification

Message `OSPF max age LSA, rid router-id , area area-id , LSA type lsa-type , LSA id lsa-id , LSA rid lsa-router-id`

Explanation An LSA has reached its maximum age.

The *router-id* is the router ID of the device.

The *area-id* is the OSPF area.

The *lsa-type* is the type of LSA.

The *lsa-id* is the LSA ID.

The *lsa-router-id* is the LSA router ID.

Message Level Notification

Message `OSPF LSDB overflow, rid router-id , limit num`

Explanation A Link State Database Overflow (LSDB) condition has occurred.

The *router-id* is the router ID of the device.

The *num* is the number of LSAs.

Message Level Notification

Message `OSPF LSDB approaching overflow, rid router-id , limit num`

Explanation The software is close to an LSDB condition.

The *router-id* is the router ID of the device.

The *num* is the number of LSAs.

Message Level Notification

Message `OSPF intf rcvd bad pkt Bad Checksum, rid ip-addr , intf addr ip-addr , pkt size num , checksum num , pkt src addr ip-addr , pkt type type`

Explanation The device received an OSPF packet that had an invalid checksum.

The rid *ip-addr* is device's router ID.

The intf addr *ip-addr* is the IP address of the interface that received the packet.

The pkt size *num* is the number of bytes in the packet.

The checksum *num* is the checksum value for the packet.

The pkt src addr *ip-addr* is the IP address of the neighbor that sent the packet.

The pkt type *type* is the OSPF packet type and can be one of the following:

- hello
- database description
- link state request
- link state update
- link state acknowledgement
- unknown (indicates an invalid packet type)

Message Level Notification

Message `OSPF intf rcvd bad pkt Bad Packet type, rid ip-addr , intf addr ip-addr , pkt size num , checksum num , pkt src addr ip-addr , pkt type type`

Explanation The device received an OSPF packet with an invalid type.

The parameters are the same as for the Bad Checksum message. The pkt type *type* value is "unknown", indicating that the packet type is invalid.

Message Level Notification

Message `OSPF intf rcvd bad pkt Unable to find associated neighbor, rid ip-addr , intf addr ip-addr , pkt size num , checksum num , pkt src addr ip-addr , pkt type type`

Explanation The neighbor IP address in the packet is not on the device's list of OSPF neighbors.

The parameters are the same as for the Bad Checksum message.

Message Level Notification

Message `OSPF intf rcvd bad pkt Invalid packet size, rid ip-addr , intf addr ip-addr , pkt size num , checksum num , pkt src addr ip-addr , pkt type type`

Explanation The device received an OSPF packet with an invalid packet size.

The parameters are the same as for the Bad Checksum message.

Message Level Notification

Syslog messages OSPFv3

Message `OSPFv3 Memory Overflow`

Explanation OSPF has run out of memory.

Message Level Alert

Message `OSPFv3 LSA Overflow, LSA Type = lsa-type`

Explanation Indicates an LSA database overflow.

The *lsa-type* parameter indicates the type of LSA that experienced the overflow condition. The LSA type is one of the following:

- 1 - Router
- 2 - Network
- 3 - Summary

- 4 - Summary
- 5 - External

Message Level Alert

Message `OSPFv3 interface state changed,rid router-id , intf addr ip-addr , state ospf-state`

Explanation Indicates that the state of an OSPF interface has changed.

The *router-id* is the router ID of the device.

The *ip-addr* is the interface's IP address.

The *ospf-state* indicates the state to which the interface has changed and can be one of the following:

- down
- loopback
- waiting
- point-to-point
- designated router
- backup designated router
- other designated router
- unknown

Message Level Notification

Message `OSPFv3 virtual intf state changed, rid router-id , area area-id , nbr ip-addr , state ospf-state`

Explanation Indicates that the state of an OSPF virtual routing interface has changed.

The *router-id* is the router ID of the router the interface is on.

The *area-id* is the area the interface is in.

The *ip-addr* is the IP address of the OSPF neighbor.

The *ospf-state* indicates the state to which the interface has changed and can be one of the following:

- down
- loopback
- waiting
- point-to-point
- designated router
- backup designated router
- other designated router
- unknown

Message Level Notification

Message `OSPFv3 nbr state changed, rid router-id , nbr addr ip-addr , nbr rid nbr-router-Id , state ospf-state`

Explanation Indicates that the state of an OSPF neighbor has changed.

The *router-id* is the router ID of the device.

The *ip-addr* is the IP address of the neighbor.

The *nbr-router-id* is the router ID of the neighbor.

The *ospf-state* indicates the state to which the interface has changed and can be one of the following:

- down
- attempt
- initializing
- 2-way
- exchange start
- exchange
- loading
- full
- unknown

Message Level Notification

Message `OSPFv3 virtual nbr state changed, rid router-id , nbr addr ip-addr , nbr rid nbr-router-id , state ospf-state`

Explanation Indicates that the state of an OSPF virtual neighbor has changed.

The *router-id* is the router ID of the device.

The *ip-addr* is the IP address of the neighbor.

The *nbr-router-id* is the router ID of the neighbor.

The *ospf-state* indicates the state to which the interface has changed and can be one of the following:

- down
- attempt
- initializing
- 2-way
- exchange start
- exchange
- loading
- full
- unknown

Message Level Notification

Message `OSPFv3 intf config error, rid router-id , intf addr ip-addr , pkt src addr src-ip-addr , error type error-type , pkt type pkt-type`

Explanation Indicates that an OSPF interface configuration error has occurred.

The *router-id* is the router ID of the device.

The *ip-addr* is the IP address of the interface on the device.

The *src-ip-addr* is the IP address of the interface from which the device received the error packet.

The *error-type* can be one of the following:

- bad version
- area mismatch
- unknown NBMA neighbor
- unknown virtual neighbor

- authentication type mismatch
- authentication failure
- network mask mismatch
- hello interval mismatch
- dead interval mismatch
- option mismatch
- unknown

The *packet-type* can be one of the following:

- hello
- database description
- link state request
- link state update
- link state ack
- unknown

Message Level Notification

Message `OSPFv3 virtual intf config error, rid router-id , intf addr ip-addr , pkt src addr src-ip-addr , error type error-type , pkt type pkt-type`

Explanation Indicates that an OSPF virtual routing interface configuration error has occurred.

