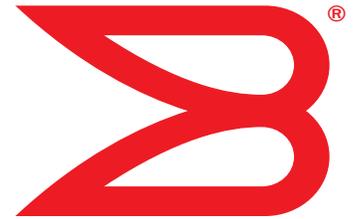


53-1003028-01
30 September, 2013



Multi-Service IronWare

Administration Guide

Supporting Multi-Service IronWare R05.6.00

BROCADE

Copyright © 2013 Brocade Communications Systems, Inc. All Rights Reserved.

ADX, AnyIO, Brocade, Brocade Assurance, the B-wing symbol, DCX, Fabric OS, ICX, MLX, MyBrocade, OpenScript, VCS, VDX, and Vyatta are registered trademarks, and HyperEdge, The Effortless Network, and The On-Demand Data Center are trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries. Other brands, products, or service names mentioned may be trademarks of their respective owners.

Notice: This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any equipment, equipment feature, or service offered or to be offered by Brocade. Brocade reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. This informational document describes features that may not be currently available. Contact a Brocade sales office for information on feature and product availability. Export of technical data contained in this document may require an export license from the United States government.

The authors and Brocade Communications Systems, Inc. shall have no liability or responsibility to any person or entity with respect to any loss, cost, liability, or damages arising from the information contained in this book or the computer programs that accompany it.

The product described by this document may contain "open source" software covered by the GNU General Public License or other open source license agreements. To find out which open source software is included in Brocade products, view the licensing terms applicable to the open source software, and obtain a copy of the programming source code, please visit <http://www.brocade.com/support/oscd>.

Brocade Communications Systems, Incorporated

Corporate and Latin American Headquarters
Brocade Communications Systems, Inc.
130 Holger Way
San Jose, CA 95134
Tel: 1-408-333-8000
Fax: 1-408-333-8101
E-mail: info@brocade.com

Asia-Pacific Headquarters
Brocade Communications Systems China HK, Ltd.
No. 1 Guanghua Road
Chao Yang District
Units 2718 and 2818
Beijing 100020, China
Tel: +8610 6588 8888
Fax: +8610 6588 9999
E-mail: china-info@brocade.com

European Headquarters
Brocade Communications Switzerland Sàrl
Centre Swissair
Tour B - 4ème étage
29, Route de l'Aéroport
Case Postale 105
CH-1215 Genève 15
Switzerland
Tel: +41 22 799 5640
Fax: +41 22 799 5641
E-mail: emea-info@brocade.com

Asia-Pacific Headquarters
Brocade Communications Systems Co., Ltd. (Shenzhen WFOE)
Citic Plaza
No. 233 Tian He Road North
Unit 1308 - 13th Floor
Guangzhou, China
Tel: +8620 3891 2000
Fax: +8620 3891 2111
E-mail: china-info@brocade.com

Document History

Title	Publication number	Summary of changes	Date
<i>Multi-Service IronWare Administration Guide</i>	53-1003028-01	Release 05.5.00c document updated with enhancements in Release 05.6.00	30 September, 2013

Contents

About This Document

In this chapter	xiii
Audience	xiii
Supported hardware and software	xiv
Supported software	xiv
Document conventions	xv
Text formatting	xv
Notes, cautions, and danger notices	xv
Notice to the reader	xvi
Related publications	xvi
Getting technical help or reporting errors	xvii

Chapter 1

Getting Started with the Command Line Interface

Logging on through the CLI	2
On-line help	2
Command completion	3
Scroll control	3
Line editing commands	3
EXEC commands	4
Global level	5
CONFIG commands	5
Accessing the CLI	8
Single user in CONFIG mode	9
Multi-user conflict during deletion of group configuration (or stanza)	9
Navigating among command levels	10
CLI command structure	10
Searching and filtering output	11
Allowable characters for LAG names	15
CLI parsing enhancement	16
Syntax shortcuts	16
Saving configuration changes	16

Chapter 2

Configuring Basic Parameters

Enabling and disabling interactivity for scripts	21
Entering system administration information	22
Configuring Simple Network Management (SNMP) traps	23

Specifying an SNMP trap receiver	23
Specifying a single trap source	24
Setting the SNMP trap holddown time	25
Disabling SNMP traps	25
Configuring SNMP ifIndex	26
On Brocade NetIron CES and Brocade NetIron CER only	26
On Brocade NetIron XMR and Brocade MLX series only	26
SNMP scalability optimization	27
Configuring SNMP throughput optimization	27
Configuring SNMP load throttling	29
Configuring optical monitoring	29
Displaying optical monitoring thresholds	30
Displaying media information	32
Optics compatibility checking	33
Disabling transceiver type checking	34
Designating an interface as the packet source	34
Configuring an interface as the source for all Telnet packets	34
Cancelling an outbound Telnet session	35
Configuring an interface as the source for all SSH packets	35
Configuring an interface as the source for all TFTP packets	35
Configuring an interface as the source for all TACACS or TACACS+ packets	36
Configuring an interface as the source for all RADIUS packets	36
Setting the system clock	37
DST “change” notice for networks using US time zones	38
Creating a command alias	38
Removing an alias	38
Displaying a list of all configured alias	39
Limiting broadcast, multicast, or unknown unicast rates	39
Limiting broadcasts	39
Limiting multicasts	39
Limiting unknown unicasts	40
Configuring CLI banners	40
Setting a message of the day banner	40
Setting a privileged EXEC CLI level banner	41
Displaying a message on the console when an incoming Telnet session is detected	41
Configuring terminal display	41
Checking the length of terminal displays	42
Enabling or disabling routing protocols	42

Displaying and modifying default settings for system parameters	43
Enabling or disabling layer 2 switching	48
Configuring static MAC addresses	49
Changing the MAC age time	50
Configuring static ARP entries	50
Configuring system max values	51
Configuring CAM size for an IPv4 multicast group	55
Configuring CAM size for an IPv6 multicast group	56
Configuring profiles with a zero-size IPv4 or IPv6 ACL	57
Maintaining system-max configuration with available system resources	57
Configuration time	58
Bootup time	58
L2 elements	60
L3 elements	60
VPLS elements	60
Miscellaneous elements	61
Monitoring dynamic memory allocation	61
Switch fabric fault monitoring	62
Displaying switch fabric information	62
Displaying switch fabric module information	64
Powering a switch fabric link on or off manually	64
Powering a switch fabric module off automatically on failure	64
Auto-tune enhancement	65
Switch fabric log messages	65
Switch fabric utilization monitoring	67
Verifying an image checksum	67
Displaying information for an interface for an Ethernet port	68
Displaying the full port name for an Ethernet interface	69
Displaying statistics information for an Ethernet port	73
Monitoring Ethernet port statistics in real time	73
Displaying recent traffic statistics for an Ethernet port	78
Configuring SNMP to revert ifType to legacy values	79
Configuring snAgentConfigModuleType to return original values	80
Preserving interface statistics in SNMP	80
Disabling CAM table entry aging	81
Data integrity protection	81

	Configuring Detection Parameters	81
	Showing Status	82
	Commands	83
Chapter 3	Telemetry Solutions	
	About telemetry solutions.	89
	Limitations	89
	Configuration examples	90
	Configuration example 1	90
	Configuration example 2	91
	Configuration example 3	93
	Configuring	94
Chapter 4	Remote Network Monitoring	
	Basic management	97
	Viewing system information	97
	Viewing configuration information	98
	Viewing port statistics	98
	Viewing STP statistics	98
	Clearing statistics.	98
	RMON support.	98
	Statistics (RMON group 1).	99
	History (RMON group 2).	101
	Alarm (RMON group 3).	102
	Event (RMON group 9).	102
Chapter 5	Continuous System Monitor	
	Continuous system monitor overview	104
	Event monitoring.	104
	Event monitoring overview	104
	Event types.	105
	Displaying event information.	105
	Histogram information	106
	Histogram information overview	106
	Displaying CPU histogram information	107
	Displaying buffer histogram information	109
	Displaying memory histogram information	112
	NP memory error monitoring	114
	NP memory error monitoring overview	114
	NP memory error monitoring: basic configuration.	114
	Port CRC error monitoring test.	116
	Port CRC error monitoring overview	116
	Port CRC error monitoring: basic configuration	116
	Commands	118

Chapter 6

Operations, Administration, and Maintenance (OAM)

IEEE 802.1ag Connectivity Fault Management (CFM)	143
Ethernet OAM capabilities	143
IEEE 802.1ag purpose	143
IEEE 802.1ag provides hierarchical network management . . .	145
Mechanisms of Ethernet IEEE 802.1ag OAM	145
Fault detection (continuity check message)	145
Fault verification (Loopback messages)	146
Fault isolation (Linktrace messages)	146
Configuring IEEE 802.1ag CFM	147
Enabling or disabling CFM	147
Creating a Maintenance Domain	147
Setting Maintenance Domain parameters	148
Creating Maintenance Associations	148
Tag-type configuration	149
Configuring a CCM interval for a Maintenance Association (MA)	149
Configuring local ports	150
Configuring Remote MEPs	151
Setting the Remote Check Start-Delay	151
Specifying MIP creation policy	151
Y.1731 performance management	152
About Y.1731	152
Y. 1731 show commands	154
CFM monitoring and show commands	156
Sending linktrace messages	156
Sending loopback messages	156
Displaying CFM configurations	158
Displaying connectivity statistics	160
Sample configuration for a customer's domain	160
Configuring CFM using Provider Bridges	162
Displaying the connectivity status in a customer's domain	167
Sample configuration for a customer domain using MPLS VLL	168
Achieving end-to-end connectivity between CE1 and CE2	169
Monitoring the status of devices in a VPLS network in a Provider's Maintenance Domain	178
Verifying connectivity in a VPLS network using IEEE 802.1ag	181
Verifying connectivity in a VPLS network using IEEE 802.1ag Loopback	183
Support for IEEE 802.1ag CFM for Provider Bridges (PB) and Provider Backbone Bridges (PBB)	185
IEEE 802.3ah EFM-OAM	186
EFM- OAM protocol	187
Process overview	188
Link monitoring process	189

Enabling and disabling EFM-OAM	190
Enabling an interface to accept remote loopback.	191
Display information	192
Ping	196
Executing ping	197
Executing ping VRF	198
Executing ping IPv6	198
Trace route	200
Executing traceroute	200
Executing traceroute VRF	200
Executing traceroute IPv6	201
Trace-I2 protocol	201
IPv6 Traceroute over an MPLS network.	203
LSP ping and traceroute	208
Overview	208
LSP ping operation.	208
LSP traceroute operation	208
MPLS echo request	208
MPLS echo reply	209
LSP ping TLVs	210
LSP FEC types	210
Redundant RSVP LSPs	210
One-to-one Fast ReRoute (FRR) LSPs	211
FRR bypass LSPs	211
Transit-originated detour	211
LSP reoptimization.	211
PHP behavior	212
Using the LSP ping and Traceroute commands.	212
Displaying LSP ping and traceroute statistics	217
CFM monitoring for ISID	219
Configuring CFM monitoring for ISID	219
Link MA	223
Port status TLV	227
Remote Defect Indication	229
Frame Loss Measurement	231
LMM over VLAN	231
LMM over VPLS	231
Configuration considerations and limitations	231
Supported configurations	232
LMM configurations common for VLAN and VPLS.	233
Configuration examples.	235
Syslog messages	237
One-way Delay Measurement.	238
Configuration considerations	238
One-way Delay Measurement	239
One-way Delay Measurement transmission.	239
One-way Delay Measurement reception.	239
Use Cases.	239
Supported configurations	240

Configuration procedure	241
Configuration examples	245
Show commands	248
Syslog messages	249
Synthetic loss measurement	250
Configuration considerations	250
Commands	251
Configuration examples	252
Show commands	257
Syslog messages	259

Chapter 7

Network Time Protocol

Network Time Protocol (NTP) overview	261
How NTP works	263
NTP server	263
NTP client	263
NTP peer	263
NTP broadcast server	263
NTP broadcast client	264
Synchronizing time	265
Configuring NTP	265
Changing to the NTP mode	265
Configuring the NTP client	267
Configuring the NTP peer	268
Configuring NTP on an interface	268
Show commands	270
Displaying NTP status	270
Displaying NTP associations	271
Displaying NTP associations details	272
Configuration Examples	274

Chapter 8

Network Configuration Protocol

NETCONF protocol introduction	277
Platforms	278
Related documentation	278
NETCONF in client/server architecture	278
RPC request	279
RPC reply	279
RPC and error handling	280
CLI and SSH subsystem	280
Recommendations for NETCONF	281
Basic NETCONF operations	281
Initial connection	281
<get> operation	283
<get-config> operation	286
<edit-config> operation	295
Closing sessions	297
NETCONF commands and specifications	297

	Data models and mapping	301
Chapter 9	Foundry Direct Routing and CAM Partition Profiles for the NetIron XMR and the Brocade MLX Series	
	Configuring Density Mode for the 2x100G CAM	303
	Configuring IPv6 host CAM mode	303
	Configuring IPv6 host drop CAM limit	304
	CAM partition profiles	304
	Supernet CAM partition sharing	312
	Displaying CAM partition	312
	Displaying CAM Partition for IPv6 VPN	316
	Output from show CAM partition usage command	316
	Displaying CAM information	319
	Displaying IPv6 VPN CAM information	322
	Show cam v6acl	323
	Displaying IPv6 host drop CAM limit	323
	Show IFL CAM ISID partition	324
	Configuring CAM partition size	324
	CAM overflow logging	325
	Configuring minimum logging interval and threshold value	325
Appendix A	Using Syslog	
	Displaying Syslog messages	329
	Configuring the Syslog service	330
	Displaying the Syslog configuration	330
	Configuring an encrypted syslog server	334
	Displaying the configured server connections	335
	Current active SSL syslog server: 10.25.105.201:60514	336
	Ascending or descending option for show log command	336
	Disabling or re-enabling Syslog	336
	Specifying a Syslog server	337
	Specifying an additional Syslog server	337
	Disabling logging of a message level	337
	Changing the number of entries for the local buffer	338
	Changing the log facility	338
	Displaying the interface name in Syslog messages	339
	Clearing the Syslog messages from the local buffer	340
	Logging all CLI commands to Syslog	340
	Syslog messages	340
Appendix B	Global and Address Family Configuration Levels	
	Accessing the address family configuration level	383
	Backward compatibility for existing BGP4 and IPv4 IS-IS configurations	384

	Global BGP4 commands and BGP4 unicast route commands	384
Appendix C	Commands That Require a Reload	
Appendix D	NIAP-CCEVS	
	NIAP-CCEVS certified Brocade equipment and Ironware releases	389
	Web management access to NIAP-CCEVS certified Brocade equipment	390
	Warning: local user password changes	391
Appendix E	Acknowledgements	
	Cryptographic software	393
	MPL 1.1	393
	OpenSSL license	393
	Cryptographic software	394
Appendix F	NP Memory Errors	

About This Document

In this chapter

- Audience. xiii
- Supported hardware and software. xiv
- Document conventions xv
- Notice to the reader xvi
- Related publications xvi
- Getting technical help or reporting errors xvii

Audience

This document is designed for system administrators with a working knowledge of Layer 2 and Layer 3 switching and routing.

If you are using a Brocade device, you should be familiar with the following protocols if applicable to your network – IP, RIP, OSPF, BGP, ISIS, IGMP, PIM, MPLS, and VRRP.

Supported hardware and software

The following hardware platforms are supported by this release of this guide:

TABLE 1 Supported devices

Brocade NetIron XMR Series	Brocade MLX Series	NetIron CES 2000 and NetIron CER 2000 Series
Brocade NetIron XMR 4000	Brocade MLX-4	Brocade NetIron CES 2024C
Brocade NetIron XMR 8000	Brocade MLX-8	Brocade NetIron CES 2024F
Brocade NetIron XMR 16000	Brocade MLX-16	Brocade NetIron CES 2048C
Brocade NetIron XMR 32000	Brocade MLX-32	Brocade NetIron CES 2048CX
	Brocade MLXe-4	Brocade NetIron CES 2048F
	Brocade MLXe-8	Brocade NetIron CES 2048FX
	Brocade MLXe-16	Brocade NetIron CER 2024C
	Brocade MLXe-32	Brocade NetIron CER-RT 2024C
		Brocade NetIron CER 2024F
		Brocade NetIron CER-RT 2024F
		Brocade NetIron CER 2048C
		Brocade NetIron CER-RT 2048C
		Brocade NetIron CER 2048CX
		Brocade NetIron CER-RT 2048CX
		Brocade NetIron CER 2048F
		Brocade NetIron CER-RT 2048F
		Brocade NetIron CER 2048FX
		Brocade NetIron 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 *Multi-Service IronWare R05.6.00 Release Notes*.

Document conventions

This section describes text formatting conventions and important notice formats used in this document.

Text formatting

The narrative-text formatting conventions that are used are as follows:

bold text	Identifies command names
	Identifies the names of user-manipulated GUI elements
	Identifies keywords
	Identifies text to enter at the GUI or CLI
<i>italic text</i>	Provides emphasis
	Identifies variables
	Identifies document titles
<code>code text</code>	Identifies CLI output

For readability, command names in the narrative portions of this guide are presented in bold: for example, **show version**.

Notes, cautions, and danger notices

The following notices and statements are used in this manual. They are listed below in 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.



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.

Notice to the reader

This document may contain references to the trademarks of the following corporations. These trademarks are the properties of their respective companies and corporations.

These references are made for informational purposes only.

Corporation	Referenced Trademarks and Products
Microsoft Corporation	Internet Explorer
Mozilla Corporation	Mozilla Firefox
Sun Microsystems	Java Runtime Environment

Related publications

For the latest edition of these documents, which contain the most up-to-date information, see Documentation at <http://www.brocade.com/ethernetproducts>

- *Multi-Service IronWare Administration Guide*
- *Multi-Service IronWare Security Configuration Guide*
- *Multi-Service IronWare Switching Configuration Guide*
- *Multi-Service IronWare Routing Configuration Guide*
- *Multi-Service IronWare Traffic Management Configuration Guide*
- *Multi-Service IronWare Multicast Configuration Guide*
- *Multi-Service IronWare Multiprotocol Label Switch (MPLS) Configuration Guide*
- *Multi-Service IronWare Software Defined Networking (SDN) Guide*
- *Brocade MLX Series and NetIron Family YANG Guide*
- *Brocade MLX Series and NetIron XMR Series Diagnostic Reference*
- *Unified IP MIB Reference*
- *Multi-Service IronWare Software Upgrade Procedures for Brocade MLX Series and NetIron Family devices*
- *Brocade MLXe Series Installation Guide*
- *Brocade MLX Series and Brocade NetIron XMR Installation Guide*
- *Brocade NetIron CES 2000 Series and Brocade NetIron CER 2000 Series Hardware Installation Guide*

Getting technical help or reporting errors

To contact Technical Support, go to <http://www.brocade.com/services-support/index.page> for the latest e-mail and telephone contact information.

In this chapter

Getting Started with the Command Line Interface

Table 2 displays the individual devices and the command line features they support.

TABLE 2 Supported command line features

Features supported	Brocade NetronXMR Series	Brocade MLX Series	Brocade Netron CES 2000 Series BASE package	Brocade Netron CES 2000 Series ME_PREM package	Brocade Netron CES 2000 Series L3_PREM package	Brocade Netron CER 2000 Series Base package	Brocade Netron CER 2000 Series Advanced Services package
On-Line Help	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Command Completion	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Scroll Control	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Line Editing Commands	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Accessing the CLI	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Single User in CONFIG Mode	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Multi-User Conflict During Deletion Of Group Configuration (Or Stanza)	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Searching and Filtering Output	Yes	Yes	Yes	Yes	Yes	Yes	Yes
CLI Parsing	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Syntax Shortcuts	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Web Management Interface	Yes	Yes	No	No	No	No	No

This chapter presents information to help you become familiar with the Brocade command line interface (CLI).

As with other devices, you can manage a Brocade using any of the following applications:

- **Command Line Interface (CLI)** – a text-based interface accessible directly from a PC or terminal attached to the management module's serial (Console) port or 10BaseT/100BaseTX Ethernet (management) port, or from a Telnet connection to the PC or terminal.

1 Logging on through the CLI

- **Web Management Interface** – a GUI-based management interface accessible through an HTTP (web browser) connection.

NOTE

The following interface cards are not supported by the front panel of the Web Management Interface: BR-MLX-100Gx2-X, NI-MLX-1Gx48-T, BR-MLX-10GX4-X-ML

- Brocade Network Advisor – an optional SNMP-based standalone GUI application.

This user guide describes how to configure the features using the CLI.

NOTE

This user guide assumes that an IP address and default gateway have been assigned to the Brocade device when it was installed. If you need to assign an IP address or default gateway to the device, refer to the Brocade Installation guides.

Logging on through the CLI

After an IP address is assigned to the Brocade device's management port, you can access the CLI through a PC or terminal attached to the management module's serial (Console) port or 10BaseT/100BaseTX Ethernet (management) port, or from a Telnet or SSH connection to the PC or terminal.

You can initiate a local Telnet, SSH or SNMP connection by specifying the management port's IP address.

The commands in the CLI are organized into the following levels:

- **User EXEC** – Lets you display information and perform basic tasks such as pings and traceroutes.
- **Privileged EXEC** – Lets you use the same commands as those at the User EXEC level plus configuration commands that do not require saving the changes to the system-config file.
- **CONFIG** – Lets you make configuration changes to the device. To save the changes across software reloads and system resets, you need to save them to the system-config file. The CONFIG level contains sub-levels for individual ports, for VLANs, for routing protocols, and other configuration areas.

NOTE

By default, the Brocade devices have all management access disabled, except for console port management. To create access, you must configure Enable passwords or local user accounts, or you can configure the device to use a RADIUS or TACACS or TACACS+ server for authentication.

On-line help

To display a list of available commands or command options, enter “?” or press Tab. If you have not entered part of a command at the command prompt, all the commands supported at the current CLI level are listed. If you enter part of a command, then enter “?” or press Tab, the CLI lists the options you can enter at this point in the command string.

If you enter an invalid command, a message appears indicating the command was unrecognized.

Example

```
Brocade(config)# router ip
Unrecognized command
```

Command completion

The CLI supports command completion, so you do not need to enter the entire name of a command or option. As long as you enter enough characters of the command or option name to avoid ambiguity with other commands or options, the CLI understands what you are typing.

Scroll control

By default, the CLI uses a page mode to paginate displays that are longer than the number of rows in your terminal emulation window. For example, if you display a list of all the commands at the global CONFIG level but your terminal emulation window does not have enough rows to display them all at once, the page mode stops the display and lists your choices for continuing the display.

Example

```
aaa
access-list
all-client
arp
banner
base-mac-addr
boot
some lines omitted for brevity...

default-vlan-id
enable
enable-acl-counter
end
exit
--More--, next page: Space, next line: Return key, quit: Control-c
```

The software provides the following scrolling options:

- Press the Space bar to display the next page (one screen at a time).
- Press the Return or Enter key to display the next line (one line at a time).
- Press Ctrl-C cancel the display.

Line editing commands

The CLI supports the following line editing commands. To enter a line-editing command, use the CTRL-key combination for the command by pressing and holding the CTRL key, then pressing the letter associated with the command.

TABLE 3 CLI line-editing commands

Ctrl-key combination	Description
Ctrl-A	Moves to the first character on the command line.
Ctrl-B	Moves the cursor back one character.

TABLE 3 CLI line-editing commands (Continued)

Ctrl-key combination	Description
Ctrl-C	Escapes and terminates command prompts and ongoing tasks (such as lengthy displays), and displays a fresh command prompt.
Ctrl-D	Deletes the character at the cursor.
Ctrl-E	Moves to the end of the current command line.
Ctrl-F	Moves the cursor forward one character.
Ctrl-K	Deletes all characters from the cursor to the end of the command line.
Ctrl-L; Ctrl-R	Repeats the current command line on a new line.
Ctrl-N	Enters the next command line in the history buffer.
Ctrl-P	Enters the previous command line in the history buffer.
Ctrl-U; Ctrl-X	Deletes all characters from the cursor to the beginning of the command line.
Ctrl-W	Deletes the last word you typed.
Ctrl-Z	Moves from any CONFIG level of the CLI to the Privileged EXEC level; at the Privileged EXEC level, moves to the User EXEC level.

EXEC commands

There are two different levels of EXEC commands, the *User Level* and the *Privileged Level*.

User level

The User level commands are at the top of the CLI hierarchy. These are the first commands that you have access to when connected to the device through the CLI. When you first connect to the Brocade, you will see the User level prompt, similar to the following.

```
Brocade>
```

The “Brocade” part of the prompt is configurable. Your system may display a different string.

At this level, you can view basic system information and verify connectivity but cannot make any changes to the device configuration. To make changes to the configuration, you must move to other levels of the CLI hierarchy, such as the Privileged EXEC level.

Privileged EXEC level

Commands at the Privileged EXEC level enable you to transfer and store software images and configuration files between the network and the system, and review the configuration.

You reach this level by entering the **enable** [password] or **enable** username password at the User EXEC level.

Example

```
Brocade> enable
```

or

```
Brocade> enable user1 mypassword
```

After entering the enable command, you see the following prompt.

```
Brocade#
```

The prompt indicates that you are at the Privilege EXEC level.

When you are at the Privilege EXEC level, you can enter commands that are available at that level. It is also at this level where you enter the **configure terminal** command to Global Configuration level.

Global level

The global CONFIG level allows you to globally apply or modify parameters for ports on the device. You reach this level by entering **configure terminal** at the privileged EXEC level.

```
Brocade> enable  
Brocade# configure terminal
```

The prompt changes to the Global Configuration level.

```
Brocade(config)#
```

NOTE

For configuration files which are copied to device running, or startup config via TFTP/SCP, entering a blank comment line or ! (exclamation mark denotes a comment line) followed only by blank spaces, in any of the global config sublevels, resets the mode to global config level.

CONFIG commands

CONFIG commands modify the configuration of a Brocade device. When you are at the Global Configuration level, you can enter commands to configure the features in a Brocade. This section describes the CONFIG CLI levels.

Redundancy level

This redundancy level allows you to configure redundancy parameters for redundant management modules. You reach this level by entering the **redundancy** command at the global CONFIG level.

Interface level

The interface level allows you to assign or modify specific port parameters on a specific port. You reach this level by entering the following at the global CONFIG level:

- **interface ethernet** *slot/port*
- **interface loopback** *num*
- **interface management** *portnum*
- or **interface ve** *num*
- **interface tunnel** *tunnel_id*
- **interface group-ve** *vlan_group_id*

LAG level

The LAG level allows you to change parameters for statically-configured LAG groups. You reach this level by entering a **LAG** command with the appropriate port parameters.

Router RIP level

The RIP level allows you to configure parameters for the RIP routing protocol. You reach this level by entering the **router rip** command at the global CONFIG level.

Router OSPF level

The OSPF level allows you to configure parameters for the OSPF routing protocol. You reach this level by entering the **router ospf** command at the global CONFIG level.

BGP level

The BGP level allows you to configure Border Gateway Protocol version 4 (BGP4) features. You reach this level by entering the **router bgp** command at the global CONFIG level.

Global BGP and BGP4 unicast address family level

The global BGP and BGP4 unicast address family levels are present only on devices that support IPv6. The global BGP level allows you to configure the BGP routing protocol. The BGP4 unicast address family level allows you to configure a BGP4 unicast route. For backward compatibility, you can currently access BGP4 unicast address family commands at both global BGP configuration and BGP4 unicast address family configuration levels. Therefore, the global BGP and BGP4 unicast address family commands are documented together.

You reach the global BGP level by entering the **router bgp** command at the global CONFIG level. You reach the BGP4 unicast address family level by entering the **address-family ipv4 unicast** command at the global BGP level.

BGP4 multicast address family level

The BGP4 multicast address family level allows you to configure BGP4 multicast routes. You reach this level by entering the **address-family ipv4 multicast** command at the global BGP, BGP4 unicast address family, or IPv6 BGP unicast address family levels.

Router DVMRP level

The DVMRP level allows you to configure details for the DVMRP multicast protocol. You reach this level by entering the **router dvmrp** command at the global CONFIG level.

Router PIM level

The PIM level allows you to configure parameters for the Protocol Independent Multicast (PIM) routing protocol. You reach this level by entering the **router pim** command at the global CONFIG level.

Route Map level

The Route Map level allows you to configure parameters for a BGP4 route map. You reach this level by entering the **route-map name** command at the global CONFIG level.

Router VRRP level

The VRRP level allows you to configure parameters for the Virtual Router Redundancy Protocol (VRRP). You reach this level by entering the **router vrrp** command at the global CONFIG level, then entering the **ip vrrp vrid num** command at the interface configuration level.

Router VRRPE level

The VRRPE level allows you to configure parameters for VRRP Extended. You reach this level by entering the **router vrrp-extended** command at the global CONFIG level, then entering the **ip vrrp-extended vrid num** command at the interface configuration level.

VLAN level

Policy-based VLANs allow you to assign VLANs to a protocol, port, or 802.1q tags.

You reach this level by entering the **vlan vlan-id** command at the Global CONFIG Level.

Ethernet Service Instance (ESI) level

Ethernet Service Instance (ESI) allow you to assign an ESI to a protocol, or port.

Metro ring level

Metro rings provide Layer 2 connectivity and fast failover in ring topologies.

You reach this level by entering the **metro-ring ring-id** command at the VLAN CONFIG Level.

VSRP level

The VSRP level allows you to configure parameters for the Virtual Switch Redundancy Protocol (VSRP). You reach this level by entering the **vsrp vrid num** command at the VLAN configuration level, then entering the **vsrp vrid num** command at the VLAN configuration level.

Topology group level

A topology group enables you to control the Layer 2 protocol configuration and Layer 2 state of a set of ports in multiple VLANs based on the configuration and states of those ports in a single master VLAN. One instance of the Layer 2 protocol controls all the VLANs.

You reach this level by entering the **topology-group group-id** command at the Global CONFIG Level.

802.1X port security level

The 802.1X port security level allows you to configure the 802.1X port security. You reach this level by entering the **dot1x-enable** command at the Global level.

MAC port security level

The MAC port security level allows you to configure the port security feature. You reach this level by entering the **port security** command at the at the Global or Interface levels.

Accessing the CLI

The CLI can be accessed through both serial and Telnet connections. For initial log on, you must use a serial connection. Once an IP address is assigned, you can access the CLI through Telnet.

Once connectivity to the device is established, you will see the a prompt.

```
Brocade>
```

When accessing the CLI through Telnet, you maybe prompted for a password. By default, the password required is the password you enter for general access at initial setup. You also have the option of assigning a separate password for Telnet access with the **enable telnet password password** command, found at the Global Level.

At initial log on, all you need to do is type **enable** at the prompt, then press Return. You only need to enter a password after a permanent password is entered at the Global CONFIG Level of the CLI.

To reach the Global CONFIG Level, the uppermost level of the CONFIG commands, enter the following commands

Brocade> enable	User Level commands
Brocade# configure terminal	Privileged Level-EXEC commands
Brocade(config)#	Global Level-CONFIG commands

You can then reach all other levels of the CONFIG command structure from this point.

The CLI prompt will change at each level of the CONFIG command structure, to easily identify the current level.

```
Brocade> User Level EXEC Command
Brocade# Privileged Level EXEC Command
Brocade(config)#Global Level CONFIG Command
Brocade(config-if-e10000-5/1)#Interface Level CONFIG Command
Brocade(config-lbif-1)#Loopback Interface CONFIG Command
Brocade(config-ve-1)#Virtual Interface CONFIG Command
Brocade(config-trunk-4/1-4/8)# trunk group CONFIG Command
Brocade(config-if-e10000-tunnel)#IP Tunnel Level CONFIG Command
Brocade(config-bgp-router)#BGP Level CONFIG Command
Brocade(config-dvmrp-router)#DVMRP Level CONFIG Command
Brocade(config-ospf-router)#OSPF Level CONFIG Command
Brocade(config-isis-router)#IS-IS Level CONFIG Command
Brocade(config-pim-router)#PIM Level CONFIG Command
Brocade(config-redundancy)#Redundant Management Module CONFIG Command
Brocade(config-rip-router)#RIP Level CONFIG Command
Brocade(config-port-80)#Application Port CONFIG Command
Brocade(config-bgp-routemap Map_Name)#Route Map Level CONFIG Command
Brocade(config-vlan-1)#VLAN Port-based Level CONFIG Command
Brocade(config-vlan-ataalk-proto)#VLAN Protocol Level CONFIG Command
```

NOTE

The CLI prompt at the interface level includes the port speed. The speed is one of the following:

Brocade(config-if-e100-5/1)# - The interface is a 10/100 port.

Brocade(config-if-e1000-5/1)# - The interface is a Gigabit port.

For simplicity, the port speeds sometimes are not shown in example Interface level prompts in this manual.

Single user in CONFIG mode

By default, more than one user can enter the CONFIG mode of a device CLI, which is accessed through the **configure terminal** command. While in CONFIG mode, users can override another user's configuration changes.

You can configure a device to allow only one user to be in CONFIG mode at any one time. Other users who try to enter that mode in will be denied. To allow only one user to enter CONFIG mode, enter the following command.

```
Brocade#configure terminal
Brocade(config)# single-config-user
Brocade(config)# write memory
```

Syntax: [no] single-config-user

After the **single-config-user** command is issued, the device will not allow more than one user to enter CONFIG mode. However, if you run the command while more than one user is in CONFIG mode, the other users continue to be in CONFIG mode and can potentially override each other's configuration changes. Only users who try to enter the CONFIG mode after the command is issued are prevented from entering CONFIG mode. If a user is already in that mode and another user tries to enter CONFIG mode after the **single-config-user** command is issued, the following error is displayed.

```
Brocade#configure terminal
Single user config mode is being enforced. Config mode is being used by
<session-type> session.
```

where *session-type* can be one of the following:

- console
- telnet *number*
- SSH *number*

Multi-user conflict during deletion of group configuration (or stanza)

By default, a user may delete a group configuration, even if another user is simultaneously in that mode. You can disable this feature by issuing the **enable multi-user-mode-deletion** command.

To allow only one user to delete group configurations, enter the following command.

```
Brocade#configure terminal
Brocade(config)# enable multi-user-mode-deletion
Brocade(config)# write memory
```

1 CONFIG commands

When a user attempts to delete a group configuration from the CLI, and another user is already within that group configuration, the user who tries to delete a group configuration in that mode will be denied and will receive the following error message.

Session 1:

```
Brocade(config)# vlan 10
Brocade(config-vlan-10)#
```

Session 2:

```
Brocade(config)# no vlan 10
"Error: Cannot undo the configuration as {console|telnet|SSH} session is
using this mode."
```

Syntax: [no] enable multi-user-mode-deletion

Use the **no** form of this command will allow multiple users the ability to delete group configurations.

NOTE

This feature will not work on commands that are issued from the WEB management and the SNMP management.

Navigating among command levels

To reach other CLI command levels, you need to enter certain commands. At each level there is a launch command that allows you to move either up or down to the next level.

CLI command structure

Many CLI commands may require textual or numeral input as part of the command.

Required or optional fields

These fields are either required or optional depending on how the information is bracketed. For clarity, a few CLI command examples are explained below.

Example

Syntax: [no] deny redistribute *value* all | bgp | rip | static address *ip-addr ip-mask*
[match-metric *value* | set-metric *value*]

When an item is in italics, the information requested is a variable and required.

When an item is not bracketed with “[]” symbols, the item is a required keyword.

When an item is bracketed with “[]” symbols, the information requested is optional.

Optional fields

When two or more options are separated by a vertical bar, “ | “, you must enter one of the options as part of the command.

Example

Syntax: priority normal | high

For example, the “normal | high” entry in the Syntax above means that priority can be either priority normal or priority high. The command in the syntax above requires that you enter either normal or high as part of the command.

List of available options

To get a quick display of available options at a CLI level or for the next option in a command string, enter a question mark (?) at the prompt or press TAB.

Example

To view all available commands at the user EXEC level, enter the following or press TAB at the User EXEC CLI level.

```
Brocade> ?
enable
exit
fastboot
ping
show
stop-trace-route
traceroute
```

You also can use the question mark (?) with an individual command, to see all available options or to check context.

Example

Enter the following to view possible **copy** command options.

```
Brocade# copy ?
  flash
  running-config
  startup-config
  tftp
Brocade# copy flash ?
  tftp
```

Searching and filtering output

You can filter CLI output from **show** commands and at the **-More-** prompt. You can search for individual characters, strings, or construct complex regular expressions to filter the output.

Searching and filtering output from show commands

You can filter output from **show** commands to display lines containing a specified string, lines that do not contain a specified string, or output starting with a line containing a specified string. The search string is a regular expression consisting of a single character or string of characters. You can use special characters to construct complex regular expressions. Refer to [“Using special characters in regular expressions”](#) on page 14 for information on special characters used with regular expressions.

Displaying lines containing a specified string

The following command filters the output of the **show interface** command for port 3/11 so it displays only lines containing the word “Internet”. This command can be used to display the IP address of the interface.

```
Brocade# show interface e 3/11 | include Internet
  Internet address is 192.168.1.11/24, MTU 1518 bytes, encapsulation ethernet
```

Syntax: *show-command* | **include** *regular-expression*

NOTE

The vertical bar (|) is part of the command.

Note that the regular expression specified as the search string is case sensitive. In the example above, a search string of “Internet” would match the line containing the IP address, but a search string of “internet” would not.

Displaying lines that do not contain a specified string

The following command filters the output of the **show who** command so it displays only lines that do not contain the word “closed”. This command can be used to display open connections to the device.

```
Brocade# show who | exclude closed
Console connections:
    established
    you are connecting to this session
    2 seconds in idle
Telnet connections (inbound):
  1    established, client ip address 192.168.9.37
      27 seconds in idle
Telnet connection (outbound):
SSH connections:
```

Syntax: *show-command | exclude regular-expression*

Displaying lines starting with a specified string

The following command filters the output of the **show who** command so it displays output starting with the first line that contains the word “SSH”. This command can be used to display information about SSH connections to the Brocade.

```
Brocade# show who | begin SSH
SSH connections:
  1    established, client ip address 192.168.9.210
      7 seconds in idle
  2    closed
  3    closed
  4    closed
  5    closed
```

Syntax: *show-command | begin regular-expression*

Searching and filtering output at the --More-- prompt

The --More-- prompt is displayed when output extends beyond a single page. From this prompt, you can press the Space bar to display the next page, the Return or Enter key to display the next line, or Ctrl-C or Q to cancel the display. You can also search and filter output from this prompt.

Example

```

Brocade# ?
  append          Append one file to another
  attrib          Change file attribute
  boot            Boot system from bootp/tftp server/flash image
  cd              Change current working directory
  chdir           Change current working directory
  clear           Clear table/statistics/keys
  clock           Set clock
  configure       Enter configuration mode
  copy            Copy between flash, tftp, config/code
  cp              Copy file commands
  debug           Enable debugging functions (see also 'undebug')
  delete          Delete file on flash
  dir             List files
  dm              test commands
  dot1x           802.1X
  erase           Erase image/configuration files from flash
  exit            Exit Privileged mode
  fastboot        Select fast-reload option
  force-sync-standby Sync active flash (pri/sec/mon/startup config/lp images)
                  to standby
  format          Format Auxiliary Flash card
  hd              Hex dump
  ipc             IPC commands
--More--, next page: Space, next line: Return key, quit: Control-c

```

At the --More-- prompt, you can press the forward slash key (/) and then enter a search string. The device displays output starting from the first line that contains the search string, similar to the **begin** option for **show** commands.

Example

```

--More--, next page: Space, next line: Return key, quit: Control-c
/telnet

```

The results of the search are displayed.

```

searching...
telnet          Telnet by name or IP address
terminal        Change terminal settings
traceroute      TraceRoute to IP node
undelete        Recover deleted file
whois           WHOIS lookup
write           Write running configuration to flash or terminal

```

To display lines containing only a specified search string (similar to the **include** option for **show** commands) press the plus sign key (+) at the --More-- prompt and then enter the search string.

```

--More--, next page: Space, next line: Return key, quit: Control-c
+telnet

```

The filtered results are displayed.

```

filtering...
telnet          Telnet by name or IP address

```

To display lines that do not contain a specified search string (similar to the **exclude** option for **show** commands) press the minus sign key (-) at the --More-- prompt and then enter the search string.

1 CONFIG commands

```
--More--, next page: Space, next line: Return key, quit: Control-c  
-telnet
```

The filtered results are displayed.

```
filtering...  
sync-standby      Sync active flash (pri/sec/mon/startup config/lp images)  
                  to standby if different  
terminal          Change terminal settings  
traceroute        TraceRoute to IP node  
undelete          Recover deleted file  
whois             WHOIS lookup  
write             Write running configuration to flash or terminal
```

As with the commands for filtering output from **show** commands, the search string is a regular expression consisting of a single character or string of characters. You can use special characters to construct complex regular expressions. Refer to the next section for information on special characters used with regular expressions.

Using special characters in regular expressions

You use a regular expression to specify a single character or multiple characters as a search string. In addition, you can include special characters that influence the way the software matches the output against the search string. These special characters are listed in the following table.

TABLE 4 Special characters for regular expressions

Character	Operation
.	The period matches on any single character, including a blank space. For example, the following regular expression matches “aaz”, “abz”, “acz”, and so on, but not just “az”: a.z
*	The asterisk matches on zero or more sequential instances of a pattern. For example, the following regular expression matches output that contains the string “abc”, followed by zero or more Xs: abcX*
+	The plus sign matches on one or more sequential instances of a pattern. For example, the following regular expression matches output that contains “de”, followed by a sequence of “g”s, such as “deg”, “degg”, “deggg”, and so on: deg+
?	The question mark matches on zero occurrences or one occurrence of a pattern. For example, the following regular expression matches output that contains “dg” or “deg”: de?g NOTE: Normally when you type a question mark, the CLI lists the commands or options at that CLI level that begin with the character or string you entered. However, if you enter Ctrl-V and then type a question mark, the question mark is inserted into the command line, allowing you to use it as part of a regular expression.
^	A caret (when not used within brackets) matches on the beginning of an input string. For example, the following regular expression matches output that begins with “deg”: ^deg
\$	A dollar sign matches on the end of an input string. For example, the following regular expression matches output that ends with “deg”: deg\$

TABLE 4 Special characters for regular expressions (Continued)

Character	Operation
<code>_</code>	<p>An underscore matches on one or more of the following:</p> <ul style="list-style-type: none"> • <code>,</code> (comma) • <code>{</code> (left curly brace) • <code>}</code> (right curly brace) • <code>(</code> (left parenthesis) • <code>)</code> (right parenthesis) • The beginning of the input string • The end of the input string • A blank space <p>For example, the following regular expression matches on “100” but not on “1002”, “2100”, and so on.</p> <pre>_100_</pre>
<code>[]</code>	<p>Square brackets enclose a range of single-character patterns.</p> <p>For example, the following regular expression matches output that contains “1”, “2”, “3”, “4”, or “5”:</p> <pre>[1-5]</pre> <p>You can use the following expression symbols within the brackets. These symbols are allowed only inside the brackets.</p> <ul style="list-style-type: none"> • <code>^</code> – The caret matches on any characters except the ones in the brackets. For example, the following regular expression matches output that does not contain “1”, “2”, “3”, “4”, or “5”: <pre>[^1-5]</pre> <ul style="list-style-type: none"> • <code>-</code> – The hyphen separates the beginning and ending of a range of characters. A match occurs if any of the characters within the range is present. Refer to the example above.
<code> </code>	<p>A vertical bar separates two alternative values or sets of values. The output can match one or the other value.</p> <p>For example, the following regular expression matches output that contains either “abc” or “defg”:</p> <pre>abc defg</pre>
<code>()</code>	<p>Parentheses allow you to create complex expressions.</p> <p>For example, the following complex expression matches on “abc”, “abcabc”, or “defg”, but not on “abcdefgdefg”:</p> <pre>((abc)+) ((defg)?)</pre>

If you want to filter for a special character instead of using the special character as described in the table above, enter “\” (backslash) in front of the character. For example, to filter on output containing an asterisk, enter the asterisk portion of the regular expression as “*”.

```
Brocade# show ip route bgp | include \*
```

Allowable characters for LAG names

When creating a LAG name, you can use spaces in a file or subdirectory name if you enclose the name in double quotes. For example, to specify a subdirectory name that contains spaces, enter a string such as the following: “a long subdirectory name”. The maximum length for a string is 64 characters.

The following characters are valid in file names:

- All upper and lowercase letters
- All digits

Any of the following special characters are valid:

1 CONFIG commands

- \$
- %
- '
- -
- _
- .
- @
- ~
- ^
- !
- (
-)
- {
- }
- ^
- #
- &

CLI parsing enhancement

The response to an invalid keyword, the command returns to the cursor will include all valid content up to where the error was made. The prompt will only delete the invalid keyword “proc” and return to a prompt with the command “Brocade# show”. This will allow the user to continue typing from the point of failure, rather than having to type out the entire command again.

```
Brocade# show proc
Unrecognized command
Brocade# show
```

Syntax shortcuts

A command or parameter can be abbreviated as long as enough text is entered to distinguish it from other commands at that level. For example, given the possible commands **copy tftp...** and **config tftp...**, possible shortcuts are **cop tftp** and **con tftp** respectively. In this case, **co** does not properly distinguish the two commands.

Saving configuration changes

You can make configuration changes while the device is running. The type of configuration change determines whether or not it becomes effective immediately or requires a save to flash (**write memory**) and reset of the system (**reload**), before it becomes active.

This approach in adopting configuration changes:

- Allows you to make configuration changes to the operating or running configuration of the device to address a short-term requirement or validate a configuration without overwriting the permanent configuration file, the startup configuration, that is saved in the system flash, and;

- Ensures that dependent or related configuration changes are all cut in at the same time.

In all cases, if you want to make the changes permanent, you need to save the changes to flash using the **write memory** command. When you save the configuration changes to flash, this will become the configuration that is initiated and run at system boot.

NOTE

Most configuration changes are dynamic and thus do not require a software reload. If a command requires a software reload to take effect, the documentation states this.

Modifying startup and running configuration file manually

When you manually modify a **startup-config** or **running-config** file, ensure that you do not delete the ! (exclamation mark) from any of the lines in the configuration file.

NOTE

For configuration files which are copied to device running, or startup config via TFTP/SCP, entering a blank comment line or ! (exclamation mark denotes a comment line) followed only by blank spaces, in any of the global config sublevels, resets the mode to global config level.

1 CONFIG commands

Configuring Basic Parameters

Table 5 displays the individual Brocade devices and the basic parameters they support.

TABLE 5 Supported Brocade basic parameters features

Features supported	Brocade Netron XMR Series	Brocade MLX Series	Brocade Netron CES 2000 Series BASE package	Brocade Netron CES 2000 Series ME_PREM package	Brocade Netron CES 2000 Series L3_PREM package	Brocade Netron CER 2000 Series Base package	Brocade Netron CER 2000 Series Advanced Services package
Interactivity mode prompt	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Simple Network Management (SNMP) traps	Yes	Yes	Yes	Yes	Yes	Yes	Yes
SNMP ifIndex	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Optical Monitoring	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Optics Compatibility Checking	Yes	Yes	No	No	No	No	No
New encryption code for passwords, authentication keys, and community strings	Yes	Yes	No	No	No	No	No
Designating an Interface as the source for packets	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Setting the systemclock	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Command Alias	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Limiting broadcast, multicast, or unknown-unicast rates	Yes	Yes	Yes	Yes	Yes	Yes	Yes
CLI banners	Yes	Yes	Yes	Yes	Yes	Yes	Yes

2 Configuring Basic Parameters

TABLE 5 Supported Brocade basic parameters features (Continued)

Features supported	Brocade Netiron XMR Series	Brocade MLX Series	Brocade Netiron CES 2000 Series BASE package	Brocade Netiron CES 2000 Series ME_PREM package	Brocade Netiron CES 2000 Series L3_PREM package	Brocade Netiron CER 2000 Series Base package	Brocade Netiron CER 2000 Series Advanced Services package
Terminal display	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Modifying system parameter default settings	Yes	Yes	Yes	Yes	Yes	Yes	Yes
System low memory prevention and reporting	Yes	Yes	No	No	No	No	No
Layer 2 switching	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MAC age time	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Static ARP entries	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Configurable CAM size for IPv4 and IPv6 multicast entries	Yes	Yes	No	No	No	No	No
Switch fabric fault monitoring	Yes	Yes	No	No	No	No	No
Automatic Tuning of Links Between Line Modules and SFMs	^a Yes	^a Yes	No	No	No	No	No
Switch fabric utilization monitoring	Yes	Yes	No	No	No	No	No
Verifying an image checksum	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Displaying information for an interface for an Ethernet port	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Displaying statistics Information for an Ethernet port	Yes	Yes	Yes	Yes	Yes	Yes	Yes

TABLE 5 Supported Brocade basic parameters features (Continued)

Features supported	Brocade Netiron XMR Series	Brocade MLX Series	Brocade Netiron CES 2000 Series BASE package	Brocade Netiron CES 2000 Series ME_PREM package	Brocade Netiron CES 2000 Series L3_PREM package	Brocade Netiron CER 2000 Series Base package	Brocade Netiron CER 2000 Series Advanced Services package
Show Pause Frame Statistics	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Real-time monitoring	Yes	Yes	Yes	Yes	Yes	Yes	Yes
SNMP scalability	Yes	Yes	Yes	Yes	Yes	Yes	Yes
HW aging	Yes	Yes	Yes	Yes	Yes	Yes	Yes

a. 8x10G, 4x40G, 24x10G and 2x100G cards.

This chapter describes how to configure basic system parameters.

The Brocade device is configured with default parameters to allow you to begin using the basic features of the system immediately. However, many advanced features, such as VLANs or routing protocols for the device, must first be enabled at the system (global) level before they can be configured.

You can find system level parameters at the Global CONFIG level of the CLI.

NOTE

Before assigning or modifying any device parameters, you must assign the IP subnet (interface) addresses for each port.

Enabling and disabling interactivity for scripts

[Table 6](#) lists certain configuration and action commands that are interactive by default.

Because these commands require a user response, confirmation, or result in multiple changes across the system before the device can complete the configuration changes, they cannot be used in scripts as they are. You can, however, disable the interactive behavior by entering the **prompt** command.

Syntax: [no] prompt

The **no prompt** command will only disable the confirmation prompt for commands in configuration mode. Commands executed in the EXEC mode will continue to prompt for confirmation.

Entering the **no prompt** command allows you to use the commands and actions that are listed in [Table 6](#) within scripts without difficulty. After running a script, you can re-enable the default interactive behavior by entering the **prompt** command.

2 Entering system administration information

TABLE 6 Interactive commands

Command type	Command
Configuration	cluster-l2protocol-forward route-only rate-limit policy-map spanning-tree pms disable enable violation deny violation restrict violation shutdown
Action	reboot-standby reset reload switchover power-off lp <i>all</i> <i>slot</i> power-off power-supply <i>index</i> <i>forced</i> hitless-reload mp <i>primary</i> <i>secondary</i> lp <i>primary</i> <i>secondary</i> power-supply monitoring clear <i>all</i> <i>index</i> boot system flash <i>primary</i> <i>secondary</i>

Default behavior for certain configuration commands:

```
Brocade(config)#route-only
This will change the route-only behavior at the global level.
Are you sure? (enter 'y' or 'n'): y
Global 'route-only' committed.
```

Disabling default behavior to allow for script use:

```
Brocade(config)#no prompt
Brocade(config)#no route-only
Global 'no route-only' committed.
```

Re-enabling default behavior:

```
Brocade(config)#prompt
Brocade(config)#route-only
This will change the route-only behavior at the global level.
Are you sure? (enter 'y' or 'n'): y
Global 'route-only' committed.
```

Entering system administration information

You can configure a system name, contact, and location for the Brocade device and save the information locally in the configuration file for future reference. The information is not required for system operation but recommended. When you configure a system name, it replaces the default system name in the CLI command prompt.

To configure a system name, contact, and location, enter commands such as the following.

```

Brocade(config)# hostname home
home(config)# snmp-server contact Suzy Sanchez
home(config)# snmp-server location Centerville
home(config)# end
home# write memory

```

The system name you configure **home** replaces the system name Brocade.

Syntax: [no] hostname *string*

Syntax: [no] snmp-server contact *string*

Syntax: [no] snmp-server location *string*

The name, contact, and location each can be up to 255 alphanumeric characters. The text strings can contain blanks. The SNMP text strings do not require quotation marks when they contain blanks but the host name does.

Configuring Simple Network Management (SNMP) traps

This section explains how to do the following:

- Specify an SNMP trap receiver.
- Specify a source address and community string for all traps that the Brocade device sends.
- Change the holddown time for SNMP traps.
- Disable individual SNMP traps. (All traps are enabled by default.)
- Disable traps for CLI access that is authenticated by a local user account, a RADIUS server, or a TACACS or TACACS+ server.

Specifying an SNMP trap receiver

You can specify a trap receiver to ensure that all SNMP traps sent by the Brocade device go to the same SNMP trap receiver or set of receivers, typically one or more host devices on the network. When you specify the host, you also specify a community string. The Brocade device sends all the SNMP traps to the specified hosts and includes the specified community string. Administrators can therefore filter for traps from a Brocade device based on IP address or community string. The number of SNMP Trap receivers that can be configured is limited to 32.

If the string is in the clear format, the system will internally encrypt it. When you display or save the configuration, the encrypted string is used.

To specify an SNMP trap receiver, enter a command such as the following.

```

Brocade(config)# snmp-server host 10.2.2.2 version v2c mypublic port 200

```

The command adds trap receiver 10.2.2.2 and designates the UDP port that will be used to receive traps.

Syntax: [no] snmp-server host *ip-addr* version [v1|v2c|v3] *string* [port *value*]

The *ip-addr* parameter specifies the IP address of the trap receiver.

2 Configuring Simple Network Management (SNMP) traps

The `v1`, `v2c`, or `v3` parameter indicates which version of SNMP is used.

The `string` parameter specifies an SNMP community string configured on the Brocade device. It is not used to authenticate access to the trap host, but it is a useful method for filtering traps on the host. For example, if you configure each of your Brocade devices that use the trap host to send a different community string, you can easily distinguish among the traps from the devices based on the community strings.

By default, `string` is encrypted. If you want `string` to be in clear text, insert a `0` preceding `string`.

Example

```
Brocade(config)# snmp-server host 10.2.2.2 version v2c 0 mypublic port 200
```

The software adds a prefix to the string in the configuration. For example, the following portion of the code has the encrypted code “2”.

```
snmp-server host 10.2.2.2 version v2c 12 $Si2^=d
```

The prefix can be one of the following:

- 0 = the key string is not encrypted and is in clear text
- 1 = the key string uses simple encryption (only for Brocade NetIron CES and Brocade NetIron CER)
- 2 = the key string uses base64 encryption format (only for Brocade NetIron XMR and Brocade MLX series)

The `port value` parameter specifies the UDP port that will be used to receive traps. This parameter allows you to configure several trap receivers in a system. With this parameter, Brocade Network Advisor and another network management application can coexist in the same system. The Brocade devices can be configured to send copies of traps to more than one network management application.

Specifying a single trap source

You can specify a single trap source to ensure that all SNMP traps sent by the Brocade device use the same source IP address. When you configure the SNMP source address, you specify the Ethernet port, loopback interface, or virtual routing interface that is the source for the traps. The Brocade device then uses the lowest-numbered IP address configured on the port or interface as the source IP address in the SNMP traps it sends.

Identifying a single source IP address for SNMP traps provides the following benefits:

- If your trap receiver is configured to accept traps only from specific links or IP addresses, you can simplify configuration of the trap receiver by configuring the Brocade device to always send the traps from the same link or source address.
- If you specify a loopback interface as the single source for SNMP traps, SNMP trap receivers can receive traps regardless of the states of individual links. Thus, if a link to the trap receiver becomes unavailable but the receiver can be reached through another link, the receiver still receives the trap, and the trap still has the source IP address of the loopback interface.

To configure the Brocade device to send all SNMP traps from the first configured IP address on port 4/11, enter the following commands.

```
Brocade(config)# snmp-server trap-source ethernet 4/11
Brocade(config)# write memory
```

Syntax: `[no] snmp-server trap-source loopback num | ethernet slot/port | ve num`

The *num* parameter is a loopback interface or virtual routing interface number.

If you do not configure this command, the device will use the device router ID as the source IP address of the notification packet. The router ID of the device can be obtained from the “show ip” command output.

To specify a loopback interface as the device’s SNMP trap source, enter following commands.

```
Brocade(config)# int loopback 1
Brocade(config-lbif-1)# ip address 10.0.0.1/24
Brocade(config-lbif-1)# exit
Brocade(config)# snmp-server trap-source loopback 1
```

The commands configure loopback interface 1, gives it IP address 10.0.0.1/24, then designate it as the SNMP trap source for the Brocade device. Regardless of the port the Brocade uses to send traps to the receiver, the traps always arrive from the same source IP address.

Setting the SNMP trap holddown time

When a Brocade device starts up, the software waits for Layer 2 convergence (STP) and Layer 3 convergence (OSPF) before beginning to send SNMP traps to external SNMP servers. Until convergence occurs, the Brocade device might not be able to reach the servers, in which case the messages are lost.

By default, the Brocade device uses a one-minute holddown time to wait for the convergence to occur before starting to send SNMP traps. After the holddown time expires, the Brocade device sends the traps, including traps such as “cold start” or “warm start” that occur before the holddown time expires.

You can change the holddown time to a value from one second to ten minutes.

To change the holddown time for SNMP traps, enter a command such as the following at the global CONFIG level of the CLI.

```
Brocade(config)# snmp-server enable traps holddown-time 30
```

The command changes the holddown time for SNMP traps to 30 seconds. The Brocade device waits 30 seconds to allow convergence in STP and OSPF before sending traps to the SNMP trap receiver.

Syntax: `[no] snmp-server enable traps holddown-time secs`

The *secs* parameter specifies the number of seconds (1 – 600). The default is 60.

Disabling SNMP traps

The Brocade device comes with SNMP trap generation enabled by default for all traps.

NOTE

By default, all SNMP traps are enabled at system startup.

You can selectively disable one or more of the following traps:

- SNMP authentication key
- Temperature

2 Configuring SNMP ifIndex

- Power supply failure
- Fan failure
- Cold start
- Link up
- Link down
- Bridge new root
- Bridge topology change
- Module insert
- Module remove
- Redundant module
- Metro-ring
- MPLS
- BGP4
- OSPF
- VRRP
- VSRP

To stop link down occurrences from being reported, enter the following command.

```
Brocade(config)# no snmp-server enable traps link-down
```

Syntax: `[no] snmp-server enable traps trap-type`

A list of traps is available in the *Unified IP MIB Reference*.

Configuring SNMP ifIndex

This section explains how ifIndex values are assigned on Brocade devices.

On Brocade NetIron CES and Brocade NetIron CER only

On the Brocade NetIron CES and Brocade NetIron CER, the system automatically assign 64 indexes to each module on the device. This value is not configurable.

On Brocade NetIron XMR and Brocade MLX series only

On Brocade NetIron XMR and Brocade MLX series devices, SNMP Management Information Base (MIB) uses Interface Index (ifIndex) to assign a unique value to each port on a module or slot. The number of indexes that can be assigned per module is 20, 40, or 64, depending on the number of ports on the module.

Enter the following to change the number of indexes per module.

```
Brocade(config)# snmp-server max-ifindex-per-module 40
```

Syntax: `[no] snmp-server max-ifindex-per-module [20 | 40 | 64]`

20 is the default.

You cannot change the maximum ifIndex per module to a number less than the number of ports.

After this command is issued the following are generated:

- “System: IfIndex assignment was changed” is logged in the Syslog.
- The snTrapIfIndexAssignmentChanged trap is generated.

Configuration notes for the Brocade NetIron XMR and Brocade MLX series

Note the following if you are upgrading the software on the Brocade NetIron XMR and Brocade MLX series:

- If you are running an earlier version of the software and you will not be installing the NI-MLX-1Gx48-T module, you do not need to change your ifIndex allocation scheme. The current definition is maintained. The maximum ifIndex per module can remain at 20 or 40.
- If you are running an earlier version of the software and you will be installing the NI-MLX-1Gx48-T module on you Brocade MLX series, you must configure the maximum ifIndex per module to 64. **You must change the ifIndex allocation before installing the NI-MLX-1Gx48-T module**; otherwise, the module status remains in the Offline state.
- If you have a new Brocade MLX series (no previous software installed), but will not be installing an NI-MLX-1Gx48-T module, it is recommended that you configure the maximum ifIndex per module to 64 to avoid future ifIndex problems in case an NI-MLX-1Gx48-T module is installed in the future.
- If you have a new Brocade MLX series (no previous software installed), and you will be installing an NI-MLX-1Gx48-T module, you **must** configure the maximum ifIndex per module to 64; otherwise, the module remains in the Offline state.

SNMP scalability optimization

To ensure that SNMP requests are responded to promptly and that SNMP loads do not impact other device activities, the Brocade device speeds SNMP tasks and limit their effects on the CPU by a combination of throughput optimization and load throttling.

Configuring SNMP throughput optimization

SNMP throughput is optimized on the Brocade device through a combination of SNMP value caching, conditional yielding by the SNMP agent, and acceptance of incoming packets during queue processing.

Configuring SNMP value caching

To accelerate SNMP communication between the management module (MP) and interface module (LP), the LP will return values for multiple ports when it receives a request from the MP. These values will be cached on the MP to reduce MP/LP communication time.

The SNMP value caching mechanism is enabled by default. To disable the caching mechanism, use the following command at the MPLS configuration level of the CLI.

```
Brocade(config-mpls)#snmp-server cache disable
```

Syntax: `[no] snmp-server cache disable`

Use the **no** form of the command to enable SNMP value caching.

To configure the maximum length of time that a cached SNMP port value will be considered valid by the MP, use the following command at the MPLS configuration level of the CLI.

```
Brocade (config-mpls)#snmp-server cache max-age 10000
```

Syntax: `[no] snmp-server cache max-age milliseconds`

- The preceding example sets the maximum age for cached SNMP port values to 10,000 milliseconds.
- The *milliseconds* parameter is the maximum number of milliseconds a value will be considered valid before the MP discards the value and requests a new one from the LP. The range for this parameter is from 50 through 15000.
- Use the **no** form of this command to return the cache aging limit to the default value of 100 milliseconds.

To configure the maximum size of the SNMP port value cache, use the following command at the MPLS configuration level of the CLI.

```
Brocade (config-mpls)#snmp-server cache size 100
```

Syntax: `[no] snmp-server cache size kilobytes`

- The preceding example increases the maximum size for the SNMP port cache to 100 KB.
- The *kilobytes* parameter is the maximum memory size for the SNMP port cache.
- Use the **no** form of this command to return the maximum cache size to the default value of 32 KB.

SNMP agent yielding behavior

When an SNMP agent yields CPU control unconditionally between processing of queued packets, it can result in low throughput for packets which are processed quickly. To increase throughput for these packets, the SNMP agent in the Brocade device yields CPU control between packets only when the agent has controlled the CPU for more than 10 milliseconds.

SNMP queue processing

To ensure that SNMP packets are not dropped, the SNMP task on the Brocade device continues to accept newly received SNMP packets from the IP stack while processing the SNMP queue.

Configuring SNMP load throttling

To ensure that high SNMP loads do not interfere with the performance of the device, the Brocade device limits the percentage of CPU time that can be occupied by SNMP processing. This limit is not imposed when the CPU is idle.

NOTE

This command tries to fix the maximum percentage of time SNMP task can run in a non-idle system environment. This implies that SNMP task can't run for more than the specified percentage of time if the system is having zero idle time. But this constraint is checked only between processing of 2 SNMP PDU's. If the processing of a single SNMP PDU takes longer time then we may overrun the maximum limit.

This command also tries to fix the minimum percentage of time SNMP task can run in a non-idle system environment. But if there is another task which is continuously hogging the CPU and SNMP is not getting time to run then we may under run the specified limit.

To configure the maximum percentage of CPU time that can be used by SNMP processing, use the following command at the configuration level of the CLI.

```
Brocade(config)#snmp-server cpu max-non-idle-utilization 25
```

Syntax: `[no] snmp-server cpu max-non-idle-utilization percent`

- The preceding example raises the maximum percentage of non-idle CPU time to be used by SNMP processing to 25%.
- The *percent* parameter is the maximum percentage of non-idle CPU time to be used by SNMP processing. The range for this parameter is from 1 through 25.
- Use the **no** form of this command to return the SNMP non-idle CPU time maximum to the default value of 10%.

Configuring optical monitoring

You can configure your Brocade 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.

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

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

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

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

```
Brocade(config)# interface ethernet 1/1 to 1/2
Brocade(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.

2 Configuring optical monitoring

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

```
Brocade#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           Normal      Normal      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 7](#) describes each of these status values.

TABLE 7 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.

```
Brocade#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

2 Displaying media information

```
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
Brocade#
```

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.

```
Brocade#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.

```
Brocade#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 Brocade 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 29 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.

```
Brocade(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.

Designating an interface as the packet source

The software uses the lowest-numbered IP address configured on an interface as the source IP address for all Telnet, SSH, NTP, TFTP, TACACS or TACACS+, or RADIUS packets originated from the Brocade device.

You can specify the source interface for one or more of these types of packets.

Configuring an interface as the source for all Telnet packets

Identifying a single source IP address for Telnet packets provides the following benefits:

- If your Telnet server is configured to accept packets only from specific links or IP addresses, you can simplify configuration of the Telnet server by configuring the Brocade device to always send the Telnet packets from the same link or source address.
- If you specify a loopback interface as the single source for Telnet packets, Telnet servers can receive the packets regardless of the states of individual links. Thus, if a link to the Telnet server becomes unavailable but the client or server can be reached through another link, the client or server still receives the packets, and the packets still have the source IP address of the loopback interface.

To specify the lowest-numbered IP address configured on a loopback interface as the device's source for all Telnet packets, enter commands such as the following.

```
Brocade(config)# int loopback 2
Brocade(config-lbif-2)# ip address 10.0.0.2/24
Brocade(config-lbif-2)# exit
Brocade(config)# ip telnet source-interface loopback 2
```

The commands configure loopback interface 2, assign IP address 10.0.0.2/24 to it, then designate it as the source for all Telnet packets from the Brocade device.

Syntax: [no] ip telnet source-interface ethernet *portnum* | loopback *num* | ve *num*

The following commands configure an IP interface on an Ethernet port and designate the address port as the source for all Telnet packets from the Brocade device.

```
Brocade(config)# interface ethernet 1/4
Brocade(config-if-e10000-1/4)# ip address 10.157.22.110/24
Brocade(config-if-e10000-1/4)# exit
Brocade(config)# ip telnet source-interface ethernet 1/4
```

Cancelling an outbound Telnet session

If you want to cancel a Telnet session from the console to a remote Telnet server (for example, if the connection is frozen), you can terminate the Telnet session by performing the tasks listed below.

1. At the console, press Ctrl-^ (Ctrl-Shift-6).
2. Press the X key to terminate the Telnet session.

Pressing Ctrl-^ twice in a row causes a single Ctrl-^ character to be sent to the Telnet server. After you press Ctrl-^, pressing any key other than X or Ctrl-^ returns you to the Telnet session.

Configuring an interface as the source for all SSH packets

You can configure the Brocade device to use the lowest-numbered IP address configured on a loopback interface, virtual routing interface, or Ethernet port as the source for the SSH packets it sends.