The *router-id* is the router ID of the device.

The *ip-addr* is the IP address of the interface on the device.

The *src-ip-addr* is the IP address of the interface from which the device received the error packet.

The *error-type* can be one of the following:

- bad version
- area mismatch
- unknown NBMA neighbor
- unknown virtual neighbor
- authentication type mismatch
- authentication failure
- network mask mismatch
- hello interval mismatch
- dead interval mismatch
- option mismatch
- unknown

The *packet-type* can be one of the following:

- hello
- database description
- link state request
- link state update
- link state ack

- unknown

Message Level Notification

Message `OSPFv3 intf authen failure, rid router-id , intf addr ip-addr , pkt src addr src-ip-addr , error type error-type , pkt type pkt-type`

Explanation Indicates that an OSPF interface authentication failure has occurred.

The *router-id* is the router ID of the device.

The *ip-addr* is the IP address of the interface on the device.

The *src-ip-addr* is the IP address of the interface from which the device received the authentication failure.

The *error-type* can be one of the following:

- bad version
- area mismatch
- unknown NBMA neighbor
- unknown virtual neighbor
- authentication type mismatch
- authentication failure
- network mask mismatch
- hello interval mismatch
- dead interval mismatch
- option mismatch
- unknown

The *packet-type* can be one of the following:

- hello
- database description
- link state request
- link state update
- link state ack
- unknown

Message Level Notification

Message `OSPFv3 virtual intf authen failure,rid router-id , intf addr ip-addr , pkt src addr src-ip-addr , error type error-type , pkt type pkt-type`

Explanation Indicates that an OSPF virtual routing interface authentication failure has occurred.

The *router-id* is the router ID of the device.

The *ip-addr* is the IP address of the interface on the device.

The *src-ip-addr* is the IP address of the interface from which the device received the authentication failure.

The *error-type* can be one of the following:

- bad version
- area mismatch
- unknown NBMA neighbor

- unknown virtual neighbor
- authentication type mismatch
- authentication failure
- network mask mismatch
- hello interval mismatch
- dead interval mismatch
- option mismatch
- unknown

The *packet-type* can be one of the following:

- hello
- database description
- link state request
- link state update
- link state ack
- unknown

Message Level Notification

Message `OSPFv3 intf rcvd bad pkt, rid router-id , intf addr ip-addr , pkt src addr src-ip-addr , pkt type pkt-type`

Explanation Indicates that an OSPF interface received a bad packet.

The *router-id* is the router ID of the device.

The *ip-addr* is the IP address of the interface on the device.

The *src-ip-addr* is the IP address of the interface from which the device received the authentication failure.

The *packet-type* can be one of the following:

- hello
- database description
- link state request
- link state update
- link state ack
- unknown

Message Level Notification

Message `OSPFv3 virtual intf rcvd bad pkt, rid router-id , intf addr ip-addr , pkt src addr src-ip-addr , pkt type pkt-type`

Explanation Indicates that an OSPF interface received a bad packet.

The *router-id* is the router ID of the device.

The *ip-addr* is the IP address of the interface on the device.

The *src-ip-addr* is the IP address of the interface from which the device received the authentication failure.

The *packet-type* can be one of the following:

- hello

- database description
- link state request
- link state update
- link state ack
- unknown

Message Level Notification

Message `OSPFv3 intf retransmit, rid router-id , intf addr ip-addr , nbr rid nbr-router-id ,pkt type is pkt-type , LSA type lsa-type ,LSA id lsa-id , LSA rid lsa-router-id`

Explanation An OSPF interface on the device has retransmitted a Link State Advertisement (LSA).

The *router-id* is the router ID of the device.

The *ip-addr* is the IP address of the interface on the device.

The *nbr-router-id* is the router ID of the neighbor router.

The *packet-type* can be one of the following:

- hello
- database description
- link state request
- link state update
- link state ack
- unknown

The *lsa-type* is the type of LSA.

The *lsa-id* is the LSA ID.

The *lsa-router-id* is the LSA router ID.

Message Level Notification

Message `OSPFv3 virtual intf retransmit, rid router-id , intf addr ip-addr , nbr rid nbr-router-id , pkt type is pkt-type , LSA type lsa-type ,LSA id lsa-id , LSA rid lsa-router-id`

Explanation An OSPF interface on the device has retransmitted a Link State Advertisement (LSA).

The *router-id* is the router ID of the device.

The *ip-addr* is the IP address of the interface on the device.

The *nbr-router-id* is the router ID of the neighbor router.

The *packet-type* can be one of the following:

- hello
- database description
- link state request
- link state update
- link state ack
- unknown

The *lsa-type* is the type of LSA.

The *lsa-id* is the LSA ID.

The *lsa-router-id* is the LSA router ID.

Message Level Notification

Message `OSPFv3 originate LSA, rid router-id , area area-id , LSA type lsa-type , LSA id lsa-id , LSA router id lsa-router-id`

Explanation An OSPF interface has originated an LSA.

The *router-id* is the router ID of the device.

The *area-id* is the OSPF area.

The *lsa-type* is the type of LSA.

The *lsa-id* is the LSA ID.

The *lsa-router-id* is the LSA router ID.

Message Level Notification

Message `OSPFv3 max age LSA, rid router-id , area area-id , LSA type lsa-type , LSA id lsa-id , LSA rid lsa-router-id`

Explanation An LSA has reached its maximum age.

The *router-id* is the router ID of the device.

The *area-id* is the OSPF area.

The *lsa-type* is the type of LSA.

The *lsa-id* is the LSA ID.

The *lsa-router-id* is the LSA router ID.

Message Level Notification

Message `OSPFv3 LSDB overflow, rid router-id , limit num`

Explanation A Link State Database Overflow (LSDB) condition has occurred.

The *router-id* is the router ID of the device.

The *num* is the number of LSAs.