For example, to specify an Ethernet port as the interface whose lowest-numbered IP address will be the source address for the SSH packets originated from the Brocade device, enter commands such as the following.

```
Brocade(config)# interface ethernet 1/5
Brocade(config-if-e10000-1/5)# ip address 10.157.22.111/24
Brocade(config-if-e10000-1/5)# exit
Brocade(config)# ip ssh source-interface ethernet 1/5
```

The commands configure Ethernet port 1/5, assign IP address 10.157.22.111/24 to it, then designate it as the source interface.

Syntax: [no] ip ssh source-interface ethernet *portnum* | loopback *num* | ve *num*

Configuring an interface as the source for all TFTP packets

You can configure the Brocade device to use the lowest-numbered IP address configured on a loopback interface, virtual routing interface, or Ethernet port as the source for all TFTP packets it sends.

For example, to specify the lowest-numbered IP address configured on a virtual routing interface as the Brocade device's source for all TFTP packets, enter commands such as the following.

```
Brocade(config)# int ve 1
Brocade(config-vif-1)# ip address 10.0.0.3/24
Brocade(config-vif-1)# exit
Brocade(config)# ip tftp source-interface ve 1
```

2 Designating an interface as the packet source

The commands configure virtual routing interface 1, assign IP address 10.0.0.3/24 to it, then designate the address as the source address for all TFTP packets.

Syntax: `[no] ip tftp source-interface ethernet portnum | loopback num | ve num`

The default is the lowest-numbered IP address configured on the port through which the packet is sent. The address therefore changes, by default, depending on the port.

Configuring an interface as the source for all TACACS or TACACS+ packets

You can configure the Brocade device to use the lowest-numbered IP address configured on a loopback interface, virtual routing interface, or Ethernet port as the source for the TACACS or TACACS+ packets it sends.

For example, to specify a virtual routing interface as the interface whose lowest-numbered IP address will be the source address for the TACACS or TACACS+ packets originated from the Brocade device, enter commands such as the following.

```
Brocade(config)# int ve 1
Brocade(config-vif-1)# ip address 10.0.0.3/24
Brocade(config-vif-1)# exit
Brocade(config)# ip tacacs source-interface ve 1
```

The commands configure virtual routing interface 1, assign IP address 10.0.0.3/24 to it, then designate it as the source interface.

Syntax: `[no] ip tacacs source-interface ethernet portnum | loopback num | ve num`

Configuring an interface as the source for all RADIUS packets

You can configure the Brocade device to use the lowest-numbered IP address configured on a loopback interface, virtual routing interface, or Ethernet port as the source for the RADIUS packets it sends.

For example, to specify an Ethernet port as the interface whose lowest-numbered IP address will be the source address for the RADIUS packets originated from the Brocade device, enter the following commands.

```
Brocade(config)# interface ethernet 1/5
Brocade(config-if-e10000-1/5)# ip address 10.157.22.111/24
Brocade(config-if-e10000-1/5)# exit
Brocade(config)# ip radius source-interface ethernet 1/5
```

The commands configure Ethernet port 1/5, assign IP address 10.157.22.111/24 to it, then designate it as the source interface.

Syntax: `[no] ip radius source-interface ethernet portnum | loopback num | ve num`

Setting the system clock

The Brocade device allows you to manually set the system clock. Using the **clock set** command starts the system clock with the time and date you specify. The time counter setting is retained across power cycles.

To set the system time and date to 10:15:05 on October 15, 2005, enter the following command.

```
Brocade# clock set 10:15:05 10-15-05
```

Syntax: **[no] clock set** *hh:mm:ss mm-dd-yy | mm-dd-yyyy*

By default, the Brocade device does not change the system time for daylight savings time. To enable daylight savings time, enter the following command.

```
Brocade# clock summer-time
```

Syntax: **[no] clock summer-time**

You can configure the Brocade device to adjust the time for any one-hour offset from GMT or for one of the following U.S. time zones:

- US Pacific (default)
- Alaska
- Aleutian
- Arizona
- Central
- East-Indiana
- Eastern
- Hawaii
- Michigan
- Mountain
- Pacific
- Samoa

The default is US Pacific.

To change the time zone to Australian East Coast time (which is normally 10 hours ahead of GMT), enter the following command.

```
Brocade(config)# clock timezone gmt gmt+10
```

Syntax: **[no] clock timezone gmt gmt | us time-zone**

You can enter one of the following values for *time-zone*:

- US time zones (**us**): alaska, aleutian, arizona, central, east-indiana, eastern, hawaii, michigan, mountain, pacific, samoa.
- GMT time zones (**gmt**): gmt+12, gmt+11, gmt+10...gmt+01, gmt+00, gmt-01...gmt-10, gmt-11, gmt-12.

DST “change” notice for networks using US time zones

The new Daylight Saving Time (DST) change that went into effect on March 11, 2007 affects networks in the US time zones. Because of this change, your network clock might not be correct. If your network uses US time zones, and it needs to maintain the correct time, you must enable the following command.

```
Brocade(config)# clock timezone us pacific
```

Syntax: [no] clock timezone us {pacific | eastern | central | mountain}

NOTE

This command must be configured on every device that uses the US DST.

To verify the change, use the following command.

```
Brocade(config)# show clock
```

For more information, refer to www.brocade.com.

Creating a command alias

Use the **alias** command to create an alias for a command and to save that alias within the device's configuration.

To create the alias “shro” for the **show ip routes** command, use the following command.

```
Brocade(config)# alias shro = show ip routes
Brocade(config)# write memory
```

Syntax: [no] alias [name = command]

The *name* variable is the name that you want to assign to the alias.

The *command* variable is the syntax for the command you want to create an alias for.

The **write** memory command is used to save the alias within the configuration.

Removing an alias

You can remove an alias using the **no** version of the alias command as shown in the following.

```
Brocade(config)# no alias shro
```

Alternately, you can use the **unalias** command as shown in the following.

```
Brocade(config)# unalias shro
```

Syntax: [no] unalias

If the alias you try to remove does not exist, the following error will be displayed.

```
Brocade(config)# unalias wrs
Error: Alias wrs does not exist, unalias failed
```

Displaying a list of all configured alias

The following command allows you to display a list of all configured alias.

```
Brocade# alias
#alias
      savemem      write memory
      shro         show ip routes
```

Syntax: [no] alias

Limiting broadcast, multicast, or unknown unicast rates

The Brocade device can forward all traffic at wire speed. However, some third-party networking devices cannot handle high forwarding rates for broadcast, multicast, or unknown unicast packets.

The limits are individually configurable for broadcasts, multicasts, and unknown unicasts. You can configure limits globally to apply to each individual inbound interface module. The valid range is 1 – 4294967295 packets per second. The default is 0, which disables limiting.

NOTE

These packets will be sent to the CPU for software forwarding. Brocade recommends that you use the BUM rate limiting or ACL-based rate limiting as they are performed by hardware forwarding.

Limiting broadcasts

To globally limit the number of broadcast packets a Brocade device forwards to 100,000 per second, enter the following command at the global CONFIG level of the CLI.

```
Brocade(config)# broadcast limit 100000
Brocade(config)# write memory
```

Syntax: [no] broadcast limit *number*

Limiting multicasts

To globally limit the number of multicast packets a Brocade device forwards to 120,000 per second, enter the following command at the global CONFIG level of the CLI.

```
Brocade(config)# multicast limit 120000
Brocade(config)# write memory
```

Syntax: [no] multicast limit *number*

NOTE

The multicast limit is configured at the global level, but the value you enter applies to each interface module (slot) installed on the device.

Limiting unknown unicasts

To globally limit the number of unknown unicast packets a Brocade device forwards to 110,000 per second, enter the following command at the global CONFIG level of the CLI.

```
Brocade(config)# unknown-unicast limit 110000  
Brocade(config)# write memory
```

Syntax: [no] unknown-unicast limit *number*

NOTE

Only the **unknown-unicast limit** is configured on the global level, but the value you enter applies to each interface module (slot) installed on the device.

Configuring CLI banners

The Brocade device can be configured to display a greeting message on users' terminals when they enter the Privileged EXEC CLI level or access the device through Telnet. In addition, a Brocade device can display a message on the Console when an incoming Telnet CLI session is detected.

Setting a message of the day banner

You can configure the Brocade device to display a message on a user's terminal when he or she establishes a Telnet CLI session. For example, to display the message "Welcome to Brocade!" when a Telnet CLI session is established, enter the following.

```
Brocade(config)# banner motd $(Press Return)  
Enter TEXT message, End with the character '$'.  
Welcome to Brocade! $
```

A delimiting character is established on the first line of the **banner motd** command. You begin and end the message with this delimiting character. The delimiting character can be any character except "(double-quotation mark) and cannot appear in the banner text. In this example, the delimiting character is \$(dollar sign). The text in between the dollar signs is the contents of the banner. The banner text can be up to 2047 characters long and can consist of multiple lines. To remove the banner, enter the **no banner motd** command.

Syntax: [no] banner *delimiting-character* | [motd *delimiting-character*]

NOTE

The **banner delimiting-character** command is equivalent to the **banner motd delimiting-character** command.

NOTE

The size of the MOTD banner will be restricted (truncated) to 1850 characters when using an SSH client.

Setting a privileged EXEC CLI level banner

You can configure the Brocade device to display a message when a user enters the Privileged EXEC CLI level.

Example

```
Brocade(config)# banner exec_mode # (Press Return)
Enter TEXT message, End with the character '#'.
You are entering Privileged EXEC level
Don't foul anything up! #
```

As with the **banner motd** command, you begin and end the message with a delimiting character; in this example, the delimiting character is # (pound sign). To remove the banner, enter the **no banner exec_mode** command.

Syntax: **[no] banner exec_mode delimiting-character**

Displaying a message on the console when an incoming Telnet session is detected

You can configure the Brocade device to display a message on the Console when a user establishes a Telnet session. This message indicates where the user is connecting from and displays a configurable text message.

Example

```
Brocade(config)# banner incoming $(Press Return)
Enter TEXT message, End with the character '$'.
Incoming Telnet Session! $
```

When a user connects to the CLI using Telnet, the following message appears on the Console.

```
Telnet from 10.157.22.63
Incoming Telnet Session!
```

Syntax: **[no] banner incoming delimiting-character**

To remove the banner, enter the **no banner incoming** command.

Configuring terminal display

You can configure and display the number of lines displayed on a terminal screen during the current CLI session.

The **terminal length** command allows you to determine how many lines will be displayed on the screen during the current CLI session. This command is useful when reading multiple lines of displayed information, especially those that do not fit on one screen.

To specify the maximum number of lines displayed on one page, enter a command such as the following.

```
Brocade#terminal length 15
```

Syntax: **[no] terminal length number-of-lines**

2 Enabling or disabling routing protocols

The *number-of-lines* parameter indicates the maximum number of lines that will be displayed on a full screen of text during the current session. If the displayed information requires more than one page, the terminal pauses. Pressing the space bar displays the next page.

The default for *number-of-lines* is 24. Entering a value of 0 prevents the terminal from pausing between multiple output pages:

Checking the length of terminal displays

The **show terminal** command specifies the number of lines that will be displayed on the screen as specified by the **terminal length**, **page display**, and **skip-page-display** commands. It also shows if the **enable skip-page-display** command has been configured. The **enable skip-page-display** command allows you to use the **skip-page-display** to disable the configured page-display settings.

```
Brocade(config)# show terminal
Length: 24 lines
Page display mode (session): enabled
Page display mode (global): enabled
```

Syntax: **show terminal**

Enabling or disabling routing protocols

The Multi-Service IronWare supports the following protocols:

- BGP4
- DVMRP
- IP
- IS-IS
- MPLS
- MSDP
- OSPF
- PIM
- RIP
- VRRP
- VRRPE

By default, IP routing is enabled on the Brocade device. All other protocols are disabled, so you must enable them to configure and use them.

To enable a protocol on a Brocade device, enter **router** at the global CONFIG level, followed by the protocol to be enabled. The following example shows how to enable OSPF.

```
Brocade(config)# router ospf
```

Syntax: **[no] router bgp | dvmrp | ospf | pim | rip | vrrp | vrrpe**

Displaying and modifying default settings for system parameters

The Multi-Service IronWare has default table sizes for the following parameters. The table sizes determine the maximum number of entries the tables can hold. You can adjust individual table sizes to accommodate your configuration needs:

- MAC address entries
- VLANs supported on a system
- Virtual interfaces
- Spanning tree instances
- RSTP instances
- IP cache size
- ARP entries
- IP routes
- IP ACL filter entries
- L2 ACL entries per ACL table
- Size for management port ACL
- IP subnets per port and per device
- IPv6 Multicast routes
- IPv6 PIM mcache
- Layer 4 sessions supported
- Number of VPLS's
- VPLS MAC entries
- VRF routes
- IPv6 cache
- IPv6 routes
- Number of tunnels
- Number of LAGs
- Configuration file size

The tables you can configure as well the defaults and valid ranges for each table differ depending on the Brocade device you are configuring.

NOTE

If you increase the number of subnet addresses you can configure on each port to a higher amount, you might also need to increase the total number of subnets that you can configure on the device.

NOTE

Changing the table size for a parameter reconfigures the device's memory. Whenever you reconfigure the memory on a Brocade device, you must save the change to the startup configuration file, then reload the software to place the change into effect.

2 Displaying and modifying default settings for system parameters

To display the configurable tables and their defaults and maximum values, enter the following command at any level of the CLI.

FIGURE 1 Output for the Brocade NetIron XMR and Brocade MLX series

```

Brocade#show default values
sys log buffers:50          mac age time:300 sec          telnet sessions:5
ip arp age:10 min          bootp relay max hops:4        ip ttl:64 hops
ip addr per intf:24
when multicast enabled :
igmp group memb.: 260 sec  igmp query:          125 sec
when ospf enabled :
ospf dead:40 sec          ospf hello:10 sec          ospf retrans:5 sec
ospf transit delay:1 sec
when bgp enabled :
bgp local pref.:100       bgp keep alive:60 sec       bgp hold:180 sec
bgp metric:10            bgp local as:1             bgp cluster id:0
bgp ext. distance:20     bgp int. distance:200      bgp local distance:200
when IS-IS enabled :
isis hello interval:10 sec      isis hello multiplier:3
isis port metric:10           isis priority:64
isis csnp-interval:10 sec      isis default-metric:10
isis distance:115            isis lsp-gen-interval:10 sec
isis lsp-interval:33 msec      isis lsp-refresh-interval:900 sec
isis max-lsp-lifetime:1200 sec  isis maximum-paths:4
isis retransmit-interval:5 sec  isis spf-interval:5 sec
filter change update delay:10 sec
System Parameters   Default   Maximum   Current   Actual   Bootup Revertible
mac                 131072   2097152   2097152   2097152   2097152   Yes
vlan                512     4095     4095     4095     4095     No
spanning-tree       32      128     128     128     128     No
rstp                32      128     128     128     128     No
ip-arp              8192    65536   65536   65536   65536   No
multicast-route (IPv6) 8192    153600  153600  153600  153600  Yes
pim-mcache (IPv6)   4096    4096    4096    4096    4096    Yes
ip-cache            204800  1048576 1048576 1048576 1048576  Yes
ip-route            204800  1048576 1048576 1048576 1048576  Yes
ip-subnet-port      24      128     128     128     128     No
virtual-interface   255     4095    4095    4095    4095    No
vpls-mac            8192    1000000 1000000 1000000 1000000  Yes
vpls-num            2048    16384   16384   16384   16384   No
session-limit       32768   163840  163840  163840  163840  Yes
ip-filter-sys       4096    40960   40960   40960   40960   No
mgmt-port-acl-size  20      100     100     100     100     No
l2-acl-table-entries 64      256     256     256     256     No
ipv6-cache          65536   245760  245760  245760  245760  Yes
ipv6-route          65536   245760  245760  245760  245760  Yes
vrf-route           5120    262143  262143  262143  262143  Yes
receive-cam         1024    8192    8192    8192    8192    No
ip-tunnels          256     8192    8192    8192    8192    No
lsp-out-acl-cam     0        16384   16384   16384   16384   No
trunk-num           128     256     256     256     256     No
config-file-size    8388608 16777216 16777216 16777216 16777216 No
ifl-cam             0        81920   49152   49152   49152   No
ip-source-guard-cam 0        131072  30000   30000   30000   No
ipv4-mcast-cam      8192    65536   10000   10000   10000   No
ipv6-mcast-cam      2048    16384   3500    3500    3500
No

```

FIGURE 2 Output for the Brocade NetIron CES

```

Brocade#show default values
sys log buffers:50          mac age time:300 sec          telnet sessions:5

ip arp age:10 min          bootp relay max hops:4      ip ttl:64 hops
ip addr per intf:24

when multicast enabled :
igmp group memb.:140 sec   igmp query:60 sec

when ospf enabled :
ospf dead:40 sec           ospf hello:10 sec          ospf retrans:5 sec
ospf transit delay:1 sec

when bgp enabled :
bgp local pref.:100        bgp keep alive:60 sec      bgp hold:180 sec
bgp metric:10              bgp local as:1             bgp cluster id:0
bgp ext. distance:20       bgp int. distance:200     bgp local distance:200

when IS-IS enabled :
isis hello interval:10 sec   isis hello multiplier:3
isis port metric:10          isis priority:64
isis csnp-interval:10 sec    isis default-metric:10
isis distance:115           isis lsp-gen-interval:10 sec
isis lsp-interval:33 msec    isis lsp-refresh-interval:900 sec
isis max-lsp-lifetime:1200 sec isis maximum-paths:4
isis retransmit-interval:5 sec isis spf-interval:5 sec

filter change update delay:10 sec

System Parameters   Default   Maximum   Current
mac                 56320    131072    56320
vlan                 512      4095      512
spanning-tree       32        128       32
rstp                 32        128       32
ip-arp              4096     16384     4096
multicast-route (IPv6) 1024     2048     1024
pim-mcache (IPv6)    1024     2048     1024
ip-cache            16384    32768    16384
ip-route            16384    32768    16384
ip-subnet-port       24        128       24
virtual-interface    255       1024     255
vpls-mac             512       1024     512
vpls-num             512       1024     512
session-limit        32768    32768    32768
ip-filter-sys        4096     8192     8192
mgmt-port-acl-size   20        100       20
l2-acl-table-entries 64        256       256
ipv6-cache           1024     8192     1024
ipv6-route           1024     8192     1024
vrf-route            1024     32768    1024
receive-cam          1         1         1
ip-tunnels           32        128       32
lsp-out-acl-cam      1         1         1
trunk-num            128       255       128
    
```

2 Displaying and modifying default settings for system parameters

FIGURE 3 Output for the Brocade NetIron CER device

```

Brocade#show default values
sys log buffers:50          mac age time:300 sec          telnet sessions:5

ip arp age:10 min          bootp relay max hops:4        ip ttl:64 hops
ip addr per intf:24
when multicast enabled :
igmp group memb.: 260 sec  igmp query:          125 sec

when ospf enabled :
ospf dead:40 sec          ospf hello:10 sec            ospf retrans:5 sec
ospf transit delay:1 sec
when bgp enabled :
bgp local pref.:100       bgp keep alive:60 sec        bgp hold:180 sec
bgp metric:10             bgp local as:1              bgp cluster id:0
bgp ext. distance:20      bgp int. distance:200       bgp local distance:200

when IS-IS enabled :
isis hello interval:10 sec      isis hello multiplier:3
isis port metric:10           isis priority:64
isis csnp-interval:10 sec      isis default-metric:10
isis distance:115            isis lsp-gen-interval:10 sec
isis lsp-interval:33 msec      isis lsp-refresh-interval:900 sec
isis max-lsp-lifetime:1200 sec  isis maximum-paths:4
isis retransmit-interval:5 sec  isis spf-interval:5 sec

filter change update delay:10 sec

```

System Parameters	Default	Maximum	Current	Actual	Bootup	
Revertible						
mac	65536	131072	65536	65536	65536	Yes
vlan	512	8192	512	512	512	Yes
spanning-tree	32	128	32	32	32	Yes
rstp	32	128	32	32	32	Yes
ip-arp	4096	16384	4096	4096	4096	Yes
ip-cache	290816	524288	290816	290816	290816	Yes
ip-route	290816	524288	290816	290816	290816	Yes
ip-subnet-port	24	128	24	24	24	Yes
virtual-interface	255	4095	255	255	255	Yes
vpls-mac	2048	131072	2048	2048	2048	Yes
vpls-num	128	1024	128	128	128	Yes
session-limit	32768	32768	32768	32768	32768	Yes
ip-filter-sys	4096	32768	4096	4096	4096	Yes
mgmt-port-acl-size	20	100	20	20	20	Yes
l2-acl-table-entries	64	256	64	64	64	Yes
ipv6-cache	8192	131072	8192	8192	8192	Yes
ipv6-route	8192	131072	8192	8192	8192	Yes
ip-vrf-route	1024	524288	1024	1024	1024	Yes
ip-tunnels	32	256	32	32	32	Yes
config-file-size	8388608	16777216	8388608	8388608	8388608	Yes
ip-source-guard-cam	0	131072	0	0	0	No
ip-vrf	16	128	16	16	16	Yes
ipv6-vrf-route	128	16384	128	128	128	Yes
openflow-pvlan-entries	0	2048	0	0	0	Yes

Syntax: show default values

The following table describes the system-max values of the **show default values** command for Brocade NetIron XMR and Brocade MLX series.

TABLE 8 Display of show default values for system parameters

This field...	Displays...
Default	The default value for the system-max element. This value is used in the following conditions: a) There is no system-max configured for the corresponding element. b) If the system-max element configuration is reverted at bootup time (if it is a revertible element).
Maximum	The maximum value that this element can be configured at.
Current	The most current configured value for the system-max element. If the system-max element is configured in the running system, then the value under this column will change to reflect this value. NOTE: The new value does not take affect until the node is reloaded.
Actual	The system-max value that is used by the target application of the running system. If system-max elements are reverted at bootup, then only the Actual column is affected. The Application is now using default values and will be displayed in the Actual column. Please refer to the example on the next page for more information. The Current and Bootup values are still configured on the system, and are not affected by the reversion of system-max elements at bootup.
Bootup	The system-max value that was read from the configuration when the system was booting up. If the read values are found to be acceptable, and not reverted, then the values in this column, and in the "Actual" column will have the same values. However, if the values were reverted during bootup, then the values are different for the "Revertible" elements.
Revertible	This column displays which corresponding system-max element is revertible or not. If "Yes" is displayed then the value is changed to a default value. If "No" is displayed then there no change to the value.

If system-max elements are reverted at bootup time, then the following message will display on the CLI.

```
Brocade#show default values
...
```

NOTE: All the Revertible Elements were Reverted During System Bringup.

System Parameters	Default	Maximum	Current	Actual	Bootup	Revertible
mac	131072	2097152	2097152	131072	2097152	Yes
vlan	512	4095	512	512	512	No
spanning-tree	32	128	32	32	32	No
rstp	32	128	32	32	32	No
ip-arp	8192	65536	65536	65536	65536	No
multicast-route (IPv6)	8192	153600	8192	8192	8192	Yes
pim-mcache (IPv6)	4096	4096	4096	4096	4096	Yes
ip-cache	204800	1048576	524288	204800	524288	Yes
...						

Information for the configurable tables appears under the columns shown in bold type. To simplify configuration, the command parameter you enter to configure the table is used for the table name.

2 Enabling or disabling layer 2 switching

Example : To increase the size of the IP route table

```
Brocade(config)# system-max ip-route 120000
Brocade(config)# write memory
Brocade(config)# exit
Brocade# reload
```

NOTE

If you enter a value that is not within the valid range of values, the CLI will display the valid range for you.

To increase the number of IP subnet interfaces you can configure on each port on a Brocade device to 64, enter the following commands.

```
Brocade(config)# system-max ip-subnet-port 64
Brocade(config)# write memory
Brocade(config)# exit
Brocade# reload
```

Syntax: [no] system-max ip-subnet-port *num*

The *num* parameter specifies the maximum number of subnet addresses per port. The minimum, maximum and default values for this parameter are described in [Table 9](#).

NOTE

You must reload the software for the change to take effect.

Enabling or disabling layer 2 switching

By default, Brocade devices supports routing over layer 2 switching. You can enable layer 2 switching globally or on individual port using the **no route-only** command.

NOTE

On the Brocade Netron CES and Brocade Netron CER, the **route-only** command should not be configured on untagged MPLS uplinks when using it for VPLS or VLL. Otherwise, incoming VPLS or VLL traffic is dropped.

The **no route-only** and **route-only** commands prompts you for whether or not you want to change the “route-only” behavior. You must enter **y** if you want to proceed or **n** if you do not. The prompt is displayed as shown in the following examples of the **no route-only** and **route-only** commands.

NOTE

Always perform a reload after removing a route-only config or enabling route-only. Removing or enabling the route-only option without a reload will cause multicast issues.

To enable Layer 2 switching globally, enter the following.

```
Brocade(config)# no route-only
This will change the route-only behavior at the global level.
Are you sure? (enter 'y' or 'n'): y
Global 'route-only' committed.
```

To globally disable Layer 2 switching on a Brocade device and return to the default (route-only) condition, enter commands such as the following:

```
Brocade(config)# route-only
This will change the route-only behavior at the global level.
Are you sure? (enter 'y' or 'n'): y
Global 'no route-only committed.
```

Syntax: [no] route-only

NOTE

On the Brocade NetIron XMR and Brocade MLX series devices, **route-only** is the default condition. Because **route-only** is the default condition, it will not be displayed in the configuration. If you use **no route-only** to enable switching, the **no route-only** command will be displayed in the configuration.

NOTE

On the Brocade NetIron CES device, **route-only** is disabled by default. Therefore, if **route-only** is enabled on a Brocade NetIron CES device, it will be displayed in the configuration.

To enable Layer 2 switching only on a specific interface, go to the Interface configuration level for that interface, and add the **no route-only** command. The following commands show how to enable Layer 2 switching on port 3/2.

```
Brocade(config)# interface ethernet 3/2
Brocade(config-if-e10000-3/2)# no route-only
```

Syntax: [no] route-only

To re-enable the default **route-only** condition on port 3/2, enter the **route-only** command as shown.

```
Brocade(config-if-e10000-3/2)# route-only
```

When **route-only** is enabled on a physical interface, incoming unknown unicast packets are not sent to the CPU and are dropped locally by the hardware.

NOTE

Configuring **route-only** on a physical interface affects incoming frames only. In other words, interface **route-only** disables L2 switching for incoming frames but does not disable L2 switching for outgoing frames. If the **route-only** interface is a member of a VLAN, the interface will still transmit frames received on other interfaces of that VLAN if those other interfaces still have L2 switching enabled. To prevent this from happening, make sure that any interface you have configured for **route-only** are not also members of VLANs where you are intentionally performing L2 switching.

Configuring static MAC addresses

You can assign static MAC addresses to ports of a Brocade device.

You can manually input the MAC address of a device to prevent it from being aged out of the system address table, to prevent traffic for a specific device, such as a server, from flooding the network with traffic when it is down, and to assign higher priorities to specific MAC addresses.

Static MAC addresses are configured within a specified VLAN including the default VLAN 1. Optionally you can specify a port priority (QoS).

The default and maximum configurable MAC table sizes can differ depending on the device. To determine the default and maximum MAC table sizes for your device, display the system parameter values. Refer to [“Displaying and modifying default settings for system parameters”](#) on page 43.

2 Changing the MAC age time

The ability of the CAM to store depends on the following:

- The number of source MAC address being learned by the CAM.
- The number of destination MAC addresses being forwarded by the CAM
- The distribution of the MAC address entries across ports. For example, if one port is learning all the source MAC addresses, the available of the CAM for that port will be depleted.

Example

In the following example, a static MAC address of 0000.0063.67FF with a priority of 7 is assigned to port 2 of module 1 in VLAN 200.

```
Brocade(config)# vlan 200
Brocade(config)# static-mac-address 0000.0063.67FF e 1/2 priority 7
```

Syntax: **[no] static-mac-address** *mac-addr* **ethernet** *portnum* **[priority** *number* **]**

The *mac-addr* variable specifies the MAC address that you assigning.

The **portnum** variable specifies the Ethernet port that the MAC address is being assigned to.

Using the **priority** option, you can assign a value to the *number* variable of 0 – 7.

Changing the MAC age time

The MAC age time sets the aging period for ports on the device, defining how long (how many seconds) a port address remains active in the address table.

To change the aging period for MAC addresses from the default of 300 seconds to 600 seconds.

```
Brocade(config)# mac-age-time 600
```

Syntax: **[no] mac-age-time** *age-time*

The *age-time* can be 0 or a number from 60 – 65535. The zero results in no address aging. The default is 300 (seconds).

Configuring static ARP entries

When you create a static ARP entry, the Brocade device automatically creates a static MAC entry.

NOTE

To delete the static MAC entry, you must delete the static ARP entry first.

Configuring system max values

The system max values for the several system parameters of the Brocade devices are described in [Table 9](#)

TABLE 9 System max values for Brocade NetIron XMR and Brocade MLX series devices

Parameter	Minimum value for Brocade MLX series	Maximum value for Brocade MLX series	Default value for Brocade MLX series	Minimum value for Brocade NetIron XMR	Maximum value for Brocade NetIron XMR	Default value for Brocade NetIron XMR
config-file-size	2097152	16777216	8388608	2097152	16777216	8388608
gre-tunnels	1	8192	256	1	8192	256
hw-flooding	0	4095	0	0	4095	8
ifl-cam	0	81920	0	0	81920	0
ip-arp	2048	65536	8192	2048	65536	8192
ip-cache	8192	524288	102400	8192	1048576	204800
ip-filter-system	1024	40960	4096	1024	40960	4096
ip-route	4096	524288	102400	4096	1048576	204800
ipv4-mcast-cam	0	32768	4096	0	65536	8192
ip-subnet-port	24	128	24	24	128	24
vrf-route	128	450560	5120	128	450560	5120
ipv6-cache	8192	102400	32768	8192	245760	65536
ipv6-mcast-cam	0	8192	1024	0	16384	2048
ipv6-route	4096	114688	32768	4096	245760	65536
l2-acl-table-entries	64	256	64	64	256	64
mac	4000	1048576	32768	4000	2097152	131072
mgmt-port-acl-size	1	100	20	1	100	20
subnet-broadcast-acl-cam	0	4096	0	0	4096	0
receive-cam	512	8192	1024	512	8192	1024
rstp	1	128	32	1	128	32
session-limit	1024	40960	8192	1024	163840	32768
spanning-tree	1	128	32	1	128	32
virtual-interface	40	4095	255	40	4095	255
vlan	2	4095	512	2	4095	512
vpls-mac	32	262144	2048	32	1000000	8192
vpls-num	512	4096	512	1024	16384	2048
ecmp-pram-block-size	8	32	32	8	32	32

2 Configuring system max values

TABLE 10 System max values for Brocade NetIron CES, Brocade NetIron CER, and Brocade NetIron CER-RT devices

Parameter	Minimum value for Brocade NetIron CES	Maximum value for Brocade NetIron CES	Default value for Brocade NetIron CES	Minimum value for Brocade NetIron CER	Maximum value for Brocade NetIron CER	Default value for Brocade NetIron CER	Minimum value for Brocade NetIron CER-RT	Maximum value for Brocade NetIron CER-RT	Default value for Brocade NetIron CER-RT
config-file-size	2097152	16777216	8388608	2097152	16777216	8388608	2097152	16777216	8388608
ip-arp	2048	16384	4096	2048	16384	4096	2048	16384	4096
ip-cache	4096	32768	16384	4096	524288	290816	4096	1572864	290816
ip-filter-sys	1024	32768	4096	1024	32768	4096	1024	32768	4096
ip-route	4096	32768	16384	4096	524288	290816	4096	1572864	290816
ip-subnet-port	24	128	24	24	128	24	24	128	24
l2-acl-table-entries	64	256	64	64	256	64	64	256	64
mac	4000	131072	56320	4000	131072	56320	4000	131072	65536
mgmt-port-acl-size	1	100	20	1	100	20	1	100	20
rstp	1	128	32	1	128	32	1	128	32
session-limit	1024	32768	32768	1024	32768	32768	1024	32768	32768
spanning-tree	1	128	32	1	128	32	1	128	32
virtual-interface	40	1024	255	40	4095	255	40	4095	255
vlan	512	4095	512	2	8192	512	2	8192	512
vrf	1	16	1	1	128	16	1	128	16
ipv6-cache	1024	131072	1024	1024	131072	1024	1024	262141	8192
ipv6-route	1024	131072	1024	1024	131072	8192	1024	262141	8192
vrf-route	1024	32768	1024	1024	32768	1024	1024	1572864	1024
ip-tunnels	32	128	32	32	128	32	32	256	32

NOTE

Default values are the same irrespective of the software package on the Brocade NetIron CES and Brocade NetIron CER devices.

NOTE

The maximum FIB scalability for Brocade NetIron CER and Brocade NetIron CES has been tested using an internet route mix. When using route prefixes concentrated in a narrow prefix length range, the scalability numbers will be lower. It is important to design your network keeping this in mind.

To configure system-max values, use the following command.

Syntax: [no] system-max config-file-size | gre-tunnels | ip-arp | ip-cache | ip-filter-sys | ip-route | ip-static-arp | ipv4-mcast-cam | ip-subnet-port | ip-tunnels | vrf | vrf-route | ipv6-cache | ipv6-mcast-cam | ipv6-route | l2-acl-table-entries | ifl-cam | mac | mgmt-port-acl-size | subnet-broadcast-acl-cam | receive-cam | rstp | session-limit | spanning-tree | trunk-num | virtual-interface | vlan | vpls-mac | vpls-num | ecmp-pram-block-size

The **gre-tunnels** parameter sets the maximum number of GRE tunnels. For minimum, maximum and default values for this parameter refer to [Table 9](#).