Message Level Notification

Message `OSPFv3 LSDB approaching overflow, rid router-id , limit num`

Explanation The software is close to an LSDB condition.

The *router-id* is the router ID of the device.

The *num* is the number of LSAs.

Message Level Notification

Message `OSPFv3 intf rcvd bad pkt Bad Checksum, rid ip-addr , intf addr ip-addr , pkt size num , checksum num , pkt src addr ip-addr , pkt type type`

Explanation The device received an OSPF packet that had an invalid checksum.

The rid *ip-addr* is device's device ID.

The intf addr *ip-addr* is the IP address of the interface that received the packet.

The pkt size *num* is the number of bytes in the packet.

The checksum *num* is the checksum value for the packet.

The pkt src addr *ip-addr* is the IP address of the neighbor that sent the packet.

The pkt type *type* is the OSPF packet type and can be one of the following:

- hello
- database description
- link state request
- link state update
- link state acknowledgement
- unknown (indicates an invalid packet type)

Message Level Notification

Message `OSPFv3 intf rcvd bad pkt Bad Packet type, rid ip-addr , intf addr ip-addr , pkt size num , checksum num , pkt src addr ip-addr , pkt type type`

Explanation The device received an OSPF packet with an invalid type.

The parameters are the same as for the Bad Checksum message. The pkt type *type* value is "unknown", indicating that the packet type is invalid.

Message Level Notification

Message `OSPFv3 intf rcvd bad pkt Unable to find associated neighbor, rid ip-addr , intf addr ip-addr , pkt size num , checksum num , pkt src addr ip-addr , pkt type type`

Explanation The neighbor IP address in the packet is not on the device's list of OSPF neighbors.

The parameters are the same as for the Bad Checksum message.

Message Level Notification

Message `OSPFv3 intf rcvd bad pkt Invalid packet size, rid ip-addr , intf addr ip-addr , pkt size num , checksum num , pkt src addr ip-addr , pkt type type`

Explanation The device received an OSPF packet with an invalid packet size.

The parameters are the same as for the Bad Checksum message.

Message Level Notification

Syslog messages IS-IS

Message `ISIS Memory Limit Exceeded`

Explanation IS-IS is requesting more memory than is available.

Message Level Alert

Message `ISIS ENTERED INTO OVERLOAD STATE`

Explanation The device has set the overload bit to on (1), indicating that the device's IS-IS resources are overloaded.

Message Level Notification

Message `ISIS Entered Overload State Due to overload-reason`

Explanation The device has set the overload bit to on (1), indicating that the device's IS-IS resources are Overloaded.

Reasons for the overload as expressed in the *overload-reason* variable are:

- Configuration
- Startup Configuration
- LSP Buffer Allocation Failure
- LSP Header Allocation Failure

- Maximum Number of LSPs Exceeded
- LSP Fragmentation Count Exceeded
- LSP Sequence Number Wrap Around
- LSP Option Allocation Failure
- Path Entry Allocation Failure
- Route Entry Allocation Failure

Definitions of the *overload-reason* values are described in Table 27.

Message Level Notification

Message `ISIS Exited Overload State`

Explanation The device has set the overload bit to off (0), indicating that the device's IS-IS resources are no longer overloaded.

Message Level Notification

Message `ISIS L1 ADJACENCY DOWN system-id on circuit circuit-id`

Explanation The device's adjacency with this Level-1 IS has gone down.

The *system-id* is the system ID of the IS.

The *circuit-id* is the ID of the circuit over which the adjacency was established.

Message Level Notification

Message `ISIS L1 ADJACENCY UP system-id on circuit circuit-id`

Explanation The device's adjacency with this Level-1 IS has come up.

The *system-id* is the system ID of the IS.

The *circuit-id* is the ID of the circuit over which the adjacency was established.

Message Level Notification

Message `ISIS L2 ADJACENCY DOWN system-id on circuit circuit-id`

Explanation The device's adjacency with this Level-2 IS has gone down.

The *system-id* is the system ID of the IS.

The *circuit-id* is the ID of the circuit over which the adjacency was established.

Message Level Notification

Message `ISIS L2 ADJACENCY UP system-id on circuit circuit-id`

Explanation The device's adjacency with this Level-2 IS has come up.

The *system-id* is the system ID of the IS.

The *circuit-id* is the ID of the circuit over which the adjacency was established.

Message Level Notification

Message `ISIS LSP-type LSP LSP-ID Seq sequence-number Len length LifeTime lifetime on interface-name dropped due to LSP-drop-reason`

Explanation The device has dropped the received LSP.

The *LSP-Type* can be one of the following:

- L1
- L2

The *LSP-ID* variable is in the 8 byte LSP ID value.

The *sequence-number* is a 4 byte value that is associated with each LSP ID.

The *length* is the length of the LSP PDU.

The *lifetime* is the life period of the LSP.

The *interface-name* is the name of the interface and is displayed in the following form "Ethernet 1/1".

The *LSP-drop-reason* variable describes the following reasons that the LSP was dropped:

- Adjacency not found
- Adjacency Level Mismatch
- IS Level Mismatch
- Length Too Short
- Length Too Large
- Authentication Failure
- Max Area Check Failure
- Zero Checksum
- Checksum Mismatch
- Invalid Length

Definitions of the *LSP-drop-reason* values are described in Table 27.

Message Level Notification

Message `ISIS NbrType Neighbor Hostname/systemID DOWN on interface-name due to neighbor-down-reason`

Explanation The device's Neighbor has gone down.The *NbrType* can be one of the following:

- L1
- L2
- PTPT

The *interface-name* is the name of the interface and is displayed in the following form "Ethernet 1/1".

The *neighbor-down-reason* variable can be any one of the following reasons that the Neighbor is Down:

- BFD Trigger
- Maximum Adjacencies
- User Trigger
- Hold Timer Expiry
- Adjacency ID Mismatch
- Adjacency Type Mismatch
- Interface Down
- Interface State Change

Definitions of the *neighbor-down-reason* values are described in Table 27.

Message Level Notification

Message `ISIS NbrType neighbor Hostname/systemID UP on interface-name`

Explanation The device's Neighbor has come up.