The **config-file-size** parameter sets the allowed running and startup-config file sizes. Refer to the appropriate table for your platform. For minimum, maximum and default values for this parameter refer to [Table 9](#) and [Table 10](#).

The **ifl-cam** parameter sets the maximum number of Internal Forwarding Lookup Identifiers. These are used when configuring a Local VLL for Dual Tagging. The default value for the **ifl-cam** parameter is 8K. The maximum values for this parameter are different depending on which CAM partition you have configured on your system. For minimum, maximum and default values by CAM partition for this parameter, refer to [Table 11](#).

The **ip-arp** parameter sets the maximum number of ARP entries. For minimum, maximum and default values for this parameter refer to [Table 9](#) and [Table 10](#).

The **ip-cache** parameter sets the maximum size of the IP cache. For minimum, maximum and default values for this parameter refer to [Table 9](#) and [Table 10](#).

The **ip-filter-sys** parameter sets the maximum number of IP ACL entries. For minimum, maximum and default values for this parameter refer to [Table 9](#) and [Table 10](#).

The **ip-route** parameter sets the maximum number of IP Route entries. For minimum, maximum and default values for this parameter refer to [Table 9](#) and [Table 10](#).

NOTE

There is no need to configure a system-max value for static ARP entries.

The **ip-static-arp** parameter sets the maximum number of static ARP entries. For minimum, maximum and default values for this parameter refer to [Table 9](#).

The **ipv4-mcast-cam** parameter allows you to configure the maximum CAM size for an IPv4 multicast group. For minimum, maximum and default values for this parameter refer to [Table 4.5](#). To configure the CAM size of an IPv4 multicast group, refer to “[Configuring CAM size for an IPv4 multicast group](#)” on page 55.

The **ip-subnet-port** parameter sets the maximum number of IP subnets per port. For minimum, maximum and default values for this parameter refer to [Table 9](#) and [Table 10](#).

The **vrf-route** parameter sets the maximum number of VRF routes per VRF instance. For minimum, maximum and default values for this parameter refer to [Table 9](#).

The **ipv6-cache** parameter sets the maximum size of the IPv6 cache. For minimum, maximum and default values for this parameter refer to [Table 9](#).

The **ipv6-mcast-cam** parameter allows you to configure the maximum CAM size for an IPv6 multicast group. For minimum, maximum and default values for this parameter refer to [Table 8](#). To configure the CAM size of an IPv6 multicast group, refer to “[Configuring CAM size for an IPv6 multicast group](#)” on page 56.

The **ipv6-route** parameter sets the maximum number of IPv6 routes. For minimum, maximum and default values for this parameter refer to [Table 9](#).

NOTE

The **system-max ipv6-route** command can be configured with a maximum value of 114688 on the Brocade MLX series, but the Brocade device system will only support a maximum value of 114687 for IPv6 routes.

The **l2-acl-table-entries** parameter sets the maximum number of layer-2 ACL entries per ACL table. For minimum, maximum and default values for this parameter refer to [Table 9](#) and [Table 10](#).

2 Configuring system max values

The **mac** parameter sets the maximum number of MAC entries.

For minimum, maximum and default values for this parameter refer to [Table 9](#) and [Table 10](#).

The **mgmt-port-acl-size** parameter sets the maximum size for a management port ACL.

For minimum, maximum and default values for this parameter refer to [Table 9](#) and [Table 10](#).

The **subnet-broadcast-acl-cam** parameter sets the maximum number of IP broadcast ACL CAM entries. For minimum, maximum, and default values for this parameter, refer to [Table 9](#).

The **receive-cam** parameter sets the maximum number of IP Receive ACL software CAM entries.

For minimum, maximum and default values for this parameter refer to [Table 9](#).

The **rstp** parameter sets the maximum number of RSTP instances.

For minimum, maximum and default values for this parameter refer to [Table 9](#) and [Table 10](#).

The **session-limit** parameter sets the maximum number of sessions.

For minimum, maximum and default values for this parameter refer to [Table 9](#) and [Table 10](#).

The **spanning-tree** parameter sets the maximum number of spanning-tree instances.

For minimum, maximum and default values for this parameter refer to [Table 9](#) and [Table 10](#).

The **virtual-interface** parameter sets the maximum number of virtual interfaces.

For minimum, maximum and default values for this parameter refer to [Table 9](#) and [Table 10](#).

The **vlan** parameter sets the maximum number of VLANs.

For minimum, maximum and default values for this parameter refer to [Table 9](#) and [Table 10](#).

The **vpls-mac** parameter sets the maximum number of VPLS MAC Entries.

For minimum, maximum and default values for this parameter refer to [Table 9](#).

The **vpls-num** parameter sets the maximum number of Virtual Private LAN Services (VPLS). For minimum, maximum and default values for this parameter refer to [Table 9](#)

The **ecmp-pram-block-size** parameter is used as a limiting factor when programming ECMP nexthops (v4/v6/vpnv4/vpnv6). Even though the control plane supports up to 32 nexthops per a route, the actual number of nexthops which are programmed in HW is controlled by this command. If **system-max ecmp-pram-block-size** is configured to a value lesser than the value configured for **ip load-sharing** or **ipv6 load-sharing**, or if **ip load-sharing** or **ipv6 load-sharing** is configured to a value greater than that configured for **system-max ecmp-pram-block-size**, a warning message will be displayed. For minimum, maximum and default values for this parameter refer to [Table 9](#)

TABLE 11 System max **ifi-cam** values available by CAM profile on Brocade NetIron XMR and Brocade MLX series

CAM profile	Minimum value	Maximum value	Default value
Default	0	57344	8192
ipv4	0	114688	8192
ipv6	0	131072	8192
l2-metro	0	114688	8192
mpls-l3vpn	0	114688	8192
mpls-vpls	0	114688	8192
multi-service	0	49152	8192
multi-service-2	0	81920	8192
vpn-vpls	0	114688	8192
ipv4-vpn	0	114688	8192

TABLE 11 System max **ifl-cam** values available by CAM profile on Brocade NetIron XMR and Brocade MLX series (Continued)

CAM profile	Minimum value	Maximum value	Default value
I2-metro-2	0	114688	8192
mpls-l3vpn-2	0	114688	8192
mpls-vpls-2	0	114688	8192
ipv4-ipv6	0	114688	8192
ipv4-vpls	0	114688	8192
ipv4-ipv6-2	0	81920	8192

Configuring CAM size for an IPv4 multicast group

To configure the CAM size of an IPv4 multicast group, enter the following command.

```
Brocade(config)# system-max ipv4-mcast-cam
```

Syntax: [no] system-max [ipv4-mcast-cam]

By default, no system-max parameter is configured.

The **ipv4-mcast-cam** parameter allows you to specify the maximum CAM size you want for an IPv4 multicast group.

The **decimal** parameter specifies the range that is supported for configuring the CAM size. On the Brocade NetIron XMR, the minimum value supported is 0, and the maximum value supported is 65536. The default value is 8192. On the Brocade MLX series, the minimum value supported is 0, and the maximum value supported is 32768. The default value is 4096.

Upon configuration, the Brocade system will verify the input value with the amount of CAM resources that are available. If the Brocade system is unable to allocate requested space, it will display the following error message on the Brocade NetIron XMR and Brocade MLX series.

```
Brocade(config)#system-max ipv4-mcast-cam 60000
Error - IPv4 Multicast CAM (60000) exceeding available CAM resources
Total IPv4 ACL CAM: 98304(Raw Size)
IPv4 Receive ACL CAM: 16384(Raw Size)
IPv4 Source Guard CAM: 32(Raw Size)
Reserved IPv4 Rule ACL CAM: 1024(Raw Size)
Available IPv4 Multicast CAM: 80864(Raw Size) 40432(User Size)
```

If there is not enough CAM resources available to change the cam-partition profile from IPv4 to IPv6, the following message is displayed.

```
Brocade(config)#cam-partition profile ipv4-ipv6
Error - IPv4 Receive ACL CAM (8192) exceeding available CAM resources
Total IPv4 ACL CAM: 98304(Raw Size)
IPv4 Multicast CAM: 120000(Raw Size)
IPv4 Source Guard CAM: 32(Raw Size)
Reserved IPv4 Rule ACL CAM: 1024(Raw Size)
Available IPv4 Receive ACL CAM: 0(Raw Size) 0(User Size)
Error - Failed to select this CAM profile
```

2 Configuring CAM size for an IPv6 multicast group

This error message is also displayed on the Brocade MLX series.

After you issue the `system-max` command, with `ipv4-mcast-cam` parameter included, additional information will display on the Brocade NetIron XMR and Brocade MLX series as shown in the following example.

```
Brocade(config)#system-max ipv4-mcast-cam 60000
Reload required. Please write memory and then reload or power cycle.
Failure to reload could cause system instability on failover.
Newly configured system-max will not take effect during hitless-reload.
```

NOTE

You must write this command to memory and perform a system reload for this command to take effect.

Configuring CAM size for an IPv6 multicast group

To configure the CAM size of an IPv6 multicast group, enter the following command.

```
Brocade(config)# system-max ipv6-mcast-cam
```

Syntax: [no] system-max [ipv6-mcast-cam]

By default, no `system-max` parameter is configured.

The `ipv6-mcast-cam` parameter allows you to specify the maximum CAM size you want for an IPv6 multicast group.

The **decimal** parameter specifies the range that is supported for configuring the CAM size. On the Brocade NetIron XMR, the minimum value supported is 0, and the maximum value supported is 16384. The default value is 2048. On the Brocade MLX series, the minimum value supported is 0, and the maximum value supported is 8192. The default value is 1024.

Upon configuration, the Brocade system will verify the input value with the amount of CAM resources that are available. If the Brocade system is unable to allocate requested space, it will display the following error messages on the Brocade NetIron XMR and Brocade MLX series:

On the Brocade NetIron XMR.

```
Brocade(config)# system-max ipv6-mcast-cam 15000
Error - IPV6 Multicast CAM (15000) exceeding available CAM resources
Total IPv6 ACL CAM: 32768(Raw Size)
Reserved IPv6 Rule ACL CAM: 1024(Raw Size)
Available IPv6 Multicast CAM: 31744(Raw Size) 3968(User Size)
```

On the Brocade MLX series.

```
Brocade(config)#system-max ipv6-mcast-cam 8000
Error - IPV6 Multicast CAM (8000) exceeding available CAM resources
Total IPv6 ACL CAM: 16384(Raw Size)
Reserved IPv6 Rule ACL CAM: 1024(Raw Size)
Available IPv6 Multicast CAM: 15360(Raw Size) 1920(User Size)
```

After you issue the `system-max` command, with `ipv6-mcast-cam` parameter included, additional information will display on the Brocade NetIron XMR and Brocade MLX series as shown in the following example.

```
Brocade(config)#system-max ipv6-mcast-cam 1000
Reload required. Please write memory and then reload or power cycle.
Failure to reload could cause system instability on failover.
Newly configured system-max will not take effect during hitless-reload.
```

NOTE

You must write this command to memory and perform a system reload for this command to take effect.

Configuring profiles with a zero-size IPv4 or IPv6 ACL

When a profile is configured to a zero-size IPv4 or IPv6 ACL, the minimum value of 512 for IPv4, and the minimum value of 128 for IPv6 is disabled.

If the system-max value for multicast is configured, and you select a value of 0 for IPv4 or IPv6 ACL, the system-max value will be ignored. There is minimal checking for errors. The following warning message is displayed.

```
Brocade(config)#cam-partition profile mpls-vpls-2
Warning - Changing to a profile with zero Ipv6 ACL CAM size, ignoring system-max
value check and minimum-guarantee checks.
Reload required. Please write memory and then reload or power cycle.
Failure to reload could cause system instability on failover.
Newly configured system-max will not take effect during hitless-reload.
```

NOTE

You must write this command to memory and perform a system reload for this command to take effect.

The system-max value is retained in the configuration so that it can be reused later when you want to change the CAM profile.

Maintaining system-max configuration with available system resources

When system-max values are configured, the Brocade system checks for available system resources. The system resources are required in order to maintain dynamic memory allocation. System-max values are checked at the configuration time, and at the bootup time. If there are insufficient system resources available on the Management Module, this will cause the configuration to be rejected during card bootup. On the Interface Module, insufficient system resources will lead to failure in booting up the card.

Configuration time

When system-max values are configured, the Management Module calculates the memory required to accept the value. The resulting value is checked against the Known-Available-Memory value, and calculated against the Highest Required Memory value for both the Management Module and the Interface Module.

The Known-Available-Memory is a value with the Lowest Supported Available Memory on a node. For example, if a node can accept a 1 Gigabyte LP, and a 512 MB LP, then the 512 MB LP will be used. The Highest Required Memory is a value with most amount of memory available on a node. For example, if a node has both 2 PPCR LP, and 1 PPCR LP, then the 2 PPCR LP will be used.

If the new system-max value is accepted, then the configuration will also be accepted. The following information will display.

```
Brocade(config)#system-max mac 4000
Reload required. Please write memory and then reload or power cycle.
Failure to reload could cause system instability on failover.
Newly configured system-max will not take effect during hitless-reload.
```

NOTE

You must write this command to memory and perform a system reload for this command to take effect.

If the new system-max value is not accepted, then the configuration is rejected. The following error message is printed on the console.

```
Brocade(config)# system-max ipv4 10000
ERROR: Configured System-max value cannot be accommodated.
```

Bootup time

At bootup time, the Management Module will repeat the same process as done in the Configuration time. The Management Module calculates the memory required to accept the system-max configuration. The resulting value is checked against the Known-Available-Memory value for both the Management Module and the Interface Module.

After the new system-max value is configured, there are three possible configuration outcomes. The three possible configuration outcomes are described below.

1. The configuration can be accommodated, but leaves only 10% of Available Memory

In this configuration, a check is made against 90% of Available Memory. If the difference between the Required Available Memory and the Available Memory is less than 10% of Available Memory, then the configuration is accepted. The following warning message is displayed on the console if it affects the Management Module or Interface Module.

The following warning message is displayed on the Management Module.

```
WARNING: Configured System-max Leaves less than 10 percent Free Memory
Available on MP.
```

The following warning message is displayed on the Interface Module.

```
WARNING: Configured System-max Leaves less than 10 percent Free Memory
Available on LP.
```

A syslog message showing the required memory versus the available memory is generated, and a similar warning message is displayed on the Management Module and Interface Module as shown in the following example.

```
Brocade# show log
```

```
...
```

```
Jan 17 22:55:55:N: WARN: Configured System-max Leaves less than 10 percent Free Memory Available on MP (162529285 req vs 1625292800 available)
```

```
Jan 17 22:55:55:N: WARN: Configured System-max Leaves less than 10 percent Free Memory Available on LP (162529285 req vs 1625292800 available)
```

NOTE

When the system is booted up again, the percent of free memory is discretionary and is only an estimate.

NOTE

Even if all elements are configured with the maximum allowed value, you may not see the reversion of system-max values that occur on any given Interface Module.

NOTE

Notifications and traps are sent with the same message.

2. The configuration can be easily accommodated.

In this configuration, the Management Module continues to use the configured system-max value, and send the same value to the installed Interface Modules.

3. The configuration cannot be accommodated.

If the configured system-max value cannot be used, the Management Module will locate the elements that can be reverted to a default value. These system-max elements will revert to a default value, and the following message will display on the console.

```
WARN: Configured System-max cannot be accommodated. Resetting revertible elements to default values.
```

A syslog message is generated, and a similar warning message is displayed on the Management Module and Interface Module as shown in the following example.

```
Brocade# show log
```

```
...
```

```
Jan 17 22:55:55:N: WARN: Configured System-max cannot be accommodated on MP (1625292801 req vs 1625292800 available). Resetting revertible elements to default values.
```

```
Jan 17 22:55:55:N: WARN: Configured System-max cannot be accommodated on LP (1625292801 req vs 1625292800 available). Resetting revertible elements to default values.
```

NOTE

Once the system-max have been reverted, a user might not be able to configure any system-max until configuration for some or all of the revertible system-max elements is removed using “no system-max” CLI.

NOTE
 Notifications and traps are sent with the same message.

The following tables show which elements are revertible (Yes or No) in each element category.

L2 elements

TABLE 12 L2 elements

L2 elements	Revertible: yes or no
Mac	yes
Vlan	no
Spanning-tree	no
Rstp	no

L3 elements

TABLE 13 L3 elements

L3 elements	Revertible: yes or no
Arp	no
multicast-route (for v6 only)	yes
pim-mcache	yes
ip-cache	yes
ip-route	yes
ip-subnet-port	no
virtual-interface	no

VPLS elements

TABLE 14 VPLS elements

VPLS elements	Revertible: yes or no
vpls-mac (MAX_VPLS_MAC_IN DEX)	yes
vpls-num (MAX_VPLS_NUM_IN DEX)	no

Miscellaneous elements

TABLE 15 Miscellaneous elements

Miscellaneous elements	Revertible: yes or no
session-limit	yes
ip-filter-sys	no
mgmt-port-acl-size	no
l2-acl-table-entries	no
ipv6-cache	yes
ipv6-route	yes
IPVRF MAX ROUTES	yes
mgmt-port-acl-size	no
receive-cam	no
IPGRE	no
LSP_ACL	no
SERVICE_LOOKUP	no
IP_SRC_GUARD_CAM	no
IPv4 MCAST CAM	no
IPv6 MCAST CAM	no
SERVER_TRUNKS	no
CONFIG_FILE_SIZE	no

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.

```
Brocade# 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.

2 Switch fabric fault monitoring

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.

```
Brocade# 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.

```
Brocade#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 Brocade NetIron XMR or Brocade MLX series chassis. 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 Brocade NetIron XMR or Brocade MLX series chassis using the following command. Each line represents a link between an SFM and an interface module (LP).

```

Brocade#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 16 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.

TABLE 16 CLI display of SFM link information (Continued)

This field...	Displays...
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.

```
Brocade> 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.

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

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

```
Brocade# 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.

```
Brocade(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 66) 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.

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.

Example

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

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

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

```
Brocade(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:

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

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

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

Syntax: [no] sysmon tm link auto-tune

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 Brocade 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 65), 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 64.

```
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 64), 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.

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.

```
Brocade#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.

```
Brocade#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 [sfm-number | all]`

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.

Verifying an image checksum

Use the **image-checksum** command to verify the checksum of the application, boot, or monitor images that are saved in code flash and Auxiliary Flash cards.

NOTE

The **image-checksum** command on is not applicable to a combined application image.

2 Displaying information for an interface for an Ethernet port

To check a monitor image, use the following command.

```
Brocade# image-checksum monitor
OK
```

Syntax: [no] image-checksum *file-name*

The *file-name* variable specifies the image file that you want to verify the checksum for.

The following output can be generated by this command

TABLE 17 Output from image-checksum command

Output	Description
File not found	The device failed to locate the specified file.
Failed to read file	The device failed to obtain the file length from the file system.
Not an image file	The specified file is not an image file.
File read failed	The specified file's actual length is different from the file length stored in the file system.
Checksum failed	The image has a checksum error.
OK	The checksum has been verified for the specified image file.

Displaying information for an interface for an Ethernet port

To display information for a show interface for an ethernet port, enter the following command at any CLI level.

```
Brocade# show interface ethernet 9/1
GigabitEthernet2/3 is up, line protocol is up
  STP Root Guard is disabled, STP BPDU Guard is disabled
  Hardware is GigabitEthernet, address is 0000.0098.4900 (bia 0000.0098.492a)
  Configured speed auto, actual 1Gbit, configured duplex fdx, actual fdx
  Member of VLAN 1 (untagged), 5 L2 VLANS (tagged), port is in dual mode (default
vlan), port state is Forwarding
  STP configured to ON, Priority is level0, flow control enabled
  Priority force disabled, Drop precedence level 0, Drop precedence force disabled
  arp-inspection-trust configured to OFF
  mirror disabled, monitor disabled
  Not member of any active trunks
  Not member of any configured trunks
  Port name is ->7.blade1.shelf1.access.aprd
  MTU 1544 bytes, encapsulation ethernet
  300 second input rate: 1509512 bits/sec, 713 packets/sec, 0.15% utilization
  300 second output rate: 1992071 bits/sec, 751 packets/sec, 0.20% utilization
  712896623 packets input, 204984611768 bytes, 0 no buffer
  Received 1315502 broadcasts, 53313 multicasts, 711527808 unicasts
  0 input errors, 0 CRC, 0 frame, 0 ignored
  0 runts, 29433839 giants
  NP received 712896745 packets, Sent to TM 712839428 packets
  NP Ingress dropped 57317 packets
```

```
796106728 packets output, 366570033985 bytes, 0 underruns
Transmitted 2045784 broadcasts, 32330616 multicasts, 761730328 unicasts
0 output errors, 0 collisions
NP transmitted 796106833 packets, Received from TM 796534170 packets
```

Syntax: `show interface [ethernet slot-port [to slot-port]]`

You can display information for all ports in a device by using the **show interface** command without options, or use the **ethernet slot-port** option to limit the display to a single port, or add the **to slot-port** option for a range of ports.

Displaying the full port name for an Ethernet interface

To display the full port name for an ethernet interface using the CLI, enter the following command.

```
Brocade# show interface brief slot 3
Port Link Port-State Dupl Speed Trunk Tag Priori MAC Name Type
3/1 Up Forward Full 100G None No level0 0000.0002.025c default-port
3/2 Up Forward Full 100G None No level0 0000.0002.025d default-port
```

Syntax: `show interface brief slot/port`

If the port is logically UP (meaning not LK-DISABLE or LACP-BLOCKED or OAM-DISABLE or DOT1X-BLOCKED), then:

- If the port is untagged then the L2 Port state field indicates the STP State of Port in the untagged VLAN context.
- If the port is tagged or in dual mode (both tagged and untagged), then it is marked forwarding as a single port state cannot be determined.

In case Port is logically down, L2 Port State indicates reason for Logical Port down condition (LK-DISABLE or LACP-BLOCKED or OAM-DISABLE or DOT1X-BLOCKED)

Using the **show interface brief wide** command long port names are displayed. If the **show interface brief wide** command is not used only partial names are displayed in cases of long port names.

```
Brocade# show interface brief wide
Port Link Port-State Speed Tag MAC Name
2/1 Up Forward 10G No 0000.00f7.0230 port-connected-to-chicago
2/2 DisabNone None No 0000.00f7.0231
2/3 DisabNone None No 0000.00f7.0232
2/4 DisabNone None No 0000.00f7.0233

Port Link Port-State Speed Tag MAC Name
mgmt1 Up Forward 100M Yes 0000.00f7.0200

Port Link Port-State Speed Tag MAC Name
lb1 Up N/A N/A N/A N/A
Brocade#
```

Syntax: `show interface brief wide slot/port`

2 Displaying information for an interface for an Ethernet port

TABLE 18 Display of show interface ethernet port

This field...	Displays...
Module type port# is state	The <i>module type</i> variable specifies a type of interface module, such as 10GigabitEthernet. The <i>port#</i> variable specifies the port number for the interface module. The <i>state</i> variable if the interface module is up or down.
Line protocol is status	The <i>status</i> variable specifies if the line protocol is up or down. If the interface is down due to Remote Fault, the reason is indicated as: "(remote fault)". If a port is down because of a Local Fault, the reason is indicated as: "(local fault)".
STP Root Guard is status	The <i>status</i> variable specifies if the STP Root Guard is enabled or disabled.
STP BPDU Guard is status	The <i>status</i> variable specifies if the STP BPDU Guard is enabled or disabled.
Hardware is module type	The <i>module type</i> variable specifies a type of interface module, such as #GigabitEthernet.
Address is MAC- address	The <i>MAC- address</i> variable specifies the MAC address of the port.
Configured speed and actual speed	The speed that the module has been configured to operate at, and the actual speed it is currently operating at.
Configured port speed and actual duplex	The port capacity that the module has been configured to operate at, and the actual speed it is currently operating at.
Member of VLAN # (untagged) port# L2 VLANS (tagged) Port is in dual mode/untagged/tagged mode Port state is status	The <i>VLAN# (untagged)</i> variable specifies a port that is a member of only 1 VLAN. The <i>port# L2 VLANS (tagged)</i> variable specifies a port that is a member of multiple ports and untagged. A port is in <i>dual- mode</i> specifies member VLAN ports as untagged and tagged. The default mode is dual-mode. The <i>status</i> variable identifies the flow of traffic as forwarding or disabled.
STP configured to status Priority level Flow control status	The <i>status</i> variable specifies if the STP is ON or OFF. The priority level assigned to the port-based VLAN. The priority level is on scale from 0-7. The default is 0. The <i>status</i> variable is enabled or disabled.
Priority force status	The <i>status</i> variable specifies if the priority force on a port is disabled on enabled.
Drop precedence level value	Identifies the TOS or DSCP value in the IPv4 or IPv6 packet header. The <i>value</i> variable specifies the drop precedence on a scale from 0-3. Packets that contain a DSCP value of 0 are least likely to be dropped and packets with a value of 3 are most likely to be dropped. The default value is 0.
Drop precedence force status	The <i>status</i> variable specifies the drop precedence force as enabled or disabled. Identifies the drop precedence if the force command is configured for a specific ingress port.
arp-inspection-trust configured to status	The <i>status</i> variable specifies if arp-inspection-trust feature is configured ON or OFF. The default trust setting for a port is untrusted.
Mirror status	The <i>status</i> variable specifies if the port mirror command is configured as enabled or disabled.
Monitor status	The <i>status</i> variable specifies if the port monitor command is configured as enabled or disabled.

TABLE 18 Display of show interface ethernet port (Continued)

This field...	Displays...
<i>Trunk membership</i>	The <i>Trunk membership</i> variable identifies the interface module as a member of a primary or secondary port. This specifies members of an active port or not a member of an active port.
<i>Configured trunk membership</i>	The <i>Configured trunk membership</i> variable identifies the interface module as a member of any configured trunk or not a member of a configured trunk.
<i>Port name</i>	The <i>port name</i> variable identifies the name assigned to the port.
MTU # bytes, encapsulation ethernet	Maximum Transmission Unit (MTU) refers to the size of the largest packet or frame that a known layer can pass forward. The # bytes variable refers to size of the packet or frame.
# seconds input rate: value bits/sec, value packets/sec, % utilization	The # seconds input rate refers to: <ul style="list-style-type: none"> • The value of bits received per second. • The value of packets received per second. • The % utilization specifies the port's bandwidth used by received traffic.
# seconds output rate: value bits/sec, value packets/sec, % utilization	The # seconds output rate refers to: <ul style="list-style-type: none"> • The value of bits transmitted per second. • The value of packets transmitted per second. • The % utilization specifies the port's bandwidth used by transmitted traffic.
value packets input, value bytes, value no buffer	<ul style="list-style-type: none"> • The value variable specifies the number of packets received. • The value variable specifies the number of bytes received. • The value no buffer variable specifies the total number of packets that have been discarded by the MAC device, due to temporary inability to store the packets before forwarding to the Network Processor (NP).
Received value broadcasts, value multicasts, value unicasts	The value variable specifies the amount of traffic the interface module is receiving on broadcasts, multicasts, and unicast traffic.
value input errors, value CRC, value frame, value ignored	<ul style="list-style-type: none"> • The value input errors variable specifies the number of received packets with errors. • The value CRC variable specifies the number of packets discarded by the MAC device due to detected CRC error. • The value variable specifies the number of received packets with alignment errors. • The value variable specifies the number of received packets that are discarded. <p>These parameters are not currently supported and will always display 0.</p>
value runts, value giants	<p>The value runts variable specifies the number of small packets that are less than 64 bytes.</p> <p>The value giants variable specifies the number of large packets greater than 1518 bytes.</p> <p>These parameters are not currently supported and will always display 0.</p>
NP received	The number of packets received on the Network Processor (NP).
NP transmitted	The number of packets sent from the Network Processor to the Traffic Manager (TM).
NP Ingress dropped	The number of ingress packets dropped on the Network Processor.
value packets output value bytes	<ul style="list-style-type: none"> • The value variable specifies the number of transmitted packets. • The value variable specifies the number of transmitted bytes.

2 Displaying information for an interface for an Ethernet port

TABLE 18 Display of show interface ethernet port (Continued)

This field...	Displays...
Transmitted <i>value</i> broadcasts, <i>value</i> multicasts, <i>value</i> unicasts	The <i>value</i> variable specifies the amount of traffic the interface module transmitted on broadcasts, multicasts, and unicast traffic.
<i>value</i> output errors, <i>value</i> collisions	<ul style="list-style-type: none">• The <i>value</i> variable specifies the number of transmitted packets with errors.• The <i>value</i> variable specifies the number of packets that experienced multi-access collisions. These parameters are not currently supported and will always display 0.
Network Processor transmitted <i>value</i> packets	The <i>value</i> variable specifies the number of packets transmitted from the Network Processor.
Received from Traffic Manager <i>value</i> packets	The <i>value</i> variable specifies the number of packets received by the Network Processor from the Traffic Manager.

Displaying statistics information for an Ethernet port

You can view statistical information about the traffic passing through a specified Ethernet port in one of two ways. The **monitor** commands allow you to monitor traffic statistics in real time, while the **show statistics** command provides a snapshot of the most recent traffic statistics.

Monitoring Ethernet port statistics in real time

You can monitor Ethernet traffic statistics in real time for a single port or traffic counters for all Ethernet ports using the **monitor** commands. When you execute a **monitor** command it retrieves and displays traffic statistics once per polling interval (2 seconds by default) until you pause or stop the display. The terminal window is fully occupied by the real-time display, and the command prompt is replaced by a footer listing options for pausing, canceling or modifying the display. When real-time monitoring is canceled, the command prompt is restored and the CLI resumes normal operation.

The following considerations affect the use of the **monitor** commands:

- Real-time monitor commands can be executed via Telnet, SSH, or a console session. Because of the slower communication rate in a console session, Brocade recommends executing the **monitor** commands *only* from a Telnet or SSH session. The default poll interval for telnet and SSH is 2 seconds, but the default polling interval for a console session is 8 seconds. If you execute **monitor** commands from a console session, flickering of the display may occur.
- If the **monitor** command is executed in a console session, console debug messages will not be displayed on the console screen.
- When the **monitor** command is executed via telnet or SSH, debug messages will not be displayed during execution of the command even with a **debug destination telnet session** configuration present.
- **monitor** commands, in general, display two kinds of statistics: aggregated (counted since system startup or since last cleared using a **clear** command) and delta (counted since start of this **monitor** command or since last cleared using the **c** footer option on the monitor screen).
- Resizing of the terminal window is not supported during real-time statistics display. You must stop the execution of the command before resizing the terminal window.
- Terminal display size must be at least 80 characters wide by 24 lines in order to avoid garbled or truncated display.
- Execution of the **monitor** commands is unaffected by Telnet or SSH idle timeouts; as long as the **monitor** command is running, the terminal is not idle.
- There can be a noticeable impact on CPU utilization if the polling interval (monitor refresh interval) is short and multiple sessions are simultaneously executing **monitor** commands. When monitoring takes place by way of multiple simultaneous sessions, increase the polling interval to minimize impact on the CPU. (The polling interval/refresh rate ranges from 2 to 30 seconds, with a default value of 2 seconds for SSH or telnet connections and 8 seconds for a console session.)
- When you quit the **monitor** command, the CLI command prompt will usually be displayed at the bottom of the screen. If it appears instead in the middle of the screen, clear the screen using the command **cls** before executing further commands.

Real-time monitoring of traffic statistics for a specific Ethernet port

To monitor traffic statistics for a specific Ethernet port, enter the following command at the Privileged EXEC level of the CLI.

```
Brocade# monitor statistics ethernet 1/2
```

Syntax: `monitor statistics ethernet slot/port`

The `slot/port` variable specifies the port for which you want to display statistics.

The **monitor statistics** command uses page mode display to show a detailed, port-specific traffic statistics screen which is updated every poll interval. (In the Brocade NetIron XMR and Brocade MLX series, this command also shows a second screen displaying network processor statistics.) You can modify the display using the commands shown in the footer. (Note that when you enter footer options they are not executed immediately but will be interpreted by the monitor engine during the next polling cycle.) The footer commands and their effects are described in Table 19.

TABLE 19 Footer commands for **monitor statistics** display

t	Displays the transmit/output statistics (the default) and continues the execution of the original command.
r	Displays the receive/input statistics and continues the execution of the original command.
n	Continues the execution of the command for the next available Ethernet interface. If there is no next interface available, the monitor continues to display statistics for the current interface.
p	Continues the execution of the command for the previous Ethernet interface. If there is no next interface available, the monitor continues to display statistics for the current interface.
c	Clears the current delta counters and continues the execution of the original command. To clear the aggregate counters, use the appropriate clear command.
f	Freezes the execution of the command; pauses retrieval and display of the statistics. While display is frozen, the only valid commands are s and q ; you can restart or quit the monitor, but any other command will be ignored.
s	Restarts the execution of the command; resumes retrieval and display of the statistics.
F	Decreases the polling interval (monitor refresh interval) by one second and continues the execution of the original command with the new refresh interval. This option will decrease the refresh interval until it is equal to 2 seconds, the minimum supported refresh interval value. The default value is 2 seconds. This command is not displayed in the footer of the statistics screen.
S	Increases the polling interval (monitor refresh interval) by one second and continues the execution of the original command with the new refresh interval. This option will increase the refresh interval until it is equal to 30 seconds, the maximum supported refresh interval value. The default value is 2 seconds. This command is not displayed in the footer of the statistics screen.
q or escape or ^c	Quits the execution of the command and returns to the command prompt.
u	Brocade NetIron XMR and Brocade MLX series only : Displays the first page of the multi-page display (page-up operation).
d	Brocade NetIron XMR and Brocade MLX series only : Displays the second page of the multi-page display (page-down operation).

Brocade NetIron XMR and Brocade MLX series example

Brocade# monitor statistics ethernet 4/1

Seconds: 8 poll: 8 Time: Aug 19 16:10:59

Page 1 of 2 Interface Tx Statistics Current Delta
 Ethernet 4/1 Tx interface statistics

Traffic statistics:

	Current	Delta
Out Packets	17083660926	533508
Out Octets	1093354299264	34144512
Out Unicast Packets	17083660926	533508
Out Multicast Packets	0	0
Out Broadcast Packets	0	0

Error statistics:

	Current	Delta
Out Errors	0	0
Out Discards	0	0

Tx/Rx=t/r, Page1/2=u/d, Next/Prev=n/p, Clear=c :Freeze=f/s Quit=q

Seconds: 40 poll: 8 Time: Aug 19 16:11:31

Page 2 of 2 NP Tx Statistics Current Delta
 Ethernet 4/1 Tx NP statistics

	Current	Delta
Sent to MAC Packet	17085805774	2670758
Raw Good Packet	17085805774	2670758
IPX HW Forwarded Packet	0	0
Receive from TM	17085805775	2670759
Unicast Packet	17085805774	2670758
Broadcast Packet	0	0
Multicast Packet	0	0

Error statistics :

	Current	Delta
Bad Packet Count	0	0

ACL Drop	0	0
Source Port Supress Drop	0	0
IPv4 Packet	0	0
IPv6 Packet	0	0
IPv4 Byte	0	0
IPv6 Byte	0	0

Tx/Rx=t/r, Page1/2=u/d, Next/Prev=n/p, Clear=c :Freeze=f/s Quit=q

The previous output shows the first and second pages of the detailed traffic statistics display for Ethernet port 4/1 from a Brocade NetIron XMR or Brocade MLX series, displaying transmit counters (the default).

2 Displaying statistics information for an Ethernet port

Brocade NetIron CES and Brocade NetIron CER example

```
Brocade# monitor statistics ethernet 1/2
```

```
Seconds: 26 poll: 2 Time: Aug 19 16:01:41

Current Delta
Ethernet 1/2 Tx interface statistics
Traffic statistics:
  In Packets          24847720          7738201
  In Octets          1590253440         495244864
  In Unicast Packets 24847720           7738201
  In Multicast Packets 0                  0
  In Broadcast Packets 0                  0

Error statistics:
  In Errors          0                  0
  In Discards        0                  0
```

```
Tx/Rx=t/r, Next/Prev=n/p, Clear=c :Freeze=f/s Quit=q
```

The previous output shows the detailed traffic statistics display for Ethernet port 1/2 from a Brocade NetIron CES or Brocade NetIron CER, displaying transmit counters (the default).

Real-time monitoring of traffic statistics for all Ethernet ports

To monitor summary traffic data (total packets or bytes sent and received) for all Ethernet ports (displaying up to 16 ports per screen), enter the following command at the Privileged EXEC level of the CLI.

Brocade# monitor interface traffic

```

                                Seconds: 248                Time: Mar 11 20:12:08
Interface traffic statistics:
      InPackets      Delta      OutPackets      Delta
e1/1      24615      4004      24308      3986
e1/2           0           0           0           0
e1/3           0           0           0           0
e1/4           0           0           0           0
e1/5           0           0           0           0
e1/6           0           0           0           0
e1/7           0           0           0           0
e1/8           0           0           1           1
e1/9           0           0           0           0
e1/10        0           0           0           0
e1/11        0           0           0           0
e1/12        0           0           0           0
e1/13        0           0           0           0
e1/14        0           0           0           0
e1/15        0           0           0           0
e1/16        0           0           0           0
    
```

Packets=p or Bytes=b, Delta=d or Rate=r, Clear=c, Next=n :Freeze=f/s Quit=q

Syntax: monitor interface traffic [ethernet slot/port]

The **monitor interface traffic** command uses page mode display to produce an updating statistics screen which is updated every poll interval and which can be modified using the commands shown at the bottom of the display. (Note that when you enter footer options they are not executed immediately but will be interpreted by the monitor engine during the next polling cycle.) Normally the display begins with the lowest numbered Ethernet port; the **ethernet slot/port** option starts the display instead with the specified port.

The footer commands and their effects are described in Table 20.

TABLE 20 Footer commands for **monitor interface traffic** display

p	Displays input/output packets instead of bytes and continues the execution of the original command.
b	Displays input/output bytes instead of packets and continues the execution of the original command.
d	Displays delta counters instead of rate counters and continues the execution of the original command.
r	Displays rate counters instead of delta counters and continues the execution of the original command.
c	Clears the current delta counters and continues the execution of the original command. To clear the aggregate counters, use the appropriate clear command.
n	Moves to the next group of interfaces and continues the execution of the original command.
f	Freezes the execution of the command; pauses retrieval and display of the statistics. While display is frozen, the only valid commands are s and q ; you can restart or quit the monitor, but any other command will be ignored.
s	Restarts the execution of the command; resumes retrieval and display of the statistics.

TABLE 20 Footer commands for **monitor interface traffic** display

F	Decreases the polling interval (monitor refresh interval) by one second and continues the execution of the original command with the new refresh interval. This option will decrease the refresh interval until it is equal to 2 seconds, the minimum supported refresh interval value. The default value is 2 seconds. This command is not displayed in the footer of the statistics screen.
S	Increases the polling interval (monitor refresh interval) by one second and continues the execution of the original command with the new refresh interval. This option will increase the refresh interval until it is equal to 30 seconds, the maximum supported refresh interval value. The default value is 2 seconds. This command is not displayed in the footer of the statistics screen.
q or escape or ^c	Quits the execution of the command and returns to the command prompt.

Displaying recent traffic statistics for an Ethernet port

To display information from the **show statistics** command for an Ethernet port, enter the following command at any CLI level.

```
Brocade# show statistics ethernet 9/1
PORT 9/1 Counters:
      InOctets      210753498112      OutOctets      210753550720
      InPkts        1646511726        OutPkts        1646512119
      InBroadcastPkts 0      OutBroadcastPkts 0
      InMulticastPkts 0      OutMulticastPkts 0
      InUnicastPkts  1646511726      OutUnicastPkts  1646512142
      InDiscards     0      OutDiscards     0
      InErrors       0      OutErrors       0
      InCollisions   0      OutCollisions   0
      OutLateCollisions 0
      Alignment      0      FCS             0
      InFlowCtrlPkts 0      OutFlowCtrlPkts 0
      GiantPkts      0      ShortPkts       0
      InBitsPerSec   3440829770      OutBitsPerSec   3440686411
      InPktsPerSec   3360185      OutPktsPerSec   3360085
      InUtilization  39.78%      OutUtilization  39.78%
```

Syntax: **show statistics ethernet slot/port**

The *slot/port* variable specifies the port that you want to display statistics for.

This field...	Displays...
InOctets	The total number of good octets and bad octets received.
OutOctets	The total number of good octets and bad octets transmitted.
InPkts	The total number of packets received. The count includes rejected and local packets that are not transmitted to the switching core for transmission.
OutPkts	The number of good packets received. The count includes unicast, multicast, and broadcasts packets.
InBroadcastPkts	The total number of good broadcast packets received.
OutBroadcastPkts	The total number of good broadcast packets transmitted.
InMulticastPkts	The total number of good multicast packets received.
OutMulticastPkts	The total number of good multicast packets transmitted.

This field...	Displays...
InUnicastPkts	The total number of good unicast packets received.
OutUnicastPkts	The total number of good unicast packets transmitted.
InDiscards	The total number of packets that were received and then dropped due to a lack of receive buffers.
OutDiscards	The total number of packets that were transmitted and then dropped due to a lack of transmit buffers.
InErrors	The total number of packets received that had Alignment errors or phy errors.
OutErrors	The total number of packets transmitted that had Alignment errors or phy errors.
InCollisions	The total number of packets received in which a Collision event was detected.
OutCollisions	The total number of packets transmitted in which a Collision event was detected.
OutLateCollisions	The total number of packets transmitted in which a Collision event was detected, but for which a receive error (Rx Error) event was not detected.
Alignment	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).
FCS	The Frame Checksum error.
InFlowCtrlPkts	The total number of ingress flow control packets. “N/A” indicates that the interface module does not support flow control statistics.
OutFlowCtrlPkts	The total number of egress flow control packets. “N/A” indicates that interface module does not support flow control statistics.
GiantPkts	The total number of packets for which all of the following was true: <ul style="list-style-type: none"> • The data length was longer than the maximum allowable frame size. • No Rx Error was detected. This counter is only for 10GbE interfaces.
ShortPkts	The total number of packets received for which all of the following was true: <ul style="list-style-type: none"> • The data length was less than 64 bytes. • No Rx Error was detected. • No Collision or late Collision was detected.
InBitsPerSec	The number of bits received per second.
OutBitsPerSec	The number of bits transmitted per second.
InPktsPerSec	The number of packets received per second.
OutPktsPerSec	The number of packets transmitted per second.
InUtilization	The percentage of the port’s bandwidth used by received traffic.
OutUtilization	The percentage of the port’s bandwidth used by transmitted traffic.

Configuring SNMP to revert ifType to legacy values

The ifType for all Ethernet interfaces (10/100/1G/10G) returns the value ethernetCsmacd(6) as mandated by RFC 2665. If you want ifType to return gigabitEthernet (117) or fastEther(62) for Ethernet interfaces, enter the following command.

2 Configuring snAgentConfigModuleType to return original values

```
Brocade(config)# snmp-server legacy iftype
```

Syntax: [no] snmp-server legacy iftype

When this command is configured, the values gigabitEthernet (117) or fastEther(62) are returned for ifType. If you issue a **no snmp-server legacy iftype**, ifType returns ethernetCsmacd(6) for Ethernet interfaces.

Configuring snAgentConfigModuleType to return original values

Enumeration values for snAgentConfigModuleType object in the SNMP MIB have been changed in Release 04.0.00 for the Brocade Netron XMR and Brocade MLX series to resolve enumeration conflicts with other hardware modules in the Unified IP MIB. For example, an SNMP get of the snAgentConfigModuleType of the 10x1GC module returned xmr20PortGigCopperSPModule(84). Beginning with Release 04.0.00, snAgentConfigModuleType returns fdryXmr20PortGigCopperSPModule(1084) for the 10x1GC module.

If you want snAgentConfigModuleType to return the enumeration values used before Release 04.0.00, configure the following command.

```
Brocade(config)# snmp-server legacy module-type
```

Syntax: [no] snmp-server legacy module-type

Refer to the *Unified IP MIB Reference* for details on snAgentConfigModuleType.

Preserving interface statistics in SNMP

By default, statistics for an interface is cleared from both the CLI and SNMP when the following commands are entered on the CLI:

- **clear statistics ethernet** *slot-number/port-number*
- **clear statistics** *slot-number/port-number*
- **clear rmon statistics**
- **clear statistics log** *slot-number/port-number*

If you want to preserve interface statistics in SNMP when these commands are entered, configure the following command at the Global level of the CLI.

```
Brocade(config)# snmp-server preserve-statistics
```

Syntax: [no] snmp-server preserve-statistics

For details on which interface statistics are preserved in SNMP, refer to the “Preserved interface statistics for SNMP” section of the “Supported Standard MIBs” chapter in the *Unified IP MIB Reference*.

NOTE

Statistics for an interface will be different between the CLI and SNMP if **snmp-server preserve-statistics** is configured and the clear commands listed above are executed.

Disabling CAM table entry aging

By default if no traffic hits a programmed flow-based content addressable memory (CAM) table entry, the CAM entry is removed from the system's CAM table. Depending on your network needs, however, you might have to disable the default behavior and force the system to retain CAM entries even when no traffic hits them. You can stop and start the CAM aging feature by using the **hw-aging** command in the global configuration mode.

Syntax: **hw-aging** *disable* | *enable*

The **disable** option prevents CAM entries from aging out. Even if no traffic hits a particular CAM entry, the entry remains in the CAM table.

The **enable** option returns the system to the default mode and unused CAM age out of the CAM table.

Data integrity protection

Data integrity protection provides a way to detect and report potential problems with the internal data path of the network processor. It also allows you to tune the detection and reporting of these types of problems. In addition, a show command is provided to display the status of the system.

Configuring Detection Parameters

Several parameters can be configured to support this data integrity protection: rolling window time frame and the event thresholds for ingress and egress buffer events. The configurations are applied system wide.

Rolling Window Time Frame

Data integrity protection implements a rolling window to calculate the most recent history of errors. The rolling window time frame is the period of time error events are recorded. Data integrity protection polls for events every 500 milliseconds and updates the current window.

```
Brocade(config)# system np rolling-window 10
```

Syntax: **[no] system np rolling-window** *window size*

The *window size* parameter sets the rolling window time frame. The allowable window time is 10 to 60 seconds. Setting to 0 seconds will disable error monitoring.

The **[no]** option returns the threshold to the default setting.

Event Threshold Configuration

The data integrity protection implements configurable thresholds for generating a syslog and trap. There is one threshold for ingress buffer events and one threshold for egress buffer events. Once crossed, a syslog and trap will be generated.

To prevent excessive log and traps there is a 10 minute period before another syslog or trap is generated. Setting a threshold to zero disables error detection for the monitor point on all network processors.

The default threshold values are different for ingress and egress. The ingress error count is based on the errors detected on each 32-bit word. The egress error count is based on the number of packets with one or more errors.

The **system np ingress-threshold** command configures the ingress buffer error reporting threshold.

```
Brocade(config)# system np ingress-threshold 20
```

Syntax: [no] **system np ingress-threshold** *threshold*

The threshold range is 0 to 120 events. Setting the threshold to 0 disables the monitor point for all network processors. The default setting is 20 events.

The **[no]** option returns the threshold to default.

The **system np egress-threshold** command configures the egress buffer error reporting threshold.

```
Brocade(config)# system np egress-threshold 20
```

Syntax: [no] **system np egress-threshold** *threshold*

The threshold range is 0 to 120 events. Setting the threshold to 0 disables the monitor point for all network processors. The default setting is 3 events.

The **[no]** option returns the threshold to default.

Showing Status

The **show np buffer-errors** command displays the count of error events for the rolling window.

```
Brocade# show np buffer-errors
```

Ports	Ingress		Egress	
	Current	Cumulative	Current	Cumulative
1/1- 1/24	15	37	0	0
1/25 - 1/48	0	0	0	0
2/1 - 2/2	0	0	0	0

Syntax: **show np buffer-errors**

Commands

The following commands support the features described in this chapter:

- [show statistics ethernet](#)
- [sysmon fe link auto-tune](#)
- [sysmon tm link auto-tune](#)

show statistics ethernet

Displays statistics for an ethernet port.

Syntax show statistics ethernet *slot/port*

Parameters *slot/port* Specifies the port that you want to display statistics for.

Command Modes User EXEC mode
Privileged EXEC mode
Global configuration mode

Usage Guidelines

Command Output The **show statistics ethernet** command displays the following information:

Output field	Description
InOctets	The total number of good octets and bad octets received.
OutOctets	The total number of good octets and bad octets transmitted.
InPkts	The total number of packets received. The count includes rejected and local packets that are not transmitted to the switching core for transmission.
OutPkts	The number of good packets received. The count includes unicast, multicast, and broadcasts packets.
InBroadcastPkts	The total number of good broadcast packets received.
OutBroadcastPkts	The total number of good broadcast packets transmitted.
InMulticastPkts	The total number of good multicast packets received.
OutMulticastPkts	The total number of good multicast packets transmitted.
InUnicastPkts	The total number of good unicast packets received.
OutUnicastPkts	The total number of good unicast packets transmitted.
InDiscards	The total number of packets that were received and then dropped due to a lack of receive buffers.
OutDiscards	The total number of packets that were transmitted and then dropped due to a lack of transmit buffers.
InErrors	The total number of packets received that had Alignment errors or phy errors.
OutErrors	The total number of packets transmitted that had Alignment errors or phy errors.
InCollisions	The total number of packets received in which a Collision event was detected.
OutCollisions	The total number of packets transmitted in which a Collision event was detected.
OutLateCollisions	The total number of packets transmitted in which a Collision event was detected, but for which a receive error (Rx Error) event was not detected.
Alignment	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).
FCS	The Frame Checksum error.

Output field	Description
InFlowCtrlPkts	The total number of ingress flow control packets. "N/A" indicates that the interface module does not support flow control statistics.
OutFlowCtrlPkts	The total number of egress flow control packets. "N/A" indicates that interface module does not support flow control statistics.
GiantPkts	The total number of packets for which all of the following was true: <ul style="list-style-type: none"> The data length was longer than the maximum allowable frame size. No Rx Error was detected. This counter is only for 10GbE interfaces.
ShortPkts	The total number of packets received for which all of the following was true: <ul style="list-style-type: none"> The data length was less than 64 bytes. No Rx Error was detected. No Collision or late Collision was detected.
InBitsPerSec	The number of bits received per second.
OutBitsPerSec	The number of bits transmitted per second.
InPktsPerSec	The number of packets received per second.
OutPktsPerSec	The number of packets transmitted per second.
InUtilization	The percentage of the port's bandwidth used by received traffic.
OutUtilization	The percentage of the port's bandwidth used by transmitted traffic.

Examples The following example displays traffic statistics.

```

Brocade# show statistics ethernet 9/1
PORT 9/1 Counters:
      InOctets      210753498112      OutOctets      210753550720
      InPkts        1646511726      OutPkts        1646512119
InBroadcastPkts      0      OutBroadcastPkts      0
InMulticastPkts      0      OutMulticastPkts      0
InUnicastPkts      1646511726      OutUnicastPkts      1646512142
InDiscards      0      OutDiscards      0
InErrors      0      OutErrors      0
InCollisions      0      OutCollisions      0
      Alignment      0      OutLateCollisions      0
InFlowCtrlPkts      0      FCS      0
GiantPkts      0      OutFlowCtrlPkts      0
InBitsPerSec      3440829770      ShortPkts      0
InPktsPerSec      3360185      OutBitsPerSec      3440686411
InUtilization      39.78%      OutPktsPerSec      3360085
OutUtilization      39.78%      OutUtilization      39.78%

```

History

Release	Command History
Multi-Service IronWare R05.5.00	This command was modified to display flow control packet information.

Related Commands

sysmon fe link auto-tune

Enables auto tuning on the fabric element (FE). The **no** form of this command disables auto tuning on the FE.

Syntax **sysmon fe link auto-tune**
no sysmon fe link auto-tune

Command Default This command is enabled by default.

Parameters None.

Command Modes Global configuration mode

Usage Guidelines

Examples The following example disables auto-tuning on the FE.

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

History

Release	Command History
<i>Multi-Service IronWare R05.6.00</i>	This command was introduced.

Related Commands

sysmon tm link auto-tune

Enables auto tuning on the traffic manager (TM). The **no** form of this command disables auto tuning on the TM.

Syntax **sysmon tm link auto-tune**
 no sysmon tm link auto-tune

Command Default This command is enabled by default.

Parameters None.

Command Modes Global configuration mode

Usage Guidelines

Examples The following example disables auto-tuning on the TM.

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

History

Release	Command History
<i>Multi-Service IronWare R05.6.00</i>	This command was introduced.

Related Commands

2 sysmon tm link auto-tune

Telemetry Solutions

Table 21 displays the individual Brocade devices and the Interface Parameters features they support.

TABLE 21 Supported Brocade interface parameter features

Features supported	Brocade Netron XMR Series	Brocade MLX Series	Brocade Netron CES 2000 Series BASE package	Brocade Netron CES 2000 Series ME_PREM package	Brocade Netron CES 2000 Series L3_PREM package	Brocade Netron CER 2000 Series Base package	Brocade Netron CER 2000 Series Advanced Services package
Telemetry Solutions	Yes	Yes	No	No	No	No	No

About telemetry solutions

Telemetry Solutions provides a VLAN matching capability for IPv4 and IPv6 ACLs. Telemetry Solutions also includes new types of 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.

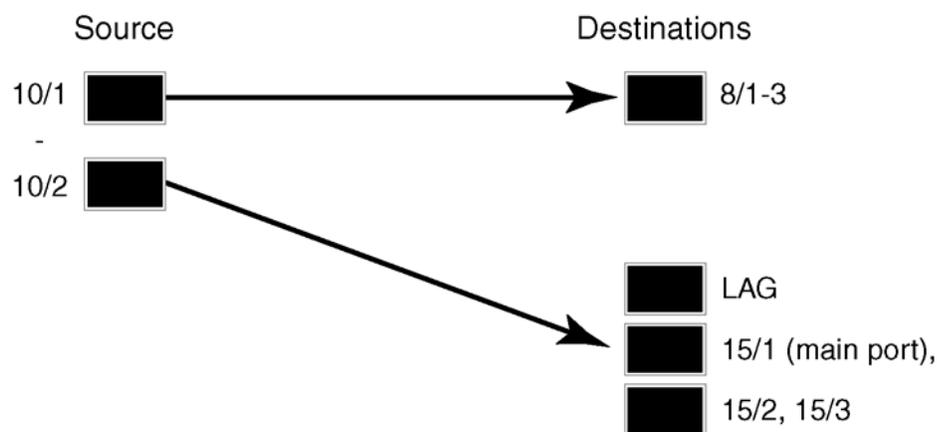
Configuration examples

NOTE

Telemetry can also be configured from SNMP. Refer to the Unified IP MIB Reference guide for more information.

Configuration example 1

FIGURE 4 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

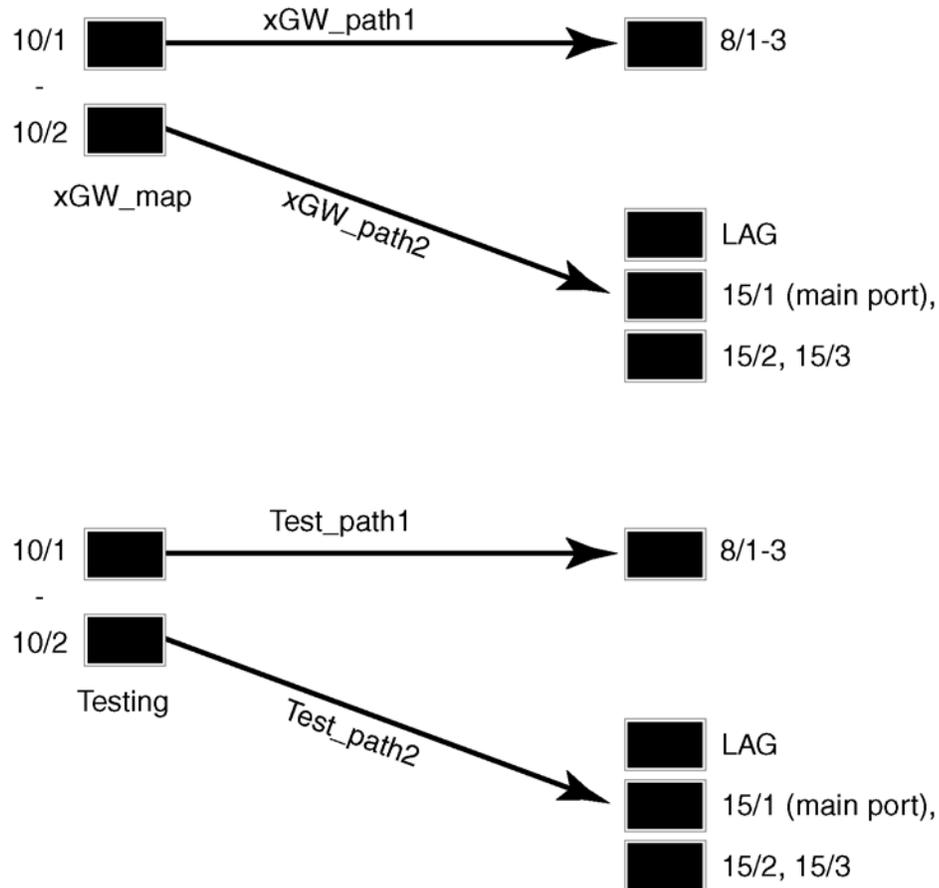
```

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 5 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

```

3 Configuration examples

```
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 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
```

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
```

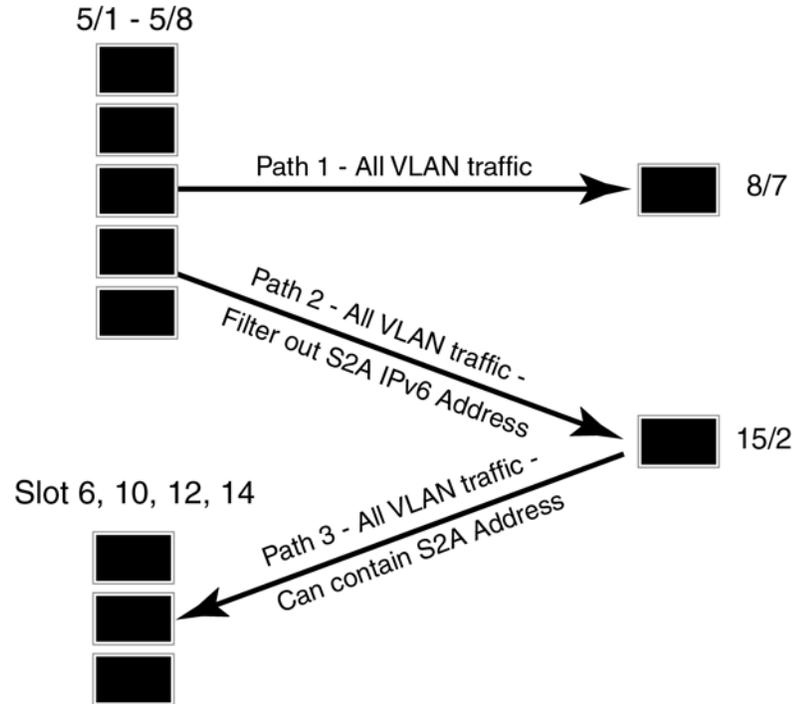
Modify destination ports (if necessary)

```
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 6 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
```

3 Configuration examples

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

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 new CLI 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 need to be created separately. In addition, both *ip policy route-map xGW_map* and *ipv6 policy route-map xGW_map* need to be configured on the interface.

3 Configuration examples

Remote Network Monitoring

Table 22 displays the individual Brocade devices and the Remote Network Monitoring features they support.

TABLE 22 Supported Brocade remote network monitoring features

Features supported	Brocade NetronXMR Series	Brocade MLX Series	Brocade Netron CES 2000 Series BASE package	Brocade Netron CES 2000 Series ME_PREM package	Brocade Netron CES 2000 Series L3_PREM package	Brocade Netron CER 2000 Series Base package	Brocade Netron CER 2000 Series Advanced Services package
Remote Network Monitoring	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Statistics (RMON Group 1)	Yes	Yes	Yes	Yes	Yes	Yes	Yes
History (RMON Group 2)	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Alarms (RMON Group 3)	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Events (RMON Group 9)	Yes	Yes	Yes	Yes	Yes	Yes	Yes

This chapter describes the remote monitoring features available on Brocade products:

- **Remote Monitoring (RMON) statistics** – All Brocade products support RMON statistics on the individual port level. Refer to “[RMON support](#)” on page 98.
- **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.

```
Brocade# 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.

```
Brocade# 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.

```
Brocade# 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 Netlon system provides limited MIB counters. Brocade 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.

```
Brocade(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 23 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). 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.

TABLE 23 Export configuration and statistics (Continued)

This line...	Displays...
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.

```
Brocade(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.

```
Brocade(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)

There are two elements to the Event Group – 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.

```
Brocade(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*

Continuous System Monitor

Table 24 displays the individual devices and the Continuous System Monitor features they support.

TABLE 24 Supported Continuous System Monitor features

Features supported	Brocade Netron XMR	Brocade MLX Series	Brocade Netron CES 2000 Series BASE package	Brocade Netron CES 2000 Series ME_PREM package	Brocade Netron CES 2000 Series L3_PREM package	Brocade Netron CER 2000 Series Base package	Brocade Netron CER 2000 Series Advanced Services package
Continuous System Monitoring (Sysmon)	Yes	Yes	Yes	Yes	Yes	Yes	Yes
System Resource Histogram	Yes	Yes	Yes	Yes	Yes	Yes	Yes
System Resource Histogram Enhancements for Memory Errors	Yes ^a	Yes	No	No	No	No	No
Automatic CRC Error Monitoring	Yes	Yes	No	No	No	No	No
Memory Error Monitoring	Yes	Yes	No	No	No	No	No

- a. Supported on the following interface cards: BR-MLX-100Gx2-X(100G), BR-MLX-10Gx24, BR-MLX-40Gx4-X, Gen-1 cards (NI-XMR-10Gx4, NI-MLX-10Gx4), Gen-1.1 cards (BR-MLX-1GCx24-X, BR-MLX-1GFx24-X, BR-MLX-10Gx4-X) and Gen-2 cards (NI-MLX-10Gx8-M, NI-MLX-10Gx8-D, BR-MLX-10Gx8-X).

This chapter contains the following sections:

- [Continuous system monitor overview](#)
- [Event monitoring](#)
- [Histogram information](#)
- [NP memory error monitoring](#)
- [Port CRC error monitoring test](#)
- [Commands](#)

Continuous system monitor overview

Continuous 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 Brocade Technical Support.

Internal logs are generated to give more information to Brocade 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](#)
- [Event types](#)
- [Displaying event information](#)

Event monitoring overview

The Sysmon monitors a number of event types periodically. Sysmon detects errors based on polling and interrupt. Polling is reading specific hardware registers. Interrupt is an instantaneous event detection by Sysmon. Sysmon continuously monitors management processor and interface processors via polling and interrupt methods. Once a threshold is reached, Sysmon logs the event in the internal Sysmon log and takes an action based on the event type. There are the following action types:

- Syslog: Generates a message in the Syslog
- Shutdown link: Disables the link between the TM and the FE

- SNMP trap: Generates an SNMP trap

By default, SYSMON is enabled to monitor and detect all 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

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 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
```

For detailed information on NP memory error monitoring, refer to [“NP memory error monitoring”](#) on page 114.

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.

```
Brocade# 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.

Brocade# 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:

Brocade# show sysmon config

EVENT	ACTION	POLL PERIOD (SEC)	THERESHOLD # (PER POLL in #POLL)	LOG BACK-OFF
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

Histogram information

This section discusses the following topics:

- [Histogram information overview](#)
- [Displaying CPU histogram information](#)
- [Displaying buffer histogram information](#)
- [Displaying memory histogram information](#)

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 across 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:

```
Brocade# 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
```

5 Histogram information

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:

```
Brocade# 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:

```
Brocade(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 25 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:

5 Histogram information

```
Brocade# 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
```

To display the buffer allocation stack for the top three tasks (in terms of buffer ownership), enter the following command:

```

Brocade(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 histogram data

To clear the buffer histogram data, enter the following command:

```
Brocade(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:

```
Brocade# 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
```

5 NP memory error monitoring

```
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 histogram data

To clear the memory histogram data, enter the following command:

```
Brocade(config)# clear memory histogram
```

Syntax: clear memory histogram

NP memory error monitoring

This sections discusses the following topics:

- [NP memory error monitoring overview](#)
- [NP memory error monitoring: basic configuration](#)

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 [Appendix F, “NP Memory Errors”](#).

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:

```
Brocade(config)# sysmon np memory-errors polling-period 10
```

To configure NP memory error monitoring to generate syslog messages, use the following command:

```
Brocade(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:

```
Brocade(config)# sysmon np memory-errors action none
```

The following example disables monitoring of memory errors on interface modules.

```
Brocade(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.

```
Brocade(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.

```
Brocade(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:

```
Brocade(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.

Port CRC error monitoring test

This section discusses the following topics:

- [Port CRC error monitoring overview](#)
- [Port CRC error monitoring: basic configuration](#)

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 counter) 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](#)"). 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

Example

To configure the port CRC error counter limit to 20, enter the following command:

```
Brocade(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:

```
Brocade(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:

```
Brocade(config)# sysmon port port-crc-test threshold 3 5
```

Syntax: `sysmon port port-crc-test threshold num-failures num-polls`

The *num-failures* variable specifies 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 threshold 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.

Example

To disable the port CRC error monitoring test, enter the following command:

```
Brocade(config)# no sysmon port port-crc-test
```

To enable the test again, enter the following command:

```
Brocade(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 26 lists the commands to transition between port CRC error monitoring test diagnostic action states.

TABLE 26 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	

Example

To disable the port CRC error monitoring test diagnostic action, enter the following command:

```
Brocade(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:

```
Brocade(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.

Example

To disable syslog for 1,000 events:

```
Brocade(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.

Commands

The following commands support the features described in this chapter:

- [clear bm histogram](#)
- [clear cpu histogram sequence](#)
- [clear memory histogram](#)
- [show bm histogram](#)
- [show cpu histogram](#)
- [show cpu histogram sequence](#)
- [show memory histogram](#)

- `show sysmon config`
- `sysmon np memory-errors`
- `sysmon port port-crc-test action`
- `sysmon port port-crc-test counter`
- `sysmon port port-crc-test log-backoff`
- `sysmon port port-crc-test polling-period`
- `sysmon port port-crc-test threshold`

5 clear bm histogram

clear bm histogram

Clears buffer histogram data.

Syntax clear bm histogram

Command default

Parameters None

Command Modes Privileged EXEC mode
Global configuration mode

Usage Guidelines

Examples The following example clears buffer histogram data.

```
Brocade(config)# clear bm histogram
```

History

Release	Command History
<i>Multi-Service IronWare R05.5.00</i>	This command was introduced

Related Commands show bm histogram

clear cpu histogram sequence

Clears CPU histogram sequential execution of task data.

Syntax `clear cpu histogram sequence`

Command default

Parameters None

Command Modes Privileged EXEC mode
Global configuration mode

Usage Guidelines

Examples The following example clears the CPU histogram sequential execution of task information.