The *NbrType* can be one of the following:

- L1

- L2
- PTPT

The *interface-name* is the name of the interface and is displayed in the following form "Ethernet 1/1".

Message Level Notification

Message Explanation `ISIS PTP ADJACENCY DOWN mac on interface portnum`

Message Level Notification

Message Explanation `ISIS PTP ADJACENCY UP mac on interface portnum`

Message Level Notification

TABLE 27 Definition of IS-IS variables

Variable	Value	Definition
neighbor-down-reason	BFD Trigger	BFD identified link failures and triggered IS-IS to clean the neighbors on that link.
	Maximum Adjacencies	IS-IS has reached the maximum number of adjacencies. Therefore, it has deleted the adjacency with the lowest SNPA address to accommodate the new adjacency.
	User Trigger	The user triggered to delete the adjacency using the clear **isis neighbor systemID** command or the **clear isis all** command.
	Hold Timer Expiry	The adjacency was deleted because there were no "hellos" received within the hold time period.
	Adjacency ID Mismatch	The adjacency was deleted because the new "hello" received from this adjacency has a different System ID.
	Adjacency Type Mismatch	The adjacency was deleted because the new "hello" received from this adjacency has a different adjacency Type.
	Interface Down	The adjacency was deleted because the interface went down.
	Interface State Change	The adjacency was deleted because the interface state has changed due to user configuration.
overload-reason	Configuration	The Overload condition was entered because of a user configuration.
	Startup Configuration	The Overload condition was entered because of the startup configuration.
	LSP Buffer Allocation Failure	The Overload condition was entered because of an LSP buffer allocation error.
	LSP Header Allocation Failure	The Overload condition was entered because of an LSP header allocation error.
	Maximum Number of LSPs Exceeded	The Overload condition was entered because the LSP count reached the maximum value.
	LSP Fragmentation Count Exceeded	The Overload condition was entered because of IS-IS trying to generate the 256th LSP fragment.
	LSP Sequence Number Wrap Around	The Overload condition was entered because the LSP numbers reached the maximum value.
	LSP Option Allocation Failure	Self LSP building failed due to an internal buffer allocation failure.
	Path Entry Allocation Failure	The SPF computation failed due to a Path Entry allocation failure.
	Route Entry Allocation Failure	The SPF computation failed due to a Route Entry allocation failure.
LSP-drop-reason	Adjacency not found	The LSP was dropped because the is no adjacency found on the interface.
	Adjacency Level Mismatch	The LSP was dropped because the adjacency is at a different level from the LSP level.

TABLE 27 Definition of IS-IS variables (continued)

Variable	Value	Definition
	IS Level Mismatch	The LSP was dropped because IS-IS is configured at a different level than the LSP level.
	Length Too Short	The LSP length is shorter than the LSP header length.
	Length Too Large	The LSP length is larger than the Maximum LSP buffer length.
	Authentication Failure	The LSP was dropped because of an authentication failure.
	Max Area Check Failure	The LSP has a Max Area Count different than the configured Max Area Count of the device.
	Zero Checksum	The LSP has a zero checksum.
	Checksum Mismatch	The LSP checksum is different than the computed checksum.
	Invalid Length	The LSP length is different than the sum of the option lengths in the LSP.

Syslog messages ITC and IPC queue usage

Message `Apr 17 05:25:19:W:ITC destination task ITC_APP_SCP : Queue Usage exceeds threshold- 80 percent of total queue length 1048576 bytes`

Explanation Indicates that the inter-task communications (ITC) destination task queue usage has exceeded the threshold value of 80 percent of the total queue length. The syslog message is generated only on the active MP module. An SNMP trap is also generated when the ITC destination task usage is above the threshold value.
ITC queue errors are sometimes seen in a highly scaled network when a system reload, system switchover, or MP Reset, and during Hitless Operating System Switchover (HLOS).

Message Level Warning

Message `Apr 17 05:26:19:I:ITC destination task ITC_APP_SCP : Queue Usage has come back to normal which is below threshold- 80 percent of total queue length 1048576 bytes`

Explanation Indicates that the ITC destination task queue usage on the active MP module is now back to normal and below the threshold value of 80 percent of the total queue length. An SNMP trap is also generated when the ITC destination task queue is normal and below the threshold value.

Message Level Informational

Message `Apr 17 05:10:28:W:ITC source task ITC_APP_CONSOLE : Retry Queue Usage exceeds threshold- 80 percent of total queue length 16384 bytes`

Explanation Indicates that the ITC source task retry queue usage has exceeded the threshold value of 80 percent of the total queue length. The syslog message is generated only on the active MP module. An SNMP trap is also generated when the ITC source task retry usage is above the threshold value.

Message Level Warning

Message `Apr 17 05:06:49:I:ITC source task Console : Retry Queue Usage has come back to normal which is below threshold- 80 percent of total queue length 16384 bytes`

Explanation Indicates that the ITC source task retry queue usage on the active MP module is now back to normal and below the threshold value of 80 percent of the total queue length. An SNMP trap is also generated when the ITC source task retry queue is normal and below the threshold value.

Message Level Informational

Message `Apr 17 05:27:10:W:IPC reliable TX Queue usage for destination slot 1 exceeds threshold- 80 percent of total queue length 1024`

Explanation Indicates that the Interprocessor Communications (IPC) reliable TX queue usage for destination slot 1 has exceeded the threshold value of 80 percent of the total queue length. The syslog message is generated only on the active MP module. An SNMP trap is also generated when the IPC reliable TX queue usage is above the threshold value.

Message Level Warning

Message `Apr 17 05:28:10:I:IPC Reliable TX Queue usage for destination slot 1 has come back to normal which is below threshold- 80 percent of total queue length 1024`

Explanation Indicates that the IPC reliable TX queue usage on the active MP module for destination slot 1 is now back to normal and below the threshold value of 80 percent of the total queue length. An SNMP trap is also generated when the IPC reliable TX queue usage is normal and below the threshold value.

Message Level Informational

Syslog messages BGP

Message `BGP4 Not enough memory available to run BGP4`

Explanation The device could not start the BGP4 routing protocol because there is not enough memory available.