```
Brocade(config)# clear cpu histogram sequence
```

History

Release	Command History
Multi-Service IronWare R05.5.00	This command was introduced

Related Commands `show cpu histogram sequence`

5 clear memory histogram

clear memory histogram

Clears memory histogram data

Syntax clear memory histogram

Command default

Parameters None

Command Modes Privileged EXEC mode
Global configuration mode

Usage Guidelines

Examples The following example clears memory histogram data.

```
Brocade(config)# clear memory histogram
```

History

Release

Command History

Multi-Service IronWare R05.5.00 This command was introduced

Related Commands show memory histogram

show bm histogram

Displays task buffer usage information.

Syntax `show bm histogram [priority threshold-value | trace]`

Parameters `priority threshold-value`
 Specifies the display of histogram information for a specific buffer priority level. 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).

`trace` Specifies the display of the buffer allocation stack of the top three tasks.

Command Modes
 User EXEC mode
 Privileged EXEC mode
 Global configuration mode

Usage Guidelines

Command Output The `show bm histogram` command displays the following information:

Output field	Description
DURATION	
SEQ IDX	
TIME	
THRESHOLD	
POOL-ID	
SIZE (KB)	
TOTAL	
FREE	
IN-USE	
APP-OWN	
Task Name	
App-Owns(Buffer)	

5 show bm histogram

Examples The following example displays buffer histogram information:

```
Brocade# 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
```

History

Release

Command History

Multi-Service IronWare R05.5.00 This command was introduced

Related Commands

clear bm histogram

show cpu histogram

Displays task CPU usage information.

Syntax `show cpu histogram { hold | wait | interrupt | timer }
[taskname name | above threshold-value | noclear]`

Parameters	hold	Specifies the display of task hold time information.
	wait	Specifies the display of task wait time information.
	interrupt	Specifies the display of task user-interrupt usage information.
	timer	Specifies the display of task sys-timer time usage information.
	taskname name	Specifies the display of histogram information for a specific task.
	above threshold-value	Specifies the display of histogram information for tasks whose maximum hold time is above the specified value.
	noclear	Specifies that histogram data should not be cleared after display. By default, information is cleared on read.

Command Modes	User EXEC mode
	Privileged EXEC mode
	Global configuration mode

Usage Guidelines

Command Output The `show cpu histogram` command displays the following information:

Output field	Description
No of bucket	
Bucket Granularity	
Last Cleared at	
Total No. of task	
Task Name	
Bkt Num	
Bkt Time (ms)	
No of Time	
HoldTime Total(s)	
HoldTime Max (ms)	
IntrptTime Total(s)	
IntrptTime Max (ms)	
TimerTime Total(s)	
TimerTime Max (ms)	

5 show cpu histogram

Output field Description

WaitTime Total(s)

WaitTime Max (ms)

Time

Examples The following example displays task hold time information:

```
Brocade# 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
```

History

Release Command History

Multi-Service IronWare R05.5.00 This command was introduced

Related Commands

show cpu histogram sequence

Displays sequential execution of task information.

Syntax `show cpu histogram sequence [taskname name | above threshold-value | trace]`

Parameters

- sequence** Specifies the display of sequential execution of task information.
- taskname *name*** Specifies the display of histogram information for a specific task.
- above *threshold-value*** Specifies the display of histogram information for tasks whose maximum hold time is above the specified value.
- trace** Specifies the display of high CPU condition task trace information.

Command Modes

- User EXEC mode
- Privileged EXEC mode
- Global configuration mode

Usage Guidelines

Command Output The `show cpu histogram sequence` command displays the following information:

Output field	Description
THRESHOLD	
DURATION	
Seq No	
Task Name	
Context	
Hold Time Max (ms)	
Start Time	
End Time	
Date	

5 show cpu histogram sequence

Examples The follow example displays sequential execution of task information:

```
Brocade# 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  
-----
```

History

Release

Command History

Multi-Service IronWare R05.5.00 This command was introduced

Related Commands

clear cpu histogram sequence

show memory histogram

Displays task memory usage information.

Syntax `show memory histogram [pool pool-id | below threshold-value | trace taskname]`

Parameters `pool pool-id` Specifies the display of memory histogram information for a specific memory pool. The valid range is 0-3, where “0” = OS, “1” = Shared, “2” = Global and “3” = User Private.

`below threshold-value`

Specifies the display of memory histogram information when available memory falls below the specified percentage (5, 10 or 20 percent).

`trace taskname` Specifies the display of high CPU condition task traces.

Command Modes

- User EXEC mode
- Privileged EXEC mode
- Global configuration mode

Usage Guidelines

Command Output The `show memory histogram` command displays the following information:

Output field	Description
DURATION	How long the histogram data is stored. 60 seconds means 60 records.
SEQ IDX	Number of records in the histogram database.
TIME	Time at which histogram is taken.
AVAIL MEM	Available memory of particular pool.
POOL	Name of the pool.
Total Memory (bytes)	Total memory size of the pool.
Used Memory (bytes)	Memory used from the pool.
Available Memory (bytes)	Available memory of particular pool.
Task Name	Name of task that is currently holding the memory.
Alloc-Number	Memory allocation count of each task.
Alloc Size (Bytes)	Total number of bytes allocated to each task.

5 show memory histogram

Examples The following example displays memory histogram information:

```
Brocade# 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
mrxp          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
-----
```

History

Release

Command History

Multi-Service IronWare R05.5.00 This command was introduced

Related Commands

clear memory histogram

show sysmon config

Displays the system monitoring configuration.

Syntax `show sysmon config`

Parameters None.

Command Modes User EXEC mode

Privileged EXEC mode

Global configuration mode

Usage Guidelines

Command Output The `show sysmon config` command displays the following information:

Output field	Description
EVENT	Error event that is monitored
ACTION	Action taken when the error event is detected.
POLL PERIOD (SEC)	Frequency of polling for the error event.
THRESHOLD	The threshold at which the ACTION is taken.
LOG BACK-OFF	Specifies the number of error events to skip before logging syslog messages again. Syslog messages are disabled for the specified number of error events.

Examples The following example displays the monitoring configuration.

```
Brocade# show sysmon config
```

EVENT	ACTION	POLL PERIOD (SEC)	THERESHOLD #(PER POLL in #POLL)	LOG BACK-OFF
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

History

Release

Multi-Service IronWare R05.6.00

Command History

This command was enhanced to display the NP memory error monitoring event configuration.

Related Commands

sysmon np memory-errors

Configures memory error monitoring and reporting on interface modules. The **no** form of this command disables memory error monitoring on interface modules.

Syntax	sysmon np memory-errors { poll-interval secs action {syslog-and-trap syslog trap none } } no sysmon np memory-errors { poll-interval secs action {syslog-and-trap syslog trap none } }	
Command Default	Monitoring of NP memory-errors is enabled by default.	
Parameters	poll-interval secs	Specifies the frequency of polling for NP memory errors. The range is from 1 through 65535. The default value is 60 seconds.
	action	Specifies the action taken when NP memory errors are detected. The default action is syslog-and-trap .
	syslog-and-trap	Generate a syslog message and a SNMP trap.
	syslog	Generate a syslog message.
	trap	Send a SNMP trap.
	none	No action; reporting of errors is disabled. In the no form of the command, specifying the action as none restores the default action (syslog-and-trap).
Command Modes	Privileged EXEC configuration mode.	
Usage Guidelines	<p>If action is configured as syslog followed by configuration as trap, the action will become syslog-and-trap.</p> <p>The poll-interval 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 either the interface or management modules, low memory, etc.</p> <p>The action parameter controls the generation of syslog messages or SNMP traps. These messages cannot be controlled by the no snmp-server enable traps command or the no logging enable command.</p> <p>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.</p>	
Examples	<p>The following example specifies polling for NP memory errors at 10 second intervals.</p> <pre>Brocade(config)# sysmon np memory-errors poll-interval 10</pre> <p>The following example disables reporting of NP memory errors.</p> <pre>Brocade(config)# sysmon np memory-errors action none</pre> <p>The following example disables monitoring of memory errors on interface modules.</p> <pre>Brocade(config)# no sysmon np memory-errors</pre> <p>The no form of the command specifying a <i>poll-interval</i> value restores the default polling interval. For example, the following command restores the polling interval to the default value (60 seconds).</p> <pre>Brocade(config)# no sysmon np memory-errors poll-interval 1000</pre>	

The **no** form of the command specifying the **action** as **syslog-and-trap**, **syslog**, or **trap** removes the specified action. The following command removes the **syslog** action.

```
Brocade(config)# no sysmon np memory-errors action syslog
```

The **no** form of the command specifying the **action** as **none** restores the default action (**syslog-and-trap**). For example:

```
Brocade(config)# no sysmon np memory-errors action none
```

History

Release

Command History

Multi-Service IronWare R05.6.00 This command was introduced

Related Commands

sysmon port port-crc-test

Enables the port CRC error monitoring test. The **no** form of this command disables the port CRC error monitoring test.

Syntax **sysmon port port-crc-test**
no sysmon port port-crc-test

Command Default The port CRC error monitoring test is enabled by default.

Parameters None.

Command Modes Global configuration mode

Usage Guidelines

Examples The following example disables the port CRC error monitoring test.

```
Brocade(config)# no sysmon port port-crc-test
```

History

Release	Command History
<i>Multi-Service IronWare R05.5.00</i>	This command was introduced

Related Commands

- sysmon port port-crc-test action**
- sysmon port port-crc-test counter**
- sysmon port port-crc-test log-backoff**
- sysmon port port-crc-test polling-period**
- sysmon port port-crc-test threshold**

sysmon port port-crc-test action

Configures the diagnostic action for the port CRC error monitoring test.

Syntax `sysmon port port-crc-test action {none | syslog | port-disable}`

Command Default The default action for the port CRC error monitoring test is **syslog**.

Parameters

none	No action.
port-disable	Disable port.
syslog	Generate a syslog message.

Command Modes Global configuration mode

Usage Guidelines

Examples The following example sets the diagnostic action to disable the port when the port CRC error limit crosses the configured threshold.

```
Brocade(config)# sysmon port port-crc-test action port-disable
```

History

Release	Command History
<i>Multi-Service IronWare R05.5.00</i>	This command was introduced

Related Commands `sysmon port port-crc-test`

sysmon port port-crc-test counter

Configures the port CRC error count limit for the configured polling period for the port CRC error monitoring test.

Syntax `sysmon port port-crc-test counter port-crc-counter less-than crc-count`

Command Default

Parameters *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.

Command Modes Global configuration mode.

Usage Guidelines

Examples The following example configures the port CRC error counter limit to 20.

```
Brocade(config)# sysmon port port-crc-test counter port-crc-counter less-than 20
```

History

Release	Command History
<i>Multi-Service IronWare R05.5.00</i>	This command was introduced

Related Commands `sysmon port port-crc-test`

sysmon port port-crc-test log-backoff

Disables syslog messages for a specified number of events, before they are logged again for the port CRC error monitoring test.

Syntax `sysmon port port-crc-test log-backoff num`

**Command
Default**

Parameters *num* Specifies the number of events to skip before logging syslog messages again. The range of values is 1 through 14,400.

**Command
Modes** Global configuration mode.

**Usage
Guidelines**

Examples The following example disables syslog for 1,000 events.

```
Brocade(config)# sysmon port port-crc-test log-backoff 1000
```

History

Release

Command History

Multi-Service IronWare R05.5.00 This command was introduced

**Related
Commands** `sysmon port port-crc-test`

5 sysmon port port-crc-test polling-period

sysmon port port-crc-test polling-period

Configures the polling period for the port CRC error monitoring test.

Syntax `sysmon port port-crc-test polling-period secs`

**Command
Default**

Parameters `secs` Specifies the polling period in seconds. The range of values is 0 through 65535. The default value is 60 seconds.

**Command
Modes** Global configuration mode.

**Usage
Guidelines**

Examples The following example configures the port CRC error monitoring test to run every 60 seconds.

```
Brocade(config)# sysmon port port-crc-test polling-period 60
```

History

Release

Command History

Multi-Service IronWare R05.5.00 This command was introduced

**Related
Commands** `sysmon port port-crc-test`

sysmon port port-crc-test threshold

Configures the threshold for diagnostic action for the port CRC error monitoring test.

Syntax `sysmon port port-crc-test threshold num-failures num-polls`

Command default The default threshold is 3 failed test runs out of 5 polls.

Parameters

<i>num-failures</i>	Specifies the number of failed test runs. The range of values is 1 through 31.
<i>num-polls</i>	Specifies the number of polls (tests). The range of values is 2 through 31.

Command Modes Global configuration mode.

Usage Guidelines

Examples The following example configures the threshold to trigger the configured diagnostic action if the test fails more than three times in five consecutive polls.

```
Brocade(config)# sysmon port port-crc-test threshold 3 5
```

History

Release	Command History
<i>Multi-Service IronWare R05.5.00</i>	This command was introduced

Related Commands `sysmon port port-crc-test`

5 sysmon port port-crc-test threshold

Operations, Administration, and Maintenance (OAM)

Table 27 displays the individual Brocade devices and the OAM features they support.

TABLE 27 Supported Brocade OAM features

Features supported	Brocade Netron XMR Series	Brocade MLX Series	Brocade Netron CES 2000 Series BASE package	Brocade Netron CES 2000 Series ME_PREM package	Brocade Netron CES 2000 Series L3_PREM package	Brocade Netron CER 2000 Series Base package	Brocade Netron CER 2000 Series Advanced Services package
IEEE 802.1ag Connectivity Fault Management (CFM)	Yes	Yes	No	Yes	No	No	Yes
IEEE 802.1ag Connectivity Fault Management (CFM) for C-VLANs and S-VLANs within an ESI	No	No	No	Yes	No	No	Yes
IEEE 802.1ag Connectivity Fault Management (CFM) for B-VLANs	No	No	No	Yes	No	No	Yes
Support for Sub-second IEEE 802.1ag Timers	Yes	Yes	No	No	No	Yes	Yes
IEEE 802.1ag on VPLS Endpoints	Yes	Yes	No	Yes	No	No	Yes
IEEE 802.1ag over VLL	Yes	Yes	No	No	No	No	No
MPLS OAM – LSP traceroute	Yes	Yes	Yes	Yes	Yes	Yes	Yes
RDI handling	Yes	Yes	Yes	Yes	Yes	Yes	Yes
CFM handling for ISID	Yes	Yes	No	No	Yes	No	Yes

TABLE 27 Supported Brocade OAM features (Continued)

Features supported	Brocade Netron XMR Series	Brocade MLX Series	Brocade Netron CES 2000 Series BASE package	Brocade Netron CES 2000 Series ME_PREM package	Brocade Netron CES 2000 Series L3_PREM package	Brocade Netron CER 2000 Series Base package	Brocade Netron CER 2000 Series Advanced Services package
Delay measurement for ISID	Yes	Yes	No	No	Yes	No	Yes
MPLS OAM – LSP traceroute	Yes	Yes	No	No	Yes	Yes	Yes
IEEE 802.3ah EFM-OAM	Yes	Yes	No	Yes	No	No	Yes
Ping	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ping within a VRF	Yes	Yes	No	No	No	No	No
Trace Route	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Trace Route within a VRF	Yes	Yes	No	No	No	No	No
IPv6 Traceroute over an MPLS network	Yes	Yes	No	Yes	Yes	No	Yes
Trace-I2 Protocol	Yes	Yes	Yes	Yes	Yes	Yes	Yes
LSP Ping and Traceroute	Yes	Yes	No	Yes	No	No	Yes
Port Status TLV	Yes	Yes	Yes	Yes	Yes	Yes	Yes
RDI handling	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Link MA	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Frame Loss Measurement	No	No	Yes	Yes	Yes	Yes	Yes
One-way Delay Measurement	No	No	Yes	Yes	Yes	Yes	Yes
Synthetic Loss Measurement	No	No	Yes	Yes	Yes	Yes	Yes

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 (CFM)

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.

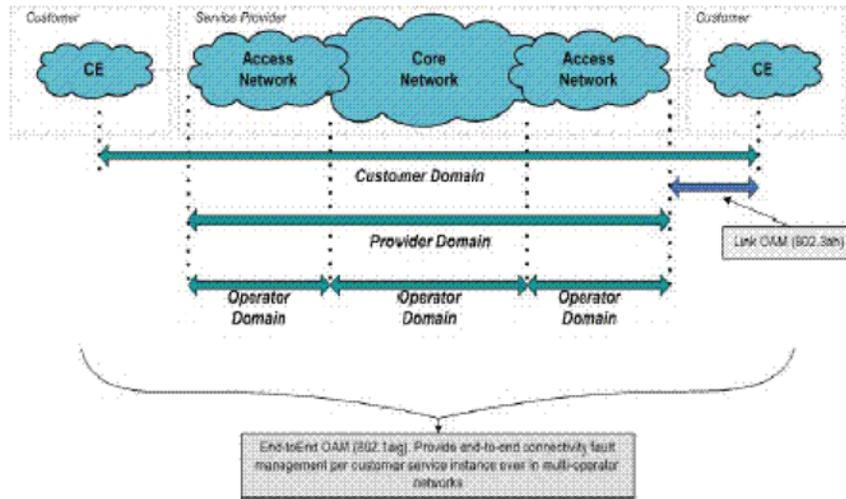
There are multiple organizations involved in a Metro Ethernet Service: Customers, Service Providers and Operators.

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 7 OAM Ethernet tools

6 IEEE 802.1ag Connectivity Fault Management (CFM)



IEEE 802.1ag provides hierarchical network management

Maintenance Domain (MD)

A Maintenance domain is part of a network controlled by a single operator. In [Figure 7](#), we have customer domain, provider domain and operator domain.

Maintenance Domain level (MD level)

The MD levels are carried on all CFM frames to identify different domains. For example, in [Figure 7](#), 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 (MA)

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

Maintenance End Point (MEP)

MEP is located on the edge of an 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 (MIP)

MIP is located within a 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 for 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; X can be 3.3 milliseconds (ms), 10ms, or 100ms, 1 second, 1 minute, or 10 minutes. 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 Brocade 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.

```
Brocade(config)#cfm-enable
Brocade(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.

6 Setting Maintenance Domain parameters

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

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

The **domain-name** *name* parameter specifies the domain name. The *name* attribute is case-sensitive.

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

The **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.

```
Brocade(config-cfm-md-VPLS-SP)# ma-name ma_1 vlan-id 30 priority 4  
Brocade(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 **tag2** 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.

```
Brocade(config)# tag-value tag1 8100
Brocade(config)# tag-value tag2 9100
Brocade(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.

```
Brocade(config-if-e1000-1/1)# tag-type tag2 ethernet 1/1
Brocade(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 (MA)

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.

6 Setting Maintenance Domain parameters

```
Brocade(config-cfm)#domain name VPLS-SP level 4
Brocade(config-cfm-md-VPLS-SP)#ma-name ma_1 vlan-id 30 priority 3
Brocade(config-cfm-md-VPLS-SP-ma-ma_1)#ccm-interval 10-second
Brocade(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 milliseconds** parameter sets the time interval between two successive CCM packets to 3.3 milliseconds.

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

The **100 milliseconds** 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.

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

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

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

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

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

Use the **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

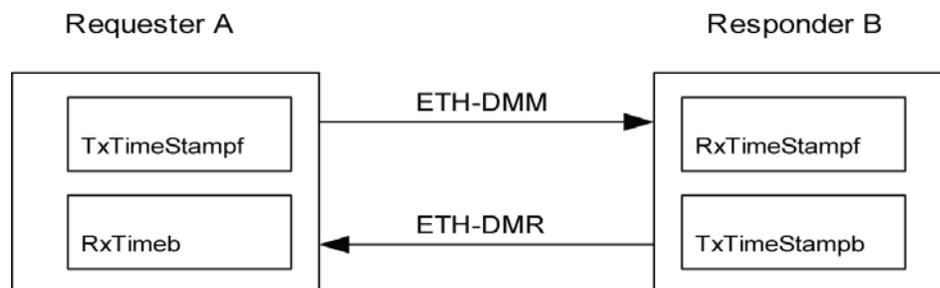
NOTE

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

About Y.1731

Figure 8 shows an ETH-DM requester and responder.

FIGURE 8 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)

This release 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

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.

```
Brocade# 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 statistic delay_measurement domain** command to display delay measurement statistics. If the command is issued gain, the output is replaced with the new values.

```

Brocade#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 :

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

Responder-B:

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

2. Issue the `cfm delay_measurement` command.

```
Brocade# 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.

```
Brocade#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.

```
Brocade# cfm linktrace domain VPLS-SP ma ma_1 src-mep 21 target-mep 1 timeout 10 t
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
          Forwarded              5/4              EgrOK
      2    0000.00fb.5378
          Not Forwarded          7/2              IgrOK              RLY_HIT
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*] [**tll** *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 **tll** *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.

```
Brocade# 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
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 src-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.

```
Brocade#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 28 Show CFM output descriptions

This field...	Displays...
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: <ul style="list-style-type: none"> 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.

TABLE 28 Show CFM output descriptions

This field...	Displays...
MIP	Displays the associated MIP
VLAN	Displays the associated VLAN.

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

```
Brocade#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 29 Show cfm brief output description

This field...	Displays...
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: <ul style="list-style-type: none"> • 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.

```

Brocade#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
    
```

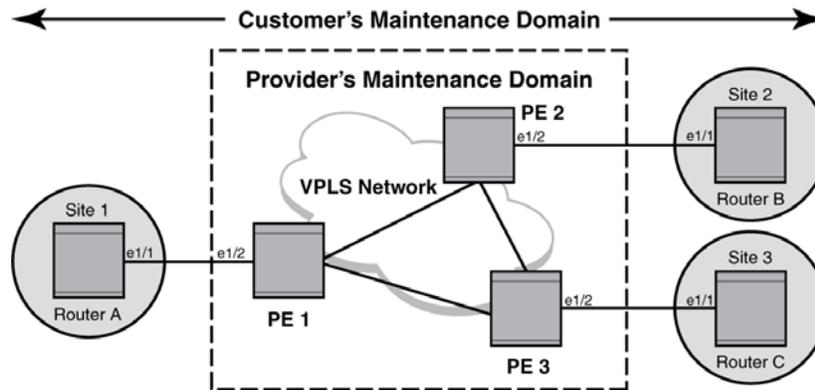
Syntax: show cfm connectivity

TABLE 30 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.
Index	The Domain Index.
Level	The level is the domain level in the range 0-7. The levels can be: <ul style="list-style-type: none"> • 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 9 Sample configuration



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](#)” on page 160.

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.

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

2. To enable CFM, enter the following command.

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

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

```
RouterA(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**.

```
RouterA(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**.

```
RouterA(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.

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

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

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

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

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

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

2. To enable CFM, enter the following command.

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

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

```
RouterB(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**.

```
RouterB(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**.

```
RouterB(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.

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

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

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

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

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

Configuring Router C

CFM configuration steps for Router C are listed below.

6 CFM monitoring and show commands

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

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

2. To enable CFM, enter the following command.

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

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

```
RouterC(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**.

```
RouterC(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**.

```
RouterC(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.

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

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

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

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

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

Provider Bridge Brocade1

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

```
Brocade1(config)# vlan30
Brocade1(config-vlan-30)# tagged ethe 1/1
```

2. Create the ESI Brocade1**vlan300** as an encapsulated SVLAN with the ESI client **Site1vlan30** by entering the following commands.

```
Brocade1(config)#esi Brocade1vlan300 encapsulation svlan
Brocade1(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.

```
Brocade1(config)# vlan300
Brocade1(config-vlan-300)# tagged ethe 1/1 ethe 1/3
```

4. To enable CFM, enter the following command.

```
Brocade1(config)#cfm-enable
```

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

```
Brocade1(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**.

```
Brocade1(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**.

```
Brocade1(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.

```
Brocade1(config-cfm-md-CUST_1-ma-ma_5)#mep 4 up port ethe 1/2
```

9. To configure the hostname as **Brocade1**, enter a command such as the following.

```
Brocade1(config)#hostname Brocade1
```

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

```
Brocade1(config)#interface ethernet 1/1
Brocade1(config-if-e10000-1/1)#port-type provider-network
Brocade1(config-if-e10000-1/1)enable
Brocade1(config-if-e10000-1/1)end
```

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

```
Brocade1(config)#interface ethernet 1/2
Brocade1(config-if-e10000-1/2)#port-type custommer-edge
Brocade1(config-if-e10000-1/2)enable
Brocade1(config-if-e10000-1/2)end
```

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

```
Brocade1(config)#interface ethernet 1/3
Brocade1(config-if-e10000-1/3)#port-type provider-network
Brocade1(config-if-e10000-1/3)enable
Brocade1(config-if-e10000-1/3)end
```

Provider Bridge Brocade2

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

```
Brocade2(config)# vlan300
Brocade2(config-vlan-300)# tagged ethe 1/1 ethe 1/3
```

2. To enable CFM, enter the following command.

```
Brocade2(config)#cfm-enable
```

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

```
Brocade2(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**.

```
Brocade2(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**.

6 CFM monitoring and show commands

```
Brocade2(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.

```
Brocade2(config-cfm-md-CUST_1-ma-ma_5)#mep 5 up port ethe 1/2
```

7. To configure the hostname as **Brocade1**, enter a command such as the following.

```
Brocade2(config)#hostname Brocade1
```

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

```
Brocade2(config)#interface ethernet 1/1
Brocade2(config-if-e10000-1/1)#port-type provider-network
Brocade2(config-if-e10000-1/1)enable
Brocade2(config-if-e10000-1/1)end
```

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

```
Brocade2(config)#interface ethernet 1/2
Brocade2(config-if-e10000-1/2)#port-type customer-edge
Brocade2(config-if-e10000-1/2)enable
Brocade2(config-if-e10000-1/2)end
```

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

```
Brocade2(config)#interface ethernet 1/3
Brocade2(config-if-e10000-1/3)#port-type provider-network
Brocade2(config-if-e10000-1/3)enable
Brocade2(config-if-e10000-1/3)end
```

Provider Bridge Brocade3

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

```
Brocade3(config)# vlan30
Brocade3(config-vlan-30)# tagged ethe 1/2
```

2. Create the ESI **Brocade3vlan300** as an encapsulated SVLAN with the ESI client **Site3vlan30** by entering the following commands.

```
Brocade3(config)#esi Brocade3vlan300 encapsulation svlan
Brocade3(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.

```
Brocade3(config)# vlan300
Brocade3(config-vlan-300)# tagged ethe 1/1 ethe 1/3
```

4. To enable CFM, enter the following command.

```
Brocade3(config)#cfm-enable
```

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

```
Brocade3(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**.

```
Brocade3(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**.

```
Brocade3(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.

```
Brocade3(config-cfm-md-CUST_1-ma-ma_5)#mep 6 up port ethe 1/2
```

9. To configure the hostname as **Brocade3**, enter a command such as the following.

```
Brocade3(config)#hostname Brocade3
Configure interface ethernet 1/1 as the provider network by entering the
following commands:
Brocade3(config)#interface ethernet 1/1
Brocade3(config-if-e10000-1/1)#port-type provider-network
Brocade3(config-if-e10000-1/1)enable
Brocade3(config-if-e10000-1/1)end
```

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

```
Brocade3(config)#interface ethernet 1/2
Brocade3(config-if-e10000-1/2)#port-type custommer-edge
Brocade3(config-if-e10000-1/2)enable
Brocade3(config-if-e10000-1/2)end
```

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

```
Brocade3(config)#interface ethernet 1/3
Brocade3(config-if-e10000-1/3)#port-type provider-network
Brocade3(config-if-e10000-1/3)enable
Brocade3(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.

```
RouterA#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,
```

6 CFM monitoring and show commands

```
RouterB#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,
```

```
RouterC#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 10 Sample configuration

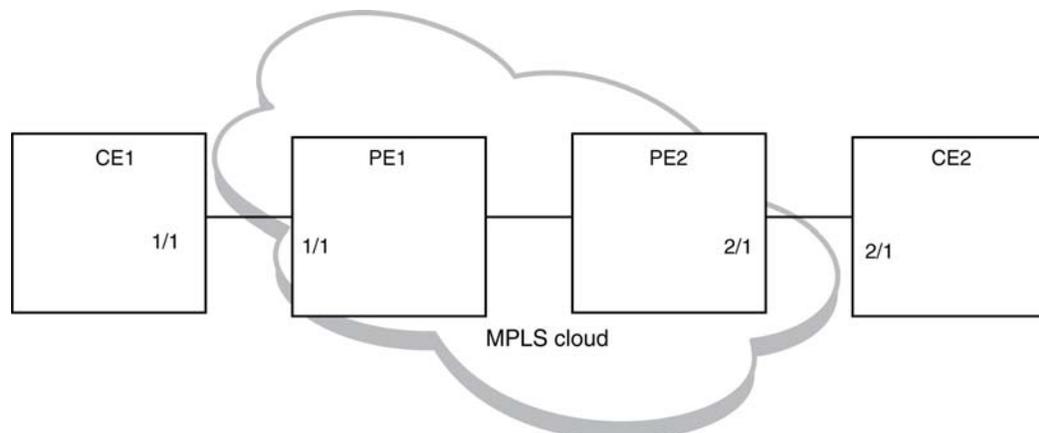


Figure 10 shows a deployment scenario where CE1 and CE2 are customer devices and PE1 and PE2 are provider routers. Port 1/1 on PE1 and port 2/1 on PE2 are VLL-end points. Port 1/1 on PE1 is connected to port 1/1 on CE1 and port 2/1 on PE2 is connected to port 2/1 on CE2.

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. The configuration for this is as follows.

Configuring CE1

1. To enable CFM, enter the following command.

```
CE1(config)#cfm-enable
```

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

```
CE1(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.

```
CE1(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.

```
CE1(config-cfm-md-CUST_1-ma-ma_5)#ccm-interval 10-second
```

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

```
CE1(config-cfm-md-CUST_1-ma-ma_5)#mep 1 down vlan 30 port ethe 1/1
```

6. 6. Configure a remote-mep.

```
CE1(config-cfm-md-CUST_1-ma-ma_5)#remote-mep 2 to 2
```

Configuring CE2

1. To enable CFM, enter the following command.

```
CE2(config)#cfm-enable
```

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

```
CE2(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.

```
CE2(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.

```
CE2(config-cfm-md-CUST_1-ma-ma_5)#ccm-interval 10-second
```

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

```
CE1(config-cfm-md-CUST_1-ma-ma_5)#mep 2 down vlan 30 port ethe 2/1
```

6. Configure a remote-mep.

```
CE1(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.

```
PE1(config)#router mpls
PE1(config-mpls)vll pe1-to-pe2 30
```

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

```
PE1(config-mpls-vll)vll-peer 10.1.1.2
```

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

```
PE1(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.

```
PE1(config-mpls-vll)vlan 30
PE1(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.

```
PE1(config)#cfm-enable
```

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

```
PE1(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.

```
PE1(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.

```
PE1(config-cfm-md-CUST_1-ma-ma_5)#ccm-interval 10-second
```

In the above 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.

```
PE1(config-cfm-md-CUST_1-ma-ma_5)# mip-creation explicit
```

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

```
PE1(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.

```
PE2(config)#router mpls
PE2(config-mpls)vll pe2-to-pe1 30
```

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

```
PE2(config-mpls-vll)vpls-peer 10.1.1.1
```

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

```
PE2(config-mpls-vll)untagged ethe 2/1
Tagged ports are configured under a VLAN ID.
```

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

```
PE2(config-mpls-vll)vlan 30
PE2(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.

- To enable CFM, enter the following command.

```
PE2(config)#cfm-enable
```

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

```
PE2(config-cfm)#domain-name CUST_1 level 7
```

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

```
PE2(config-cfm-md-CUST_1)#ma-name ma_5 vll-id 30 priority 3
```

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

```
PE2(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.

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

```
PE2(config-cfm-md-CUST_1-ma-ma_5)# mip-creation explicit
```

- To change back to default use the following command.

```
PE2(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.

```
CE1#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,
CE1#show cfm connectivity domain CUST_1 ma ma_5 rmep-id 2
Domain: CUST_1 Level: 7
```

6 CFM monitoring and show commands

```
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 31 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: <ul style="list-style-type: none"> • 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.

```
CE1#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
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-name* **src-mep** *mep-id* **target-mip** *HH:HH:HH:HH:HH:HH* | **target-mep** *mep-id* [**timeout** *timeout*] [**tll** *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** *mepid* 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 **tll** *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 [Figure 10](#), 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.

```
CE1(config)#cfm-enable
```

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

```
CE1(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.

```
CE1(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.

```
CE1(config-cfm-md-CUST_2-ma-ma_6)#ccm-interval 10-second
```

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

```
CE1 (config-cfm-md-CUST_2-ma-ma_6)#mep 3 down vlan 30 port ethe 1/1
```

6. Configure a remote-mep.

```
CE1(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.

```
PE1(config)#cfm-enable
```

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

```
PE1(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.

```
PE1(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.

```
PE1(config-cfm-md-CUST_2-ma-ma_6)#ccm-interval 10-second
```

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

```
CE-1 (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.

```
PE1(config)#cfm-enable
```

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

```
PE1(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.

```
PE1(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.

```
PE1(config-cfm-md-PROVIDER_1-ma-ma_8)#ccm-interval 10-second
CE1 (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.

```
PE2(config)#cfm-enable
```

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

```
PE2(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.

```
PE1(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.

```
PE1(config-cfm-md-PROVIDER_1-ma-ma_8)#ccm-interval 10-second
CE1 (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

- To enable CFM, enter the following command.

```
PE2(config)#cfm-enable
```
- Create a maintenance domain with a specified name PROVIDER_1 and level 4.

```
PE2config-cfm)#domain-name PROVIDER_1 level 4
```
- Create a maintenance association within a specified domain of vll-id 30 with a priority 3.