Message Level Debug

Message `BGP No of prefixes received from BGP peer ip-addr exceeds maximum prefix-limit...shutdown`

Explanation The device has received more than the specified maximum number of prefixes from the neighbor, and the device is therefore shutting down its BGP4 session with the neighbor.

Message Level Error

Message `BGP received invalid AS4_PATH attribute length (3) - entire AS4_PATH ignored`

Explanation Possible attribute length can be only even number and cannot be odd. If an attribute with odd length is received, this error is displayed.

Message Level Error

Message `BGP received invalid AS4_PATH attribute flag (0x40) - entire AS4_PATH ignored`

Explanation If the flag that describes the attribute has unacceptable values then this error is displayed.

Message Level Error

Message `BGP received invalid Confed info in AS4_PATH (@byte 43) - entire AS4_PATH ignored`

Explanation Confederation segments(AS_CONFED_SEQ/SET) must precede the (AS_SEQ/SET), if not, this error is displayed.

Message Level Error

Message `BGP received incorrect Seq type/len in AS4_PATH (@byte 41) - entire AS4_PATH ignored`

Explanation Valid segment types are (AS_SEQ/SET, AS_CONFED_SEQ/SET), any other values results in an error being displayed.

Message Level Error

Message `BGP received multiple AS4_PATH attributes - used first AS4_PATH attribute only`

Explanation When AS4_PATH is received more than one time in the update message, this error is displayed.

Message Level Error

Message `BGP No of prefixes received from BGP peer ip-addr exceeds warning limit num`

Explanation The device has received more than the allowed percentage of prefixes from the neighbor.

The *ip-addr* is the IP address of the neighbor.

The *num* is the number of prefixes that matches the percentage you specified. For example, if you specified a threshold of 100 prefixes and 75 percent as the warning threshold, this message is generated if the device receives a 76th prefix from the neighbor.

Message Level Warning

Message `BGP Peer ip-addr UP (ESTABLISHED)`

Explanation Indicates that a BGP4 neighbor has come up.

The *ip-addr* is the IP address of the neighbor's BGP4 interface with the device.

Message Level Notification

Message `BGP Peer ip-addr DOWN (IDLE)`

Explanation Indicates that a BGP4 neighbor has gone down.

The *ip-addr* is the IP address of the neighbor's BGP4 interface with the device.

Message Level Notification

Message `BGP Peer ip DOWN ( reasonrecv notif )`

Explanation

Message Level Notification

Message `Configuration (Wait for BGP)`

Explanation IS-IS is waiting for BGP convergence to complete.

Message Level Notification

Syslog messages NTP

Message `NTP server ip-addr failed to respond`

Explanation Indicates that a Network Time Protocol (NTP) server did not respond to the device's query for the current time.

The *ip-addr* indicates the IP address of the NTP server.

Message Level Warning

Message `<server | sym_active | sym_passive> association is mobilized for <ipv4 address | ipv6 address>`

Explanation Indicates the mobilization of a new NTP server, or symmetric active or symmetric passive association with the peer. The symmetric passive message is logged upon the arrival of the first NTP packet from the NTP peer, which is not statically configured.

Message Level Informational

Message `<server | sym_active | sym_passive> association is demobilized <ipv4 address | ipv6 address>`

Explanation Indicates the NTP server and symmetric active peer demobilization messages are logged when a user removes the NTP server or peer configuration. The NTP symmetric passive demobilization is logged when the NTP packet from the symmetric passive peer results in an error or timeout.

Message Level Informational

Syslog messages TCP

Message `TCP Local TCP exceeds burst-max burst packets, stopping for lockup seconds!`

Explanation The number of TCP SYN packets exceeds the *burst-max* threshold set by the **ip tcp burst** command. The device may be the victim of a TCP SYN DoS attack.

All TCP SYN packets will be dropped for the number of seconds specified by the *lockup* value. When the lockup period expires, the packet counter is reset and measurement is restarted.

Message Level Notification

Message `TCP Transit TCP in interface portnum exceeds num burst packets, stopping for num seconds!`

Explanation Threshold parameters for TCP transit (through) traffic have been configured on an interface, and the maximum burst size for TCP packets on the interface has been exceeded.

The *portnum* is the port number.

The first *num* is the maximum burst size (maximum number of packets allowed).

The second *num* is the number of seconds during which additional TCP packets will be blocked on the interface.

NOTE

This message can occur in response to an attempted TCP SYN attack.

Message Level Notification

Syslog messages DOT1X

Message `DOT1X security violation at port portnum , malicious mac address detected mac-address`

Explanation A security violation was encountered at the specified port number.

Message Level Warning

Message `DOT1X Port portnum , AuthControlledPortStatus change restricted`

Explanation

Message Level Warning

Message `DOT1X Port portnum port default vlan-id changes to vlan-id`

Explanation

Message Level Notification

Message `DOT1X Port portnum currently used vlan-id changes to vlan-id due to move to restricted vlan`

Explanation

Message Level Notification

Message `DOT1X issues software port up indication of Port portnum to other software applications`

Explanation The device has indicated that the specified port has been authenticated, but the actual port may not be active.

Message Level Notification

Message `DOT1X issues software port down indication of Port portnum to other software applications`

Explanation The device has indicated that the specified is no longer authorized, but the actual port may still be active.

Message Level Notification

Message `DOT1X Port portnum , AuthControlledPortStatus change authorized`

Explanation The status of the interface's controlled port has changed from unauthorized to authorized.

Message Level Informational

Message `DOT1X Port portnum , AuthControlledPortStatus change unauthorized`

Explanation The status of the interface's controlled port has changed from authorized to unauthorized.

Message Level Informational

Message `DOT1X Port portnum currently used vlan-id changes to vlan-id due to dot1x-RADIUS vlan assignment`

Explanation A user has completed 802.1X authentication. The profile received from the RADIUS server specifies a VLAN ID for the user. The port to which the user is connected has been moved to the VLAN indicated by *vlan-id*.

Message Level Informational

Message `DOT1X Port portnum currently used vlan-id is set back to port default vlan-id vlan-id`

Explanation The user connected to *portnum* has disconnected, causing the port to be moved back into its default VLAN, *vlan-id*.

Message Level Informational

Message `DOT1X Port portnum is unauthorized because system resource is not enough or the invalid information to set the dynamic assigned IP ACLs or MAC address filters`

Explanation 802.1X authentication could not take place on the port. This happened because strict security mode was enabled and one of the following occurred:

- Insufficient system resources were available on the device to apply an IP ACL or MAC address filter to the port
- Invalid information was received from the RADIUS server (for example, the Filter-ID attribute did not refer to an existing IP ACL or MAC address filter)

Message Level Informational

Message `DOT1X Not enough memory`

Explanation There is not enough system memory for 802.1X authentication to take place. Contact device Technical Support.

Message Level Debug

Syslog messages SNMP

Message `SNMP Auth. failure, intruder IP ip-addr`

Explanation A user has tried to open a management session with the device using an invalid SNMP community string.

The *ip-addr* is the IP address of the host that sent the invalid community string.

Message Level Informational

Message `SNMP read-only community | read-write community | contact | location | user | group | view | engineld | trap [host] [ value -str ] deleted | added | modified from console | telnet | ssh | web | snmp session`

Explanation A user made SNMP configuration changes through the Web, SNMP, console, SSH, or Telnet session.

[*value-str*] does not appear in the message if SNMP community or engineld is specified.

Message Level Informational

Syslog messages MPLS

Message `Deleting VLL name (ID number )at string port slot/port with peer IPv4 address ip-address`
Explanation Sent when PW traps are generated if the PW has been deleted, for example, when the pwRowStatus in the MIB has been set to destroy(6), the PW has been deleted by a non-MIB application, or due to auto-discovery process.
Message Level Informational

Message `MPLS Deleting VLL vll-name (ID vll-id )`
Explanation Sent when the specified VLL is being deleted.
Message Level Informational

Message `MPLS Deleting VLL vll-name (ID vll-id ) at {tagged | untagged} port slot/port`
Explanation Sent when the specified VLL at the specified tagged or untagged port is being deleted.
Message Level Informational

Message `MPLS Deleting VLL vll-name (ID vll-id ) with peer IPv4 address ip`
Explanation Sent when the specified VLL with the specified IPv4 peer is being deleted.
Message Level Informational

Message `VLL is down for table index number`
Explanation Sent when PW traps are generated if the VLL is down for one index.
Message Level Informational

Message `VLL is up for table index number`
Explanation Sent when PW traps are generated if the VLL is up for one index.
Message Level Informational

Message `VLLs are down for table indexes number through number`
Explanation Sent when PW traps are generated if the VLLs represented by sequential entries in the database are down.
Message Level Informational

Message `VLLs are up for table indexes number through number`
Explanation Sent when PW traps are generated if the VLLs represented by sequential entries in the database are up.
Message Level Informational

Message `VRF Port slot-port added to VRF name with updated port count number`
Explanation Sent when an MPLS Layer 3 VPN trap is generated if the state of an interface within the VRF changed from down to up.
Message Level Informational

Message `VRF Port slot-port deleted from VRF name with updated port count number`
Explanation Sent when an MPLS Layer 3 VPN trap is generated if the state of an interface within the VRF changed from down to up.
Message Level Informational

Message `MPLS Bypass LSP lspname using path <NULL> is down, Reason: LSP-down-reason`
Explanation Provides basic information about the event that triggered the LSP to go down. The *LSP-down-reason* string provides the information. The possible causes for the LSP down triggers include the following occurrences:

- LSP is disabled, deleted, or unconfigured.
- The LSP outgoing interface is down, disabled, or has bandwidth reduction leading to preemption.
- An LSP path error occurs that triggers the LSP, for example, there is no route to the destination.
- RSVP session reservation tears or times out at the downstream nodes.

- RSVP IGP Synchronization Neighbor down event occurs.
- Dynamic bypass deletion occurs due to no backup timeout or the disabling of the dynamic bypass.

If the system cannot determine the trigger, no information is not displayed in the Syslog message.

Message Level Notification

Message `MPLS Deleting VLL name (ID vc-id ) at {tagged I untagged} port portnum with peer IPv4 address ip`

Explanation

Message Level Notification

Message `MPLS LSP lspname switches to new active path pathame`

Explanation

Message Level Notification

Message `MPLS LSP lspname using path pathname is down, Reason: LSP-down-reason`

Explanation Provides basic information about the event that triggered the LSP to go down. The *LSP-down-reason* string provides the information. The possible causes for the LSP down triggers include the following occurrences:

- LSP is disabled, deleted, or unconfigured.
- The LSP outgoing interface is down, disabled, or has bandwidth reduction leading to preemption.
- An LSP path error occurs that triggers the LSP, for example, there is no route to the destination.
- RSVP session reservation tears or times out at the downstream nodes.
- RSVP IGP Synchronization Neighbor down event occurs.

If the system cannot determine the trigger, no information is not displayed in the Syslog message.

Message Level Notification

Message `MPLS LSP lspname using path pathname is up`

Explanation

Message Level Notification

Message `MPLS Standby LSP lspname using secondary path pathname is down, Reason: LSP-down-reason`

Explanation Provides basic information about the event that triggered the LSP to go down. The *LSP-down-reason* string provides the information. The possible causes for the LSP down triggers include the following occurrences:

- LSP is disabled, deleted, or unconfigured.
- The LSP outgoing interface is down, disabled, or has bandwidth reduction leading to preemption.
- An LSP path error occurs that triggers the LSP, for example, there is no route to the destination.
- RSVP session reservation tears or times out at the downstream nodes.
- RSVP IGP Synchronization Neighbor down event occurs.

If the system cannot determine the trigger, no information is not displayed in the Syslog message.