```
PE1(config-cfm-md-PROVIDER_1)#ma-name ma_8 vll-id 30 priority 3
```
- Set the time interval between successive Continuity Check Messages. The default is 10-seconds.

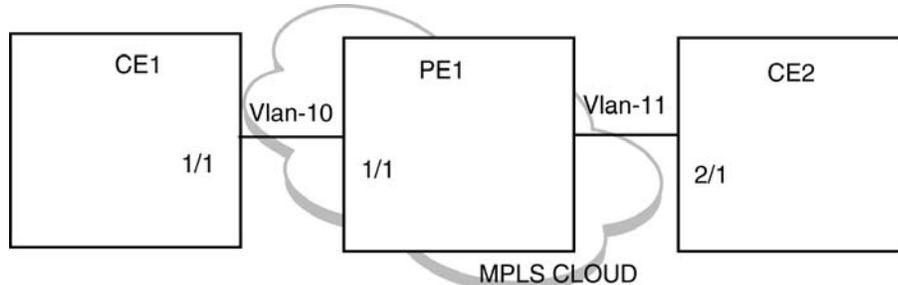
```
PE1(config-cfm-md-PROVIDER_1-ma-ma_8)#ccm-interval 10-second
```
- Configure MEP 4 down on port 1/1 and vlan 30

```
CE-1 (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

FIGURE 11 IEEE 802.1ag over 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 [Figure 11](#). 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 would not be allowed for 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.

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

```
PE1(config)#router mpls
PE1(config-mpls)vll-local test1
```
- To specify an un-tagged endpoint for a VLL instance, enter the following commands.

```
PE1(config-mpls-vll-test1)untagged ethe 1/1
```

Tagged ports are configured under a VLAN ID.

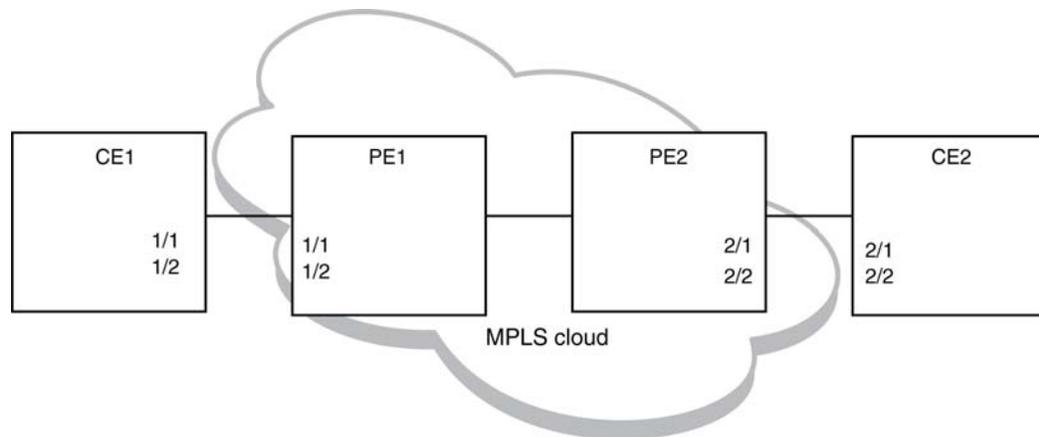
- To specify a tagged endpoint for a VLL instance, enter the following commands.

```
PE1(config-mpls-vll-test1)vlan 30
PE1(config-mpls-vll-vlan)tagged ethe 1/1
```

As in the previous case, to monitor the connectivity between CE1 and CE2, you can use “**show cfm connectivity**” commands as mentioned in the previous scenario. Also we could use either loopback or linktrace on CE1 or CE2.

LAG-support for IEEE 802.1ag-over-vll

FIGURE 12 - IEEE 802.1ag over VLL scenerio



As shown in [Figure 12](#), 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

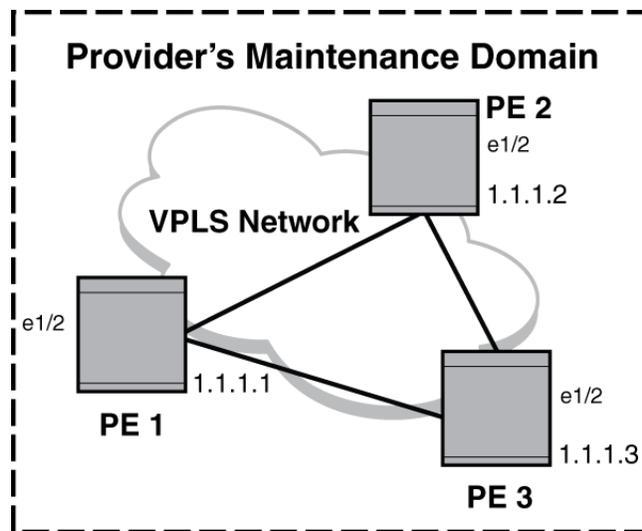
CFM provides capabilities to detect, verify, and isolate connectivity failures.

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 13](#), 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 13 VPLS cloud with CFM enabled



Configuring PE 1

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

```
PE1(config)#cfm-enable
```

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

```
PE1(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**.

```
PE1(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.

```
PE1(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.

```
PE1(config)#router mpls
PE1(config-mpls)#vpls 1 1
```

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

```
PE1(config-mpls-vpls-1)#vpls-peer 10.1.1.2
PE1(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.

```
PE1(config-mpls-vpls-1)#vlan 30
PE1(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.

```
PE2(config)#cfm-enable
```

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

```
PE2(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**.

```
PE2(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.

```
PE2(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 to PE 3.

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

```
PE2(config)#router mpls
PE2(config-mpls)vpls 1 1
```

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

```
PE2(config-mpls-vpls-1)vpls-peer 10.1.1.1
PE2(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.

```
PE2(config-mpls-vpls-1)vlan 30
PE2(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.

```
PE3(config)#cfm-enable
```

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

```
PE3(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**.

```
PE3(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.

```
PE3(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.

```
PE3(config)router mpls
PE3(cconfig-mpls)vpls 1 1
```

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

```
PE3(config-mpls-vpls-1)vpls-peer 10.1.1.1
PE3(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.

```
PE3(config-mpls-vpls-1)vlan 30
PE3(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.

```
Brocade#sh 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 32 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: <ul style="list-style-type: none"> 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.

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

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.

```
Brocade#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
```

```
Brocade#show cfm connectivity domain VPLS-SP ma ma_1 rmeip-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
      =====   =====   =====   =====   ===   =====   =====   =====   =====   =====
  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 NAME] parameter specifies the maintenance association name. By default, all defined domains are shown.

TABLE 33 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: <ul style="list-style-type: none"> 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.

TABLE 33 Output for show CFM connectivity command

This field...	Displays...
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.

```

Brocade#cfm linktrace domain VPLS-SP ma ma_1 src-mep 1 target-mep 4
Linktrace to 0000.00e2.d80a on Domain VPLS-SP, level 4: timeout 10ms, 8 hops
-----
Hops          MAC              Ingress          Ingress Action   Relay Action
              Forwarded        Egress           Egress Action    Nexthop
-----
Type Control-c to abort
  1    0000.00e2.d80a          10.1.1.1      IgrOK          RLY_HIT
              Not Forwarded
Destination 0000.00e2.d80a 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 **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.

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

```
Brocade#cfm loopback domain VPLS-SP ma ma_1 src-mep 1 target-mep 4
DOT1AG: Sending 10 Loopback to 0000.00e2.d80a, timeout 10000 msec
Type Control-c to abort
Reply from 0000.00e2.d80a: time=3ms
Reply from 0000.00e2.d80a: time<1ms
Reply from 0000.00e2.d80a: time<1ms
Reply from 0000.00e2.d80a: time<1ms
Reply from 0000.00e2.d80a: time=38ms
Reply from 0000.00e2.d80a: 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 loopback domain *name* ma *ma-name* src-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

```
Brocade#
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 208 for additional information.

Support for IEEE 802.1ag CFM for Provider Bridges (PB) and Provider Backbone Bridges (PBB)

The device supports 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 Brocade NetIron CES 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 the Brocade NetIron CES 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**Example CER1**

```

Brocade(config)#cfm-enable
Brocade(config-cfm)#domain-name customer level 7
Brocade(config-cfm-md-customer)#ma-name admin vlan 100 priority 7
Brocade(config-cfm-md-customer-ma-admin)#ccm-interval 100-ms
Brocade(config-cfm-md-customer-ma-admin)#mep 1 down port ethe 1/13

```

Example CER2

```

Brocade(config)#cfm-enable
Brocade(config-cfm)#domain-name customer level 7
Brocade(config-cfm-md-customer)#ma-name admin vlan 100 priority 7
Brocade(config-cfm-md-customer-ma-admin)#ccm-interval 100-ms

```

Example CER3

```

Brocade(config)#cfm-enable
Brocade(config-cfm)#domain-name customer level 7
Brocade(config-cfm-md-customer)#ma-name admin vlan 100 priority 7
Brocade(config-cfm-md-customer-ma-admin)#ccm-interval 100-ms
Brocade(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 OAM-EFM can be summarized under the following categories:

Discovery: Discovery is the mechanism to detect the presence of an OAM sublayer on the remote device. During the discovery process, information about OAM entities, capabilities and configuration 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 via 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

This is the first phase of the EFM-OAM. At this phase, EFM-OAM identifies network devices along with their OAM capabilities. The Discovery process relies on the Information OAMPDUs (discussed below). 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 (e.g. loopback capability).
- **OAM mode:** This is conveyed to the remote OAM entity. The mode can be either active or passive, and can also be used to determine device's functionality.
- **OAMPDU configuration:** This includes maximum OAMPDU size to delivery. In combination with the limited rate of ten frames/sec 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 1s - 10sec. The default value is 1sec.
- The Hold timer should assume the peer is dead if no packet is received for a period of 1s to 10s. 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 unsatisfactory and cannot continue.

Process overview

The discovery process allows a 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, etc.). 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 to start over again.

DTEs may either be in 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/Response PDUs.
- May send Loopback Control PDUs.

Exceptions:

- Does not respond to Variable Request PDUs from DTEs in Passive 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 (discussed in the OAMPDU section below) 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 will call 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.

Since the above conditions are severe, when they are set in the flag, the OAMPDU is not subject to normal rate limiting policy.

Remote loopback

An OAM entity can put its remote entity into loopback mode using a loopback control OAMPDU. This helps users 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. This allows to estimate if a network segment can satisfy an SLA.

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 enable command resets all link-oam parameters to default values.

By default, EFM-OAM is disabled.

To enable EFM-OAM, enter a command such as the following:

```
Brocade(config)link-oam
Brocade(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.

```
Brocade(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.

```
Brocade(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.

```
Brocade(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.

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.

```
Brocade(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 commands below display the three events that the protocol supports.

```
Brocade(config-link-oam)#ethernet 2/1 remote-failure dying-gasp action
block-interface
Brocade(config-link-oam)#ethernet 2/1 remote-failure critical-event action
block-interface
Brocade(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 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.

```
Brocade# link-oam remote-loop-back ethernet 1/1 start
```

Syntax: [no]link-oam remote loopback ethernet slot/port start/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:

```
Brocade#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:

```

Brocade#show link-oam info detail
OAM information for Ethernet port: 1/1
+link-oam mode:      active
+link status:        up
+oam status:         up
Local information
  multiplexer action: forward
  parse action:       forward
  stable:              satisfied
  state:               up
  loopback state:     disabled
  dying-gasp:         false
  critical-event:     false
  link-fault:         false
Remote information
  multiplexer action: forward
  parse action:       forward
  stable:              satisfied
  loopback support:   disabled
  dying-gasp:         false
  critical-event:     false
  link-fault:         false
OAM information for Ethernet port: 1/2
+link-oam mode:      passive
+link status:        up
+oam status:         up
Local information
  multiplexer action: forward
  parse action:       forward
  stable:              satisfied
  state:               up
  loopback state:     disabled
  dying-gasp:         false
  critical-event:     false
  link-fault:         false
Remote information
  multiplexer action: forward
  parse action:       forward
  stable:              satisfied
  loopback support:   disabled
  dying-gasp:         false
  critical-event:     false
  link-fault:         false

```

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:

```
Brocade#show link-oam info detail ethernet 1/1
OAM information for Ethernet port: 1/1
+link-oam mode:      active
+link status:       up
+oam status:        up
Local information
  multiplexer action: forward
  parse action:     forward
  stable:           satisfied
  state:            up
  loopback state:   disabled
  dying-gasp:       false
  critical-event:   false
  link-fault:       false
Remote information
  multiplexer action: forward
  parse action:     forward
  stable:           satisfied
  loopback support: disabled
  dying-gasp:       false
  critical-event:   false
  link-fault:       false
```

Syntax: show link-oam info detail [all | ethernet slot/port]

Displaying OAM statistics

To show OAM statistics, enter a command such as the following:

```
Brocade#show link-oam statistics
Ethernet Tx Pdus      Rx Pdus
1/1      507          362
1/2      358          359
1/3      480          355
1/4      0             0
1/5      0             0
1/6      0             0
1/7      0             0
1/8      0             0
1/9      0             0
1/10     0             0
```

Syntax: show link-oam statistics

Displaying detailed OAM statistics

To show detailed OAM statistics, enter a command such as the following:

```
Brocade#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
  organization 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
    organization 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
    organization 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:

```
Brocade#show link-oam statistics detail ports ethernet 1/1
OAM statistics for Ethernet port: 1/1
  Tx statistics
    information OAMPDUs:                827
    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:                682
    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
```

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.

```
Brocade# ping 10.22.2.33
```

Syntax: **ping** *ip addr | hostname | vrf instance-name* [**source** *ip addr*] [**count** *num*] [**timeout** *msec*] [**ttl** *num*] [**size** *byte*] [**quiet**] [**numeric**] [**no-fragment**] [**verify**] [**data** *1-to-4 byte hex*] [**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.

```
Brocade# ping ipv6 2001:db8:847f:a385:34dd::45
```

Syntax: `ping ipv6 ipv6-address | host-name | 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 byte hex] [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 either **eth 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 **tll** *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:

I 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.

Example

```
Brocade> 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.

Example

```
Brocade# traceroute vrf blue 10.10.10.10
```

Syntax: `traceroute vrf vrf-name ip-address`

The `vrf vrf-name` parameter is the name of the VRF for you want are running the traceroute.

The `vrf 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.

Example

To trace the path from the device to a host with an IPv6 address of 2001:db8:349e:a384::34, enter the following command.

```
Brocade> 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-I2 protocol

Trace-I2 traces introduces a new 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-I2** command is issued again.

For each hop in the path, trace-I2 displays its input/output port, 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-I2 is enabled on the Brocade devices. It can be used to trace traffic only to devices.
- The devices that will participate in the trace-I2 protocol must be assigned to a VLAN and all devices on that VLAN must be Brocade devices that support the trace-I2 protocol.
- Brocade devices, as well as other vendor devices, that do not support the trace-I2 protocol, simply forward trace-I2 packets without a reply. Hence, these devices are transparent to the trace-I2 protocol.
- The destination for the packet with the trace-I2 protocol must be a device that supports the trace-I2 protocol and the destination cannot be a client, such as a personal computer, or devices from other vendors.

Tracing a traffic path

The trace-I2 protocol is enabled on a VLAN. You can trace the traffic path of a packet by entering a command such as the following.

```
Brocade(config)#trace-l2 vlan 10 2.2.2.2
```

Syntax: [no] trace-I2 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 exist, trace-I2 collects L2 topology information which can be displayed by issuing a **trace-I2 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-I2** 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 traceI2 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-l2 vlan** command, then enter the **trace-l2 show** command as in the following example.

```
Brocade(config)#trace-l2 vlan 10
Vlan 10 L2 topology probed, use "trace-l2 show" to display
Brocade(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 *tll-exceeded* error message to the source. The next packet has a TTL (hop limit) value of 2 and the second router replies with the *tll-exceeded* error message. This process continues till the destination host receives the packets and returns an ICMPv6 Echo Reply message.

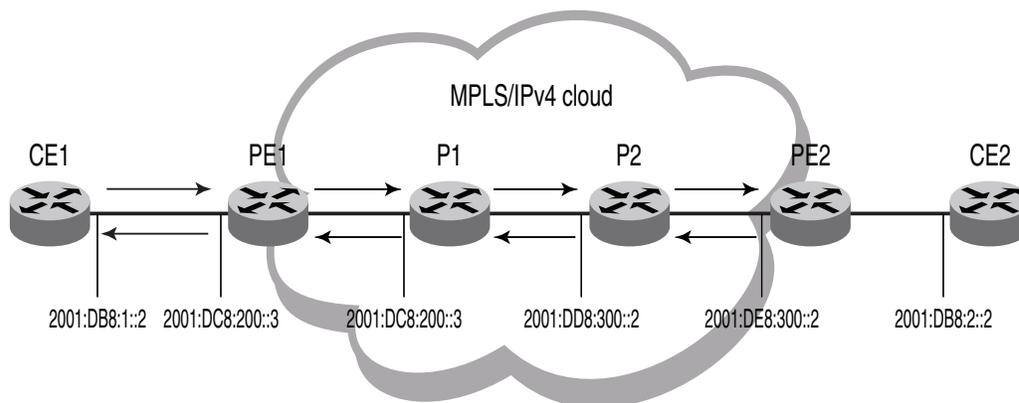
6 IPv6 Traceroute over an MPLS network

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

Figure 14 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 14 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 206 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:

```
Brocade# 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 [Figure 14](#) 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.

6 IPv6 Traceroute over an MPLS network

```
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 **ipv6 icmp 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:

```
Brocade(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-response use-lsp** command to use the configured LSPs.

To specify using LSP to forward the ICMPv6 messages with MPLS label extensions, enter the following command:

```
Brocade(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:

```
Brocade(config)# ipv6 icmp mpls-response no-label-extensions
```

To disable the IPv6 Traceroute over MPLS feature, enter the following command:

```
Brocade(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 208) 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 209) 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 208) 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 209) 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 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 34 lists the TLVs defined in RFC 3479 that are included in an echo request and reply.

TABLE 34 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
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 Brocade 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 Brocade 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 for 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 Brocade 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”](#)
- [“Executing RSVP LSP ping”](#)
- [“Executing LDP LSP traceroute”](#)
- [“Executing RSVP LSP traceroute”](#)

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.

```
Brocade)# 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.
Brocade)#
```

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 ipv4 address]`

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.

```
Brocade# 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.
Brocade)#
```

Syntax: `ping mpls rsvp lsp lsp-name | session tunnel-source-address tunnel-destination-address tunnel-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 specifies 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 an MPLS echo request from the ingress to the egress LSR.

To perform the LDP LSP traceroute operation, use the following command.

```
Brocade# 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)
Brocade#
```

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 ipv4 address]`

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-mode 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 **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 an MPLS echo request from the ingress to the egress LSR.

To perform the RSVP LSP traceroute operation, use the following command.

```
Brocade # 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)
Brocade#
```

Syntax: `traceroute mpls rsvp lsp lsp-name | session tunnel-source-address tunnel-destination-address tunnel-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-mode 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 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.

```
Brocade # 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.

6 LSP ping and traceroute

```
Brocade # 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.

1. Domain configuration.

```
Brocade(config-cfm)# domain-name ISID_domain level 7
```

Syntax: [no] domain-name *name* level *value*

2. MA configuration.

```
Brocade(config-cfm-md-ISID_domain)# ma-name ISID_2000 esi isid_1 isid 2000
priority 7
```

Syntax: [no] ma-name *name* esi *esiname* isid *isidvalue* priority *value*

3. MEP configuration.

```
Brocade(config-cfm-md-ISID_domain-ma-ISID_2000)# mep 1 down port eth 1/1
```

Syntax: [no] mep *id* dir *port* *portid*

Sample configuration

```
Brocade_BEB_1(config)#interface eth 1/1
Brocade_BEB_1(config-if-e1000-1/1)#enable
Brocade_BEB_1(config-if-e1000-1/1)#port-type backbone-network
Brocade_BEB_1(config)#esi isid_1 encapsulation isid
Brocade_BEB_1(config-esi-isid_esi_1)#isid 2000
Brocade_BEB_1(config)# esi bvlan_1 encapsulation bvlan
Brocade_BEB_1(config-esi-bvlan_1)#vlan 200
Brocade_BEB_1(config-esi-bvlan_1-vlan-200)#tagged eth1/1
Brocade_BEB_1(config-esi-bvlan_1)#esi-client isid_1
```

6 CFM monitoring for ISID

```
Brocade_BEB_2(config)#interface ethernet 1/2
Brocade_BEB_2(config-if-e1000-1/2)#enable
Brocade_BEB_2(config-if-e1000-1/2)#port-type backbone-network
Brocade_BEB_2(config)#esi isid_1 encapsulation isid
Brocade_BEB_2(config-esi-isid_esi_1)#isid 2000
Brocade_BEB_2(config)# esi bvlan_1 encapsulation bvlan
Brocade_BEB_2(config-esi-bvlan_1)#vlan 200
Brocade_BEB_2(config-esi-bvlan_1-vlan-200)#tagged ethernet 1/2
Brocade_BEB_2(config-esi-bvlan_1)#esi-client isid_esi_1
```

```
Brocade_BCB(config)#tag-value tag1 88A8
Brocade_BCB(config)#interface ethernet 1/1
Brocade_BCB(config-if-e1000-1/1)#enable
Brocade_BCB(config)#interface ethernet 1/2
Brocade_BCB(config-if-e1000-1/2)#enable
Brocade_BCB(config)#vlan 200
Brocade_BCB(config-vlan-200)#tagged eth 1/1
Brocade_BCB(config-vlan-200)#tagged eth 1/2
```

Sample configuration for ISID CFM

The below configuration shows the sample configuration for ISID CFM

```
Brocade_BEB_1(config)#cfm-enable
Brocade_BEB_1(config-cfm)#domain-name ISID_domain level 7
Brocade_BEB_1(config-cfm-md-ISID_domain)#ma-name ISID_2000 esi isid_1 isid 2000
priority 7
Brocade_BEB_1(config-cfm-md-ISID_domain-ma-ISID_2000)#ccm-interval 1-second
Brocade_BEB_1(config-cfm-md-ISID_domain-ma-ISID_2000)#mep 1 down port eth 1/1
Brocade_BEB_1(config-cfm-md-ISID_domain-ma-ISID_2000)#
```

```
Brocade_BEB_2(config)#cfm-enable BEB_2(config-cfm)#domain-name ISID_domain level
7 Brocade_BEB_2(config-cfm-md-ISID_domain)#ma-name ISID_2000 esi isid_1 isid 2000
priority 7
Brocade_BEB_2(config-cfm-md-ISID_domain-ma-ISID_2000)#ccm-interval 1-second
Brocade_BEB_2(config-cfm-md-ISID_domain-ma-ISID_2000)#mep 2 down port eth 1/2
Brocade_BEB_2(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

```
Brocade_BEB_1#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

```

```

Brocade_BEB_2#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

```

Brocade_BEB_1#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

```

```

Brocade_CER_2_2#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.

```
Brocade_BEB_1#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
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.

```
Brocade_BEB_1#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.

```
Brocade_BEB_1#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.

```
Brocade_DUT_1(config-cfm)#domain-name d7 level 7
```

Syntax: **domain-name** *name* **level** *value*

2. MA configuration.

```
Brocade_DUT_1(config-cfm-md-d7)#ma-name link link-ma priority 7
```

Syntax: **ma-name** *name* **link-ma** **priority** *value*

3. MEP configuration.

```
Brocade_DUT_1(config-cfm-md-d7-ma-link)#mep 1 down port eth 1/1
```

Syntax: **mep** *ID* **dir** **port** *portID*

4. Individual -link monitor configuration.

```
Brocade(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.

```

Brocade_DUT_1(config)#cfm-enable
Brocade_DUT_1(config-cfm)#domain-name d7 level 7
Brocade_DUT_1(config-cfm-md-d7)#ma-name link link-ma priority 7
Brocade_DUT_1(config-cfm-md-d7-ma-link)#ccm-interval 1-second
Brocade_DUT_1(config-cfm-md-d7-ma-link)#mep 1 down port eth 1/1

Brocade_DUT_2(config)#cfm-enable DUT_2(config-cfm)#domain-name d7 level 7
Brocade_DUT_2(config-cfm-md-d7)#ma-name link link-ma priority 7
Brocade_DUT_2(config-cfm-md-d7-ma-link)#ccm-interval 1-second
Brocade_DUT_2(config-cfm-md-d7-ma-link)#mep 3 down port eth 1/1
Brocade_DUT_2(config-cfm-md-d7-ma-link)#mep 4 down port eth 1/2

Brocade_DUT_3(config)#cfm-enable
Brocade_DUT_3(config-cfm)#domain-name d7 level 7
Brocade_DUT_3(config-cfm-md-d7)#ma-name link link-ma priority 7
Brocade_DUT_3(config-cfm-md-d7-ma-link)#ccm-interval 1-second
Brocade_DUT_3(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

```

Brocade_DUT_1#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

```

```

Brocade_DUT_2#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

```

```

Brocade_DUT_3#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

```

Brocade_DUT_1#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

```

```

Brocade_DUT_2#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
====  =====

```

6 CFM monitoring for ISID

1	0000.00ef.2a0b	N/A	799	1/1	1	OK	None
2	0000.0011.86d1	N/A	796	1/2	1	OK	None

```
Brocade_DUT_3#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.

```
Brocade_DUT_1#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
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.

```
Brocade_DUT_1#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. Use the following command to enable port-status-tlv at the MEP level.

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.

```

Brocade_DUT_1(config)#cfm-enable
Brocade_DUT_1(config-cfm)#domain-name customer level 7
Brocade_DUT_1(config-cfm-md-customer)#ma-name admin vlan-id 100 priority 7
Brocade_DUT_1(config-cfm-md-customer-ma-admin)#ccm-interval 1-second
Brocade_DUT_1(config-cfm-md-customer-ma-admin)#mep 1 down tlv-type
port-status-tlv port eth 1/1

Brocade_DUT_2(config)#cfm-enable
Brocade_DUT_2(config-cfm)#domain-name customer level 7
Brocade_DUT_2(config-cfm-md-customer)#ma-name admin vlan-id 100 priority 7
Brocade_DUT_2(config-cfm-md-customer-ma-admin)#ccm-interval 1-second

Brocade_DUT_3(config)#cfm-enable
Brocade_DUT_3(config-cfm)#domain-name customer level 7
Brocade_DUT_3(config-cfm-md-customer)#ma-name admin vlan-id 100 priority 7
Brocade_DUT_3(config-cfm-md-customer-ma-admin)#ccm-interval 1-second
Brocade_DUT_3(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

6 CFM monitoring for ISID

```

Brocade_DUT_1#show cfm
Domain: customer
Index: 1
Level: 7
Maintenance association: admin
Ma Index: 1
CCM interval: 1000 ms
VLAN ID: 100
Priority: 7
MEP Direction MAC PORT PORT-STATUS-TLV
=== =====
1 DOWN 0000.0011.86d1 ethe 1/1 Y

```

```

Brocade_DUT_3#show cfm
Domain: customer
Index: 1
Level: 7
Maintenance association: admin
Ma Index: 1
CCM interval: 1000 ms
VLAN ID: 100
Priority: 7
MEP Direction MAC PORT PORT-STATUS-TLV
==== =====
2 DOWN 0000.00ef.2a0b ethe 1/2 Y

```

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.

```

Brocade_DUT_1#show cfm connectivity domain customer ma admin rmep-id 2
Domain: customer Level: 7
Maintenance association: admin VLAN VLAN/VPLS/VLL ID: 100 Priority: 7
CCM interval: 1000 ms
RMEP MAC PORT Oper Age CCM RDI Port Intf Intvl Seq core
State Val Cnt Status Status Error Error Fault level
=====
2 0000.00ef.2a0b 1/1 OK 547 552 N 2 0 N N N
Brocade_DUT_1#

```

```

Brocade_DUT_3#show cfm connectivity domain customer ma admin rmep-id 1
Domain: customer Level: 7
Maintenance association: admin VLAN VLAN/VPLS/VLL ID: 100 Priority: 7
CCM interval: 1000 ms
RMEP MAC PORT Oper Age CCM RDI Port Intf Intvl Seq core
State Val Cnt Status Status Error Error Fault level
=====
1 0000.0011.86d1 1/2 OK 590 590 N 2 0 N N N

```

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.

```
Brocade_DUT_1(config)#cfm-enable
Brocade_DUT_1(config-cfm)#domain-name customer level 7
Brocade_DUT_1(config-cfm-md-customer)#ma-name admin vlan-id 100 priority 7
Brocade_DUT_1(config-cfm-md-customer-ma-admin)#ccm-interval 1-second
Brocade_DUT_1(config-cfm-md-customer-ma-admin)#mep 1 down port eth 1/1
```

```
Brocade_DUT_2(config)#cfm-enable
Brocade_DUT_2(config-cfm)#domain-name customer level 7
Brocade_DUT_2(config-cfm-md-customer)#ma-name admin vlan-id 100 priority 7
Brocade_DUT_2(config-cfm-md-customer-ma-admin)#ccm-interval 1-second
```

```
Brocade_DUT_3(config)#cfm-enable
Brocade_DUT_3(config-cfm)#domain-name customer level 7
Brocade_DUT_3(config-cfm-md-customer)#ma-name admin vlan-id 100 priority 7
Brocade_DUT_3(config-cfm-md-customer-ma-admin)#ccm-interval 1-second
Brocade_DUT_3(config-cfm-md-customer-ma-admin)#mep 2 down port ethe 1/2
```

```
Brocade_DUT_4(config)#cfm-enable
Brocade_DUT_4(config-cfm)#domain-name customer level 7
Brocade_DUT_4(config-cfm-md-customer)#ma-name admin vlan-id 100 priority 7
Brocade_DUT_4(config-cfm-md-customer-ma-admin)#ccm-interval 1-second
Brocade_DUT_4(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.

```
Brocade_DUT_1#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
Brocade_DUT_1#
```

show cfm connectivity

The **show cfm connectivity** command shows the RDI received by RMEP.

The following command shows the RDI received by RMEP

```
Brocade_DUT_3#show cfm connectivity domain customer ma admin rmep-id 1
Domain: customer Level: 7
Maintenance association: admin VLAN VLAN/VPLS/VLL ID: 100 Priority: 7
CCM interval: 1000 ms
RMEP   MAC           PORT     Oper   Age    CCM   RDI   Port   Intf   Intvl  Seq   core
      State  Val    Cnt   Status Status Error Error Fault level
=====
1      0000.0011.3301  1/2     OK     590   590   Y     0     0     N     N     N
Brocade_DUT_3#
```

```
Brocade_DUT_1#show cfm connectivity domain customer ma admin rmep-id 2
Domain: customer Level: 7
Maintenance association: admin VLAN VLAN/VPLS/VLL ID: 100 Priority: 7
CCM interval: 1000 ms
RMEP   MAC           PORT     Oper   Age    CCM   RDI   Port   Intf   Intvl  Seq   core
      State  Val    Cnt   Status Status Error Error Fault level
=====
2      0000.00ef.2a0b  1/1     OK     690   590   Y     0     0     N     N     N
Brocade_DUT_1#
```

Syntax: show cfm connectivity domain customer ma admin rmep-id id

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.

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.

```
Brocade(config-cfm)#loss-measurement lmm initiator 1
Brocade(config-cfm-loss-measurement-lmm-initiator-1)#
```

Syntax: `lmm initiator session_id`

2. LMM Initiator session configuration.

Configure the LMM session.

```
Brocade(config-cfm-loss-measurement-lmm-initiator-1)#domain mdl ma mal src-mep 1
target-mep 2
```

Syntax: `domain name ma name src-mep id target-mep id`

3. LMM session CoS configuration.

```
Brocade(config-cfm-loss-measurement-lmm-initiator-1)#Cos 1
Brocade(config-cfm-loss-measurement-lmm-initiator-1)#
```

Syntax: `Cos value`

4. LMM session Tx-interval configuration.

```
Brocade(config-cfm-loss-measurement-lmm-initiator-1)#Tx-interval 10
Brocade(config-cfm-loss-measurement-lmm-initiator-1)#
```

Syntax: `Tx-interval timer_value`

5. LMM session measurement-interval configuration.

```
Brocade(config-cfm-loss-measurement-lmm-initiator-1)#Measurement-interval 10
Brocade(config-cfm-loss-measurement-lmm-initiator-1)#
```

Syntax: `Measurement-interval timer_value`

6. LMM session threshold configuration

```
Brocade(config-cfm-loss-measurement-lmm-initiator-1)#threshold forward average
5000 maximum 10000
Brocade(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.

```
Brocade(config-cfm)#loss-measurement lmm responder 1
Brocade(config-cfm-loss-measurement-lmm-responder-1)#
```

Syntax: `loss-measurement lmm responder session_id`

2. LMM responder session configuration.

```
Brocade(config-cfm-loss-measurement-lmm-responder-1)#domain mdl ma mal src-mep 2
target-mep 1
Brocade(config-cfm-loss-measurement-lmm-responder-1)#
```

Syntax: `domain name ma name src-mep id target-mep id`

3. LMM session CoS configuration.

```
Brocade(config-cfm-loss-measurement-lmm-responder-1)#Cos 1
Brocade(config-cfm-loss-measurement-lmm-responder-1)#
```

Syntax: `Cos value`

Starting LMM session responder

Use the **start** command to start the session responder.

```
Brocade(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.

```
Brocade(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.

```
Brocade(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.