Message Level Notification

Message `MPLS VLL is down for table index n`

Explanation

Message Level Notification

Message `MPLS VLL is up for table index n`

Explanation

Message Level Notification

Message `MPLS VLLs are down for table indexes n through m`
Explanation
Message Level Notification

Message `MPLS VLLs are up for table indexes n through m`
Explanation
Message Level Notification

Message `MPLS VPLS [ID id ] peer ip is down`
Explanation Sent when a single VPLS peer is transitioning to a down state.
Message Level Notification

Message `MPLS VPLS [ID id ] peer ip is up`
Explanation Sent when a single VPLS peer is transitioning to an up state.
Message Level Notification

Message `MPLS VPLS name (ID id ) endpoint ip-address is down`
Explanation Sent when a single VPLS endpoint is transitioning to a down state.
Message Level Notification

Message `MPLS VPLS name (ID id ) endpoint ip-address is up`
Explanation Sent when a single VPLS endpoint is transitioning to an up state.
Message Level Notification

Message `MPLS VPLS for instance indices list n through m are up`
Explanation Sent when multiple VPLS instances are transitioning to an up state.
Message Level Notification

Message `MPLS VPLS for instance indices list n through m are down`
Explanation Sent when multiple VPLS instances are transitioning to a down state.
Message Level Notification

Message `MPLS VPLS peer ip associated with VC ID n is up`
Explanation Sent when a single VPLS peer is transitioning to an up state.
Message Level Notification

Message `MPLS VPLS peer ip associated with VC ID n is down`
Explanation Sent when a single VPLS peer is transitioning to a down state.
Message Level Notification

Message `MPLS VPLS peer ip associated with instances n - m list is down`
Explanation Sent when multiple VPLS instances associated with a peer are transitioning to a down state.
Message Level Notification

Message `MPLS VPLS peer ip associated with instances n - m list is up`
Explanation Sent when multiple VPLS instances associated with a peer are transitioning to an up state.
Message Level Notification

Message `MPLS VPL endpoint slot / port associated with instance indices list is down`
Explanation Sent when multiple VPLS instances associated with an endpoint is transitioning to a down state.
Message Level Notification

Message `MPLS VPL endpoint slot / port associated with instance indices list is up`
Explanation Sent when multiple VPLS instances associated with an endpoint is transitioning to an up state.
Message Level Notification

Message `MPLS VLL-Local name is down`

Explanation Sent when a single VLL-Local instance is transitioning to a down state.
Message Level Notification

Message `MPLS VLL-Local name is up`
Explanation Sent when a single VLL-Local instance is transitioning to an up state.
Message Level Notification

Message `MPLS VLL-Local for instance indices list n through m are up`
Explanation Sent when multiple VLL-Local instances are transitioning to an up state.
Message Level Notification

Message `MPLS VLL-Local for instance indices list n through m are down`
Explanation Sent when multiple VLL-Local instances are transitioning to a down state.
Message Level Notification

Message `MPLS VLL name (ID id is down`
Explanation Sent when a single VLL peer is transitioning to a down state.
Message Level Notification

Message `MPLS VLL name (ID id is up`
Explanation Sent when a single VLL peer is transitioning to an up state.
Message Level Notification

Message `MPLS VLL for instance indices list n through m are up`
Explanation Sent when multiple VLL instances are transitioning to an up state.
Message Level Notification

Message `MPLS VLL for instance indices list n through m are down`
Explanation Sent when multiple VLL instances are transitioning to a down state.
Message Level Notification

Message `Session DOWN for LSP lsp-name Reason Administratively Down`
Explanation The BFD session for the LSP specified by the *lsp-name* is down for administrative reasons.
Message Level Notification

Message `Session Up for LSP lsp-name`
Explanation The BFD session for the LSP specified by the *lsp-name* is up.
Message Level Notification

Message `Session DOWN for RSVP session session-id Reason Administratively Down`
Explanation The BFD session for the RSVP session specified by the *session-id* is down for administrative reasons.

The form of the *session-id* displayed is IPv4 tunnel endpoint or tunnel ID or extended tunnel ID. For example 10.22.22.2/3/11/11/11/1
Message Level Notification

Message `Session UP for RSVP session session-id`
Explanation The BFD session for the RSVP session specified by the *session-id* is up.

The form of the *session-id* displayed is IPv4 tunnel endpoint or tunnel ID or extended tunnel ID. For example 10.22.22.2/3/11/11/11/1
Message Level Notification

Syslog messages VRF

Message	`VRF Port portnum added to VRF name with updated port count n`
Explanation	
Message Level	Notification
Message	`VRF Port portnum deleted from VRF name with updated port count n`
Explanation	
Message Level	Notification
Message	`VRF vrf_name has been configured as management VRF.`
Explanation	Indicates that the specified VRF has been configured as a management VRF.
Message Level	Informational
Message	`VRF vrf_name has been un- configured as management VRF.`
Explanation	Indicates that the specified VRF has been removed as a management VRF.
Message Level	Informational

Syslog messages

Message	`Authentication Enabled on portnum`
Explanation	The multi-device port authentication feature was enabled on the on the specified *portnum*.
Message Level	Notification
Message	`Authenticaion Disabled on portnum`
Explanation	The multi-device port authentication feature was disabled on the on the specified *portnum*.
Message Level	Notification

Syslog messages BFD

Message	`BFD Session UP for NBR neighbor-ID on port`
Explanation	The BFD session is UP with the neighbor specified by the *neighbor-ID* on the port specified by the *port* variable.
Message Level	Notification
Message	`BFD Session DOWN for NBR neighbor-ID on port Reason Neighbor Signaled Session Down`
Explanation	The BFD session with the neighbor specified by the *neighbor-ID* on the port specified by the *port* variable is Down because the BFD neighbor has signaled the session to be down.
Message Level	Notification
Message	`BFD Session DOWN for NBR neighbor-ID on port Reason Administratively Down`
Explanation	The BFD session with the neighbor specified by the *neighbor-ID* on the port specified by the *port* variable is Down for Administrative reasons.
Message Level	Notification

Syslog messages Optics

Message `Transceiver type checking has been disabled!`
Explanation The transceiver type checking feature has been disabled. The device will continue to report incompatible transceivers through syslog messages and but will not shutdown a port that contains one.
Message Level Notification

Message `Transceiver type checking has been enabled!`
Explanation The transceiver type checking feature has been re-enabled. The feature is enabled by default and does not send the message under normal circumstances. However, if it is disabled and then re-enabled the device will send this message.
Message Level Notification

Message `Optic is not Extreme qualified ( port ) Type type-description Vendor vendor-name , Version version-num Part# part-no , Serial# serial-no`
Explanation The optic module installed in the Interface module at the port specified by the *port* variable is not Extreme qualified although the port is still operational. The Type, Vendor, Version, Part#, and Serial # of the optic module is provided.
Message Level Warning

Message `Optic is not Extreme qualified, optical monitoring is not supported ( port ) Type type-description Vendor vendor-name , Version version-num Part# part-no , Serial# serial-no`
Explanation The optic module installed in the Interface module at the port specified by the *port* variable is not Extreme qualified and will not be able to be monitored using the Optical Monitoring function. The Type, Vendor, Version, Part#, and Serial # of the optic module is provided.
Message Level Alert

Message `Optic is not capable of optical monitoring ( port ) Type type-description Vendor vendor-name , Version version-num Part# part-no , Serial# serial-no`
Explanation The optic module installed in the Interface module at the port specified by the *port* variable is not able to be monitored using the Optical Monitoring function. The Type, Vendor, Version, Part#, and Serial # of the optic module is provided.
Message Level Alert

Message `Incompatible optical trans-receiver detected on port n`
Explanation Indicates that in incompatible XFP or SFP has been installed in the port specified. A port with an incompatible optical module installed are shut down.
Message Level Alert

Syslog messages LDP

Message `MPLS LDP path vector limit mismatch for session lsrId labelSpaceId (value local vector limit ) with peer lsrId labelSpaceId (value peer vector limit )`
Explanation This notification is generated when the value of the LDP path vector limit value from the peer does not match that of the entity.
Message Level Notification

Message `MPLS LDP entity session lsrId labelSpaceId with peer lsrId labelSpaceId is up`
Explanation This notification is sent when the value of 'mplsLdpSessionState' enters the 'operational(5)' state.
Message Level Notification

Message	`MPL LDP entity session lsrId labelSpaceId with peer lsrId labelSpaceId is down`
Explanation	This notification is sent when the value of 'mplsLdpSessionState' leaves the 'operational(5)' state.
Message Level	Notification

Syslog messages DHCP

Message	`DHCPC: No DHCP service available on the network.`
Explanation	The DHCPOFFER message is not received within 16 seconds of starting the DHCP address configuration phase.
Message Level	Warning
Message	`DHCPC: Failed to renew DHCP lease on port 1/1 with IP address 10.1.1.1 mask 255.255.255.0`
Explanation	The DHCP lease cannot be renewed.
Message Level	Warning
Message	`DHCPC: Failed to configure IP address on port 1/1; with IP address 10.1.1.1,mask 255.255.255.0`
Explanation	The IP address cannot be configured without a reason.
Message Level	Warning
Message	`Failed to download image file image name`
Explanation	The image file cannot be downloaded.
Message Level	Warning
Message	`DHCPC: Failed to download configuration file image name`
Explanation	The configuration files cannot be downloaded.
Message Level	Warning

Syslog messages DHCPv6

Message	`DHCPv6: Maximum allowed 60000 delegated prefixes learned.`
Explanation	The delegated prefixes' limit has reached the maximum value at the system level.
Message Level	Warning
Message	`DHCPv6: Write to flash file to save delegated prefixes information failed.`
Explanation	Saving the delegated prefixes to flash file failed.
Message Level	Warning
Message	`DHCPv6: Maximum allowed 20000 delegated prefixes learned on interface ve 100.`
Explanation	The delegated prefixes' limit has reached the maximum value at the interface level.
Message Level	Warning

Syslog messages data integrity protection

Message	`NP CSRAM has 4 error events, exceeding configured threshold for interfaces 1/1 to 1/24.`
Explanation	A user sees this message when the CSRAM error events exceeds the configured threshold parameter for the specified port range.

Message Level Informational

Message `NP LPM has 4 error events, exceeding configured threshold for interfaces 1/1 to 1/24.`

Explanation A user receives this message when the LPM memory error events exceeds the configured threshold parameter for the specified port range.

Message Level Informational

Message `NP ingress buffer has 11 error events, exceeding configured threshold for interfaces 1/1 to 1/24.`

Explanation A user receives this message when the NP ingress buffer error events exceeds the configured threshold for the specified port range.

Message Level Informational

Message `NP egress buffer has 11 events, exceeding configured threshold for interfaces 1/1 to 1/24.`

Explanation A user receives this message when the NP egress buffer error events exceeds the configured threshold for the specified port range.

Message Level Informational

Syslog messages TCAM In-field soft repair

Message `SYSLOG: <14>Jul 23 11:02:41 sys-np-mac-224 IFSR: Soft Repair at TCAM index 0x00002fe9 of PPCR 1`

Explanation Indicates the IFSR error entry is repaired at the specified TCAM index for the PPCR 1.

Message Level Informational

Message `SYSLOG: <14>Jul 23 11:02:41 sys-np-mac-224 IFSR: Soft Repair failed at TCAM index 0x00002fe9 of PPCR 1`

Explanation Indicates the IFSR error entry failed to repair in Non-NetRoute mode at the specified TCAM index for the PPCR 1.

Message Level Informational

Message `SYSLOG: <14>Jul 23 11:02:41 sys-np-mac-224 IFSR: Soft Repair failed on PPCR 1`

Explanation Indicates the IFSR error entry failed to repair in NetRoute mode at the specified PPCR 1.

Message Level Informational

Message `SYSLOG: <14>Jul 23 11:02:41 sys-np-mac-224 IFSR: Error FIFO Overflow on PPCR 1`

Explanation Indicates that within the KBP FIFO, TCAM indices error entries are high and an FIFO overflow of entries occurred for the PPCR 1. Some error entries may have dropped.

Message Level Informational

Syslog messages NSR

Message `NSR: Successfully notified RTM6 that OSPF6 switchover complete`

Explanation OSPFv3 completes the restart process after switching over to the new master MP.

Message Level Notification