```
Brocade(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

```
BrocadeCE-1(config)# cfm
BrocadeCE-1(config-cfm)#loss-measurement lmm initiator 1
BrocadeCE-1(config-cfm-loss-measurement-lmm-initiator-1)#domain md1 ma ma1
src-mep 3 target-mep 4
BrocadeCE-1(config-cfm-loss-measurement-lmm-initiator-1)#Cos 2
BrocadeCE-1(config-cfm-loss-measurement-lmm-initiator-1)#tx-interval 10
BrocadeCE-1(config-cfm-loss-measurement-lmm-initiator-1)#measurement-interval 10
```

CE-2 Configuration

```
BrocadeCE-2(config)# cfm
BrocadeCE-2(config-cfm)# loss-measurement lmm responder 1
```

```
BrocadeCE-2(config-cfm-loss-measurement-lmm-responder-1)#domain md1 ma ma1
src-mep 4 target-mep 3
BrocadeCE-2(config-cfm-loss-measurement-lmm-responder-1)#Cos 2
```

Configuration example for VPLS tagged endpoints

PE-1 Configuration (Initiator)

```
BrocadePE-1(config-cfm)# loss-measurement lmm initiator 1
BrocadePE-1(config-cfm-loss-measurement-lmm-initiator-1)#domain md1 ma ma1
src-mep 3 target-mep 4
BrocadePE-1(config-cfm-loss-measurement-lmm-initiator-1)#Cos 2
BrocadePE-1(config-cfm-loss-measurement-lmm-initiator-1)#tx-interval 10
BrocadePE-1(config-cfm-loss-measurement-lmm-initiator-1)#measurement-interval 10
```

PE-2 Configuration (Responder)

```
BrocadePE-2(config-cfm)# loss-measurement lmm responder 1
BrocadePE-2(config-cfm-loss-measurement-lmm-responder-1)#domain md1 ma ma1
src-mep 4 target-mep 3
BrocadePE-2(config-cfm-loss-measurement-lmm-responder-1)#Cos 2
```

Configuration example for VPLS untagged endpoints

PE-1 Configuration (Initiator)

```
BrocadePE-1(config-cfm)# loss-measurement lmm initiator 1
BrocadePE-1(config-cfm-loss-measurement-lmm-initiator-1)#domain md1 ma ma1
src-mep 3 target-mep 4
BrocadePE-1(config-cfm-loss-measurement-lmm-initiator-1)#Cos 8
BrocadePE-1(config-cfm-loss-measurement-lmm-initiator-1)#tx-interval 10
BrocadePE-1(config-cfm-loss-measurement-lmm-initiator-1)#measurement-interval 10
```

PE-2 Configuration (Responder)

```
BrocadePE-2(config-cfm)# loss-measurement lmm responder 1
BrocadePE-2(config-cfm-loss-measurement-lmm-responder-1)#domain md1 ma ma1
src-mep 4 target-mep 3
BrocadePE-2(config-cfm-loss-measurement-lmm-responder-1)#Cos 8
```

Configuration example for VPLS tagged and untagged endpoints

PE-1 Configuration (Initiator)

```
BrocadePE-1(config-cfm)# loss-measurement lmm initiator 1
BrocadePE-1(config-cfm-loss-measurement-lmm-initiator-1)#domain md1 ma ma1
src-mep 3 target-mep 4
BrocadePE-1(config-cfm-loss-measurement-lmm-initiator-1)#Cos 8
BrocadePE-1(config-cfm-loss-measurement-lmm-initiator-1)#tx-interval 10
BrocadePE-1(config-cfm-loss-measurement-lmm-initiator-1)#measurement-interval 10
```

PE-2 Configuration (Responder)

```
BrocadePE-2(config-cfm)# loss-measurement lmm responder 1
BrocadePE-2(config-cfm-loss-measurement-lmm-responder-1)#domain md1 ma ma1
src-mep 4 target-mep 3
BrocadePE-2(config-cfm-loss-measurement-lmm-responder-1)#Cos 8
```

Starting LMM Sessions

Start the responder before starting the initiator.

CE-2 Configuration

```
BrocadeCE-2(config)# cfm
BrocadeCE-2(config-cfm)# loss-measurement lmm responder 1
BrocadeCE-2(config-cfm-loss-measurement-lmm-responder-1)#start now
```

CE-1 Configuration

```
BrocadeCE-1(config)# cfm
BrocadeCE-1(config-cfm)# loss-measurement lmm initiator 1
BrocadeCE-1(config-cfm-loss-measurement-lmm-initiator-1)#start now
```

Stopping LMM Sessions

Stop the initiator before stopping the responder.

CE-1 Configuration

```
BrocadeCE-1(config)# cfm
BrocadeCE-1(config-cfm)# loss-measurement lmm initiator 1
BrocadeCE-1(config-cfm-loss-measurement-lmm-initiator-1)#stop now
```

CE-2 Configuration

```
BrocadeCE-2(config-cfm)# loss-measurement lmm responder 1
BrocadeCE-2(config-cfm-loss-measurement-lmm-responder-1)#stop now
```

Clearing history statistics per session

```
BrocadeCE-1(config-cfm)# loss-measurement lmm initiator 1
BrocadeCE-1(config-cfm-loss-measurement-lmm-initiator-1)#clear-stat
```

Clearing history statistics globally

```
BrocadeCE-1(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

6 One-way Delay Measurement

```
<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:

- TxTimeStamp: 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 Brocade 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 TxTimeStamp 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 TxTimeStamp value
- RxTimef, which is the time at reception of the one-way delay measurement frame
- $\text{Frame Delay}_{\text{one-way}} = \text{RxTimef} - \text{TxTimeStamp}$

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

1. VLAN Creation.

```
Brocade(config)#vlan 20
Brocade(config-vlan-20)#tagged ethernet 1/1
```

Syntax: `vlan id`

CFM Configuration

1. Enabling CFM.

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

Syntax: `cfm-enable`

2. Domain Configuration

```
Brocade(config)#cfm-enable
Brocade(config-cfm)#domain-name md1 level 7
Brocade(config-cfm-md-md1)#
```

Syntax: `domain-name md_name [id id] level level`

3. MA Configuration

```
Brocade(config-cfm-md-md1)#ma-name ma1 vlan 20 priority 4
Brocade(config-cfm-md-md1-ma-ma1)
```

Syntax: `ma-name ma_name [id id] vlan-id vlan | vpls-id vpls priority priority`

4. MEP Configuration

```
Brocade(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

```
Brocade(config)#router mpls
Brocade(config-mpls)#vpls vpls100 100
Brocade(config-mpls-vpls-vpls100)#vlan 100
Brocade(config-mpls-vpls-vpls100-vlan-10)#tagged Ethernet 1/1
```

Syntax: `vpls vpls-name id`

Syntax: `vlan vlan-id`

Syntax: `tagged ethernet slot/port`

Creation of VLL

```
Brocade(config)#router mpls
Brocade(config-mpls)#vll vll100 100
Brocade(config-mpls-vll-vll100)#vlan 100
Brocade(config-mpls-vll-vll100-vlan-10)#tagged Ethernet 1/1
```

Syntax: `vll vll-name id`

Syntax: `vlan vlan-id`

Syntax: `tagged ethernet slot/port`

CFM Configuration

1. Enable CFM.

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

Syntax: `cfm-enable`

2. Configure the domain.

```
Brocade(config)#cfm-enable
Brocade(config-cfm)#domain-name md1 level 7
Brocade(config-cfm-md-md1)#
```

Syntax: `domain-name md_name id id level level`

3. Configure MA.

```
Brocade(config-cfm-md-md1)#ma-name ma1 vpls-id 100 priority 4
Brocade(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.

```
Brocade(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

```
Brocade(config)#cfm
Brocade(config-cfm)# oneway-dm initiator 1
Brocade (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

```
Brocade (config-cfm-oneway-dm-initiator-1)# cos 4
Brocade (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

```
Brocade(config)#cfm
Brocade(config-cfm)# oneway-dm receiver 1
Brocade(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

```
DUT (config-cfm-oneway-dm-receiver-1)# cos 4
DUT (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

```
DUT (config-cfm-oneway-dm-receiver-1)# threshold max 50
DUT (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.

```
DUT (config-cfm-oneway-dm-receiver-1)# start now
```

Syntax: `start now | after HH:MM:SS | HH:MM:SS> daily`

- a. Start the one-way delay measurement receiver session after a period of time.

```
DUT (config-cfm-oneway-dm-receiver-1)# start after 01:30:00
```

The example above will start after 1 hour and 30 minutes.

- b. Start 1DM Receiver Session exactly at given time.

```
DUT (config-cfm-oneway-dm-receiver-1)# start 09:30:00
```

The example above will be started exactly at 09:30 AM.

- c. Start 1DM Receiver Session daily at given time.

```
DUT (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.

```
DUT (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.

```
DUT (config-cfm-oneway-dm-initiator-1)# stop now
```

Syntax: `stop now | after HH:MM:SS | HH:MM:SS daily`

Stopping one-way delay measurement receiver session

A receiver session can be stopped immediately, after a specified amount of time, or at a specific time.

```
DUT (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 one-way delay measurement over VLAN

VLAN configurations:

CE-1 Configuration

```
CE-1(config)# interface ethernet 1/1
CE-1(config-if-e10000-1/1)# enable
CE-1(config-if-e10000-1/1)# exit

CE-1(config)# vlan 10
CE-1(config-vlan-10)# tagged ethernet 1/1
```

CE-2 configuration

```
CE-2(config)# interface ethernet 1/1
CE-2(config-if-e10000-1/1)# enable
CE-2(config-if-e10000-1/1)# exit

CE-2(config)# vlan 10
CE-2(config-vlan-10)# tagged ethernet 1/1
```

CFM configurations:

CE-1 configuration

```
CE-1(config)# cfm-enable
CE-1(config-cfm)# domain md1 level 7
CE-1(config-cfm-md-md1)# ma ma1 vlan 10 priority 4
CE-1(config-cfm-md-md1-ma-ma1)# mep 1 down port ethernet 1/1
```

CE-2 configuration

```
CE-1(config)# cfm-enable
CE-1(config-cfm)# domain md1 level 7
CE-1(config-cfm-md-md1)# ma ma1 vlan 10 priority 4
CE-1(config-cfm-md-md1-ma-ma1)# mep 101 down port ethernet 1/1
```

One-way delay measurement configurations:

CE-1 Configuration

```
CE-1(config)# cfm
CE-1(config-cfm)# oneway-dm initiator 1
CE-1(config-cfm-oneway-dm-initiator-1)#domain md1 ma ma1 src-mep 1 target-mep 2
CE-1(config-cfm-oneway-dm-initiator-1)#tx-interval 10
```

6 One-way Delay Measurement

CE-2 Configuration

```
CE-2(config)# cfm
CE-2(config-cfm)# oneway-dm receiver 2
CE-2 (config-cfm-oneway-dm-receiver-2)# domain mdl ma mal src-mep 2 target-mep 1
CE-2 (config-cfm-oneway-dm-receiver-2)# measurement-interval 10
```

Starting 1DM Sessions:

CE-1 Configuration

```
CE-1(config-cfm-oneway-dm-initiator-1)# start now
```

CE-2 Configuration

```
CE-2 (config-cfm-oneway-dm-receiver-2)# start now
```

Stopping 1DM Sessions:

CE-1 Configuration

```
CE-1(config-cfm-oneway-dm-initiator-1)# stop now
```

CE-2 Configuration

```
CE-2 (config-cfm-oneway-dm-receiver-2)# stop now
```

Sample Configuration one-way delay measurement over VPLS or VLL

VPLS Configurations:

PE-1 Configuration

```
PE-1(config)# interface ethernet 1/1
PE-1(config-if-e10000-1/1)# enable
PE-1(config-if-e10000-1/1)# exit

PE-1(config)# router mpls
PE-1(config-mpls)# vpls vpls100 100
PE-1(config-mpls-vpls-vpls100)# vlan 10
PE-1(config-mpls-vpls-vpls100-vlan-10)# tagged ethernet 1/1
PE-1(config-mpls-vpls-vpls100-vlan-10)# end
```

PE-2 Configuration

```
PE-2(config)# interface ethernet 1/1
PE-2(config-if-e10000-1/1)# enable
PE-2(config-if-e10000-1/1)# exit

PE-2(config)# router mpls
PE-2(config-mpls)# vpls vpls100 100
PE-2(config-mpls-vpls-vpls100)# vlan 10
PE-2(config-mpls-vpls-vpls100-vlan-10)# tagged ethernet 1/1
PE-2(config-mpls-vpls-vpls100-vlan-10)# end
```

VLL Configurations:

PE-1 Configuration

```
PE-1(config)# interface ethernet 1/1
PE-1(config-if-e10000-1/1)# enable
PE-1(config-if-e10000-1/1)# exit
```

```

PE-1(config)# router mpls
PE-1(config-mpls)# vll vll100 100
PE-1(config-mpls-vll-vll100)# vlan 10
PE-1(config-mpls-vll-vll100-vlan-10)# tagged ethernet 1/1
PE-1(config-mpls-vll-vll100-vlan-10)# end

```

PE-2 Configuration

```

PE-2(config)# interface ethernet 1/1
PE-2(config-if-e10000-1/1)# enable
PE-2(config-if-e10000-1/1)# exit

PE-2(config)# router mpls
PE-2(config-mpls)# vll vpls100 100
PE-2(config-mpls-vll-vll100)# vlan 10
PE-2(config-mpls-vll-vll100-vlan-10)# tagged ethernet 1/1
PE-2(config-mpls-vll-vll100-vlan-10)# end

```

CFM Configurations:

PE-1 Configuration

```

PE-1(config)# cfm-enable
PE-1(config-cfm)# domain md1 level 7
PE-1(config-cfm-md-md1)# ma ma1 vpls 100 priority 4
PE-1(config-cfm-md-md1-ma-ma1)# mep 1 up vlan 10 port ethernet 1/1

```

PE-2 Configuration

```

PE-2(config)# cfm-enable
PE-2(config-cfm)# domain md1 level 7
PE-2(config-cfm-md-md1)# ma ma1 vpls 100 priority 4
PE-2(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

```

CE-1(config)# cfm
CE-1(config-cfm)# oneway-dm initiator 1
CE-1(config-cfm-oneway-dm-initiator-1)#domain md1 ma ma1 src-mep 1 target-mep 2
CE-1(config-cfm-oneway-dm-initiator-1)#tx-interval 10

```

CE-2 Configuration

```

CE-2(config)# cfm
CE-2(config-cfm)# oneway-dm receiver 2
CE-2 (config-cfm-oneway-dm-receiver-2)# domain md1 ma ma1 src-mep 2 target-mep 1
CE-2 (config-cfm-oneway-dm-receiver-2)# measurement-interval 10

```

Starting one-way delay measurement sessions:

CE-1 Configuration

```

CE-1(config-cfm-oneway-dm-initiator-1)# start now

```

CE-2 Configuration

```
CE-2 (config-cfm-oneway-dm-receiver-2)# start now
```

Stopping one-way delay measurement sessions:

CE-1 Configuration

```
CE-1(config-cfm-oneway-dm-initiator-1)# stop now
```

CE-2 Configuration

```
CE-2 (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>

```
Brocade# 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.

```
Brocade# 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:

```
Brocade# 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*

```
Brocade# 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.

6 Synthetic loss measurement

```
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.

```
DUT(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.

```
DUT(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.

```
DUT(config-cfm-loss-measurement-slm-initiator-1)# domain
```

Syntax: **domain**

Domain - is used to configure the domain name

```
DUT(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.

```
DUT(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).

```
DUT(config-cfm-loss-measurement-slm-initiator-1)# tx-interval 1
```

Syntax: **measurement-interval** interval

Interval - is used to configure SLM Measurement interval (default- 15min).

```
DUT(config-cfm-loss-measurement-slm-initiator-1)# threshold forward
```

Syntax: **threshold** [forward | backward] [average | maximum] value

Default values

Threshold Forward Average 0xFFFFFFFF milli-percent

Threshold Backward Average 0xFFFFFFFF milli-percent

Threshold Forward maximum 0xFFFFFFFF milli-percent

Threshold Backward maximum 0xFFFFFFFF milli-percent

Configuration examples

Sample configuration of synthetic loss measurement over VLAN

VLAN Configurations:

DUT1 Configuration

```
DUT1(config)# vlan 2
DUT1(config-vlan-2)# tagged ethernet 1/1
```

DUT2 Configuration

```
DUT2(config)# vlan 2
DUT2(config-vlan-2)# tagged ethernet 1/1
```

CFM Configurations:

DUT1 Configuration

```
DUT1(config)# cfm-enable
DUT1(config-cfm)# domain md1 level 7
DUT1(config-cfm-md-md1)# ma ma1 vlan 10 priority 4
DUT1(config-cfm-md-md1-ma-ma1)# mep 3 down port ethernet 1/1
```

DUT2 Configuration

```
DUT2(config)# cfm-enable
DUT2(config-cfm)# domain md1 level 7
DUT2(config-cfm-md-md1)# ma ma1 vlan 2 priority 4
DUT2(config-cfm-md-md1-ma-ma1)# mep 4 down port ethernet 1/1
```

SLM Configurations:

DUT1 Configuration

```
DUT1(config)# cfm
DUT1(config-cfm)# loss-measurement slm initiator 1
DUT1(config-cfm-loss-measurement-slm-initiator-1)#domain md1 ma ma1 src-mep 3
target-mep 4
DUT1(config-cfm-loss-measurement-slm-initiator-1)#cos 2
DUT1(config-cfm-loss-measurement-slm-initiator-1)#tx-interval 1
DUT1(config-cfm-loss-measurement-slm-initiator-1)#measurement-interval 1
```

DUT2 Configuration

```
DUT2(config)# cfm
DUT2(config-cfm)# loss-measurement slm responder 1
DUT2(config-cfm-loss-measurement-slm-responder-1)#domain md1 ma ma1 src-mep 4
target-mep 3
DUT2(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)

```
DUT2(config)# cfm
DUT2(config-cfm)# cfm loss-measurement slm responder 1
DUT2(config-cfm-loss-measurement-slm-responder-1)#start now
```

DUT1 Configuration (Initiator)

```
DUT1(config)# cfm
DUT1(config-cfm)# cfm loss-measurement slm initiator 1
DUT1(config-cfm-loss-measurement-slm-initiator-1)#start now
```

Stopping synthetic loss measurement sessions:**NOTE**

Stop the initiator before stopping the responder.

DUT1 Configuration

```
DUT1(config)# cfm
DUT1(config-cfm)# cfm loss-measurement slm initiator 1
DUT1(config-cfm-loss-measurement-slm-initiator-1)#stop now
```

DUT2 Configuration

```
DUT2(config)# cfm
DUT2(config-cfm)# cfm loss-measurement slm responder 1
DUT2(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.

```
DUT1(config)# cfm
DUT1(config-cfm)# cfm loss-measurement slm initiator 1
DUT1(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.

```
DUT1(config)# cfm
DUT1(config-cfm)# cfm loss-measurement slm clear-stat
```

Sample configuration - synthetic loss measurement over VPLS**VPLS Configurations:****LER1 Configuration**

```
LER1(config)# router mpls
LER1(config-mpls)# vpls vpls100 100
LER1(config-mpls-vpls-vpls100)# vlan 10
```

6 Synthetic loss measurement

```
LER1(config-mpls-vpls-vpls100-vlan-10)# tagged ethernet 1/1
LER1(config-mpls-vpls-vpls100-vlan-10)# end
```

LER2 Configuration

```
LER2(config)# router mpls
LER2(config-mpls)# vpls vpls100 100
LER2(config-mpls-vpls-vpls100)# vlan 10
LER2(config-mpls-vpls-vpls100-vlan-10)# tagged ethernet 1/1
LER2(config-mpls-vpls-vpls100-vlan-10)# end
```

CFM Configurations:

```
LER1 Configuration
LER1(config)# cfm-enable
LER1(config-cfm)# domain md1 level 7
LER1(config-cfm-md-md1)# ma ma1 vpls 100 priority 4
LER1(config-cfm-md-md1-ma-ma1)# mep 3 up vlan 10 port ethernet 1/1
```

LER2 Configuration

```
LER2(config)# cfm-enable
LER2(config-cfm)# domain md1 level 7
LER2(config-cfm-md-md1)# ma ma1 vpls 100 priority 4
LER2(config-cfm-md-md1-ma-ma1)# mep 4 up vlan 10 port ethernet 1/1
```

Synthetic loss measurement configurations:

LER1 Configuration

```
LER1(config)# cfm
LER1(config-cfm)# loss-measurement slm initiator 1
LER1(config-cfm-loss-measurement-slm-initiator-1)#domain md1 ma ma1 src-mep 3
target-mep 4
LER1(config-cfm-loss-measurement-slm-initiator-1)#cos 2
LER1(config-cfm-loss-measurement-slm-initiator-1)#tx-interval 1
LER1(config-cfm-loss-measurement-slm-initiator-1)#measurement-interval 1
```

LER2 Configuration

```
LER2(config)# cfm
LER2(config-cfm)# loss-measurement slm responder 1
LER2(config-cfm-loss-measurement-slm-responder-1)#domain md1 ma ma1 src-mep 4
target-mep 3
LER2(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)

```
LER2(config)# cfm
LER2(config-cfm)# cfm loss-measurement slm responder 1
LER2(config-cfm-loss-measurement-slm-responder-1)#start now
```

LER1 Configuration (Initiator)

```
LER1(config)# cfm
LER1(config-cfm)# cfm loss-measurement slm initiator 1
LER1(config-cfm-loss-measurement-slm-initiator-1)#start now
```

Stopping synthetic loss measurement sessions:**NOTE**

Stop the initiator before stopping the responder.

LER1 Configuration

```
LER1(config)# cfm
LER1(config-cfm)# cfm loss-measurement slm initiator 1
LER1(config-cfm-loss-measurement-slm-initiator-1)#stop now
```

LER2 Configuration

```
LER2(config)# cfm
LER2(config-cfm)# cfm loss-measurement slm responder 1
LER2(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.

```
LER1(config)# cfm
LER1(config-cfm)# cfm loss-measurement slm initiator 1
LER1(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.

```
LER1(config)# cfm
LER1(config-cfm)# cfm loss-measurement slm clear-stat
```

Sample configuration - Synthetic loss measurement over VLL**VLL Configurations:****LER1 Configuration**

```
LER1(config)# router mpls
LER1(config-mpls)# vll vll100 100
LER1(config-mpls-vll-vll100)# vlan 10
LER1(config-mpls-vll-vll100-vlan-10)# tagged ethernet 1/1
LER1(config-mpls-vll-vll100-vlan-10)# end
```

LER2 Configuration

```
LER2(config)# router mpls
LER2(config-mpls)# vll vll100 100
LER2(config-mpls-vll-vll100)# vlan 10
LER2(config-mpls-vll-vll100-vlan-10)# tagged ethernet 1/1
LER2(config-mpls-vll-vll100-vlan-10)# end
```

CFM Configurations:**LER1 Configuration**

```
LER1(config)# cfm-enable
LER1(config-cfm)# domain md1 level 7
LER1(config-cfm-md-md1)# ma ma1 vll 100 priority 4
LER1(config-cfm-md-md1-ma-ma1)# mep 3 up vlan 10 port ethernet 1/1
```

6 Synthetic loss measurement

LER2 Configuration

```
LER2(config)# cfm-enable
LER2(config-cfm)# domain md1 level 7
LER2(config-cfm-md-md1)# ma ma1 vll 100 priority 4
LER2(config-cfm-md-md1-ma-mal)# mep 4 up vlan 10 port ethernet 1/1
```

Synthetic loss measurement configurations:

LER1 Configuration

```
LER1(config)# cfm
LER1(config-cfm)# loss-measurement slm initiator 1
LER1(config-cfm-loss-measurement-slm-initiator-1)#domain md1 ma ma1 src-mep 3
target-mep 4
LER1(config-cfm-loss-measurement-slm-initiator-1)#cos 2
LER1(config-cfm-loss-measurement-slm-initiator-1)#tx-interval 1
LER1(config-cfm-loss-measurement-slm-initiator-1)#measurement-interval 1
```

LER2 Configuration

```
LER2(config)# cfm
LER2(config-cfm)# loss-measurement slm responder 1
LER2(config-cfm-loss-measurement-slm-responder-1)#domain md1 ma ma1 src-mep 4
target-mep 3
LER2(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)

```
LER2(config)# cfm
LER2(config-cfm)# cfm loss-measurement slm responder 1
LER2(config-cfm-loss-measurement-slm-responder-1)#start now
```

LER1 Configuration (Initiator)

```
LER1(config)# cfm
LER1(config-cfm)# cfm loss-measurement slm initiator 1
LER1(config-cfm-loss-measurement-slm-initiator-1)#start now
```

Stopping synthetic loss measurement sessions:

NOTE

Stop the initiator before stopping the responder.

LER1 Configuration

```
LER1(config)# cfm
LER1(config-cfm)# cfm loss-measurement slm initiator 1
LER1(config-cfm-loss-measurement-slm-initiator-1)#stop now
```

LER2 Configuration

```
LER2(config)# cfm
LER2(config-cfm)# cfm loss-measurement slm responder 1
LER2(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.

```
LER1(config)# cfm
LER1(config-cfm)# cfm loss-measurement slm initiator 1
LER1(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.

```
LER1(config)# cfm
LER1(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 session index

```
DUT1# 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 35 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

6 Synthetic loss measurement

TABLE 35 show cfm loss-measurement slm output

Row	Definition
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 *session index*

```
DUT1# 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 row index*

```
DUT1# 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

```

DUT1# 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>.
```

6 Synthetic loss measurement

Network Time Protocol

TABLE 36 Supported platforms for NTPs

Features supported	Brocade Netron XMR	Brocade MLX Series	Brocade Netron CES 2000 Series BASE package	Brocade Netron CES 2000 Series ME_PREM package	Brocade Netron CES 2000 Series L3_PREM package	Brocade Netron CER 2000 Series Base package	Brocade Netron CER 2000 Series Advanced Services package
Network Time Protocol	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Network Time Protocol (NTP) overview

Network Time Protocol (NTP) maintains uniform time across all devices in a network. The NTP commands support the configuration of an external time server to maintain synchronization between all local clocks in a network.

NTP has a hierarchical structure. At the highest level, or stratum, are precise hardware clocks, which can synchronize to highly accurate external time reference. These hardware clock devices are known as stratum 0 devices. A stratum 1 time server obtains time directly from a hardware clock and is the most accurate reference in the NTP hierarchy. All lower stratum devices obtain time from the stratum above over a network. As the network introduces timing discrepancies, lower stratum devices are a factor less accurate.

A hierarchical structure allows the overhead of providing time to many clients to be shared among many time servers. Not all clients need to obtain time directly from a stratum 1 reference, but can utilise stratum 2 or 3 references.

NTP operates on a client-server basis. A network time client periodically requests time from a time server. The time server responds with a packet of information containing a time stamp. The time stamp is then used by the client to synchronize its system time.

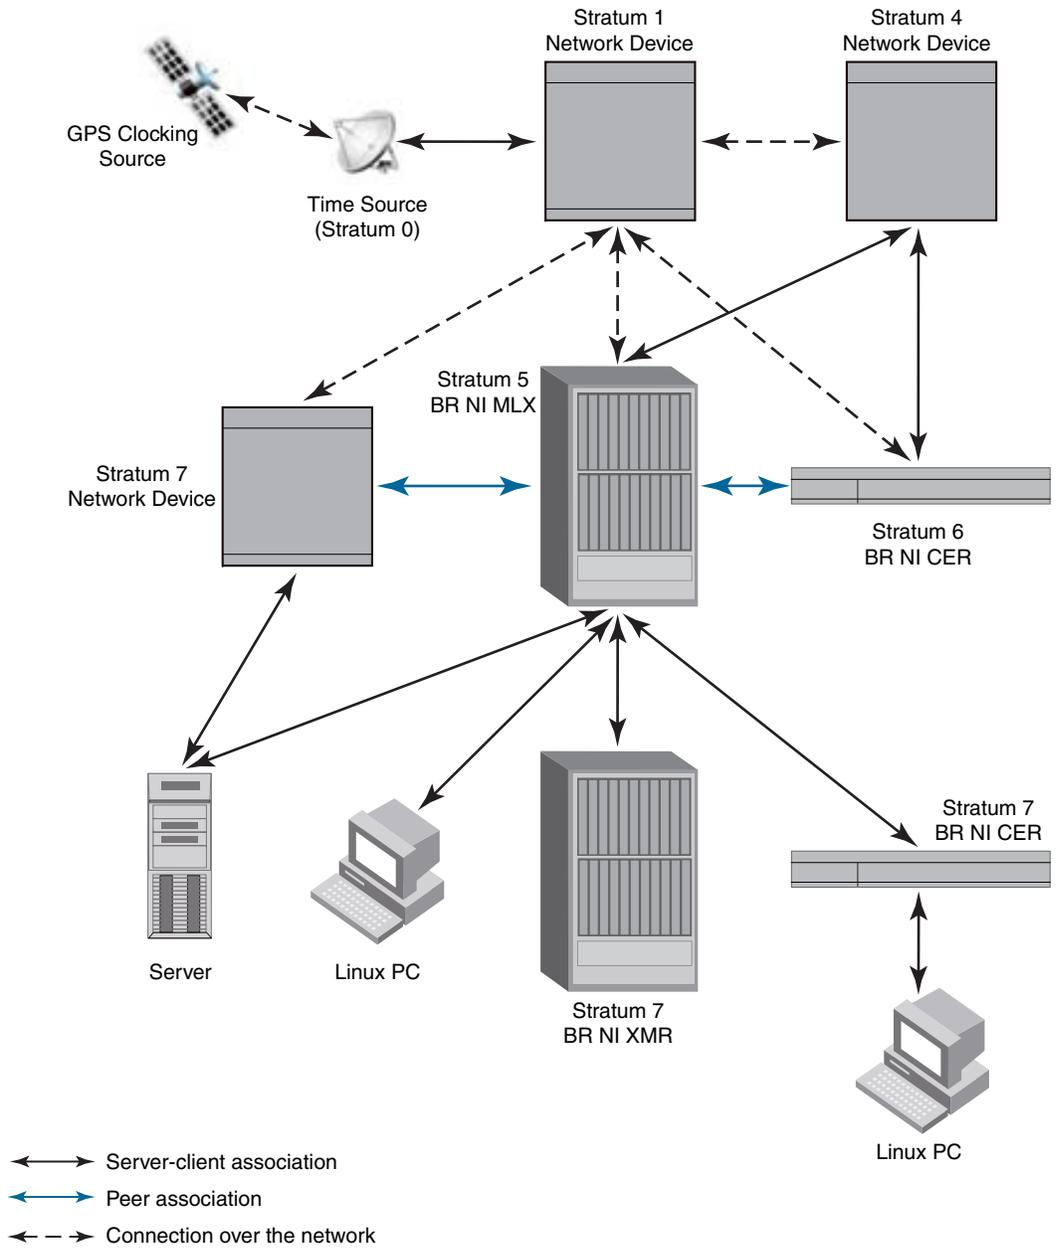
NTP uses UTC (Universal Time Coordinated) time, which is similar to GMT time. It knows nothing of local time zones or daylight-saving time. It is a function of the time client to apply an offset to the supplied time to adjust for local time. In this manner, a time server located anywhere in the world can provide synchronisation to a client located anywhere else in the world. It allows clients to utilise different time zone and daylight-saving properties.

To keep the time in your network current, it is recommended that each device have its time synchronized with at least four external NTP servers. External NTP servers should be synchronized among themselves in order to maintain time synchronization.

NOTE

Network Time Protocol (NTP) commands must be configured on each individual device.

FIGURE 15 NTP sample network configuration



How NTP works

NTP server

A **NTP server** will provide the correct network time on your device using the Network time protocol (NTP). Network Time Protocol can be used to synchronize the time on devices across a network. A NTP time server is used to obtain the correct time from a time source and adjust the local time in each connecting device.

The NTP server can operate in master mode to serve time using the local clock, when it has lost synchronization.

NTP client

An NTP client gets time responses from an NTP server or servers, and uses the information to calibrate its clock. This consists of the client determining how far its clock is off and adjusting its time to match that of the server. The maximum error is determined based on the round-trip time for the packet to be received.

NTP peer

NTP peer mode is intended for configurations where a group of devices operate as mutual backups for each other. If one of the devices lose a reference source, the time values can flow from the surviving peers to all the others.

The NTP peer can operate in:

Symmetric Active - When the peer is configured using the peer command.

Symmetric Passive - Dynamically learnt upon arrival of a NTP packet from the peer which is not configured. The symmetric passive association is removed on timeout or error.

NTP broadcast server

An NTP server can also operate in a broadcast or multicast mode. Both work similarly; broadcast servers send periodic time updates to a broadcast address, while multicast servers send periodic updates to a multicast address. Using broadcast packets can greatly reduce the NTP traffic on a network, especially for a network with many NTP clients.

The interfaces should be enabled with NTP broadcasting. The NTP broadcast server broadcasts the NTP packets periodically (every 64 sec) to subnet broadcast IP address of the configured interface.

NTP broadcast client

An NTP broadcast or multicast client listens for NTP packets on a broadcast or multicast address. When the first packet is received, it attempts to quantify the delay to the server in order to better quantify the correct time from later broadcasts. This is accomplished by a series of brief interchanges where the client and server act as a regular (non-broadcast) NTP client and server. Once these interchanges occur, the client has an idea of the network delay and thereafter can estimate the time based only on broadcast packets.

Synchronizing time

After the system peer is chosen, the system time is synchronized using one of the following ways based on the time difference with system peer:

- < 128 msec - The system clock is adjusted slowly towards the system peer time reference time.
- > 128 msec and < 1000 sec - The system clock is stepped to the system peer reference time and the NTP state information is cleared.
- > 1000 sec - NTP is operationally disabled. The admin should set the time to the current UTC time.

Configuration considerations of NTP

- NTP multicast server, client, and manycast client functionalities are not supported.
- While upgrading from R05.2.00 or lower versions to R05.3.00, the SNTP configuration will be ignored.
- On reboot or MP switchover, all the NTP state information will be lost and time synchronization will start from fresh. The time synchronized to real time clock is retained across reboot and MP switchover.
- The following SNTP MIB objects are not supported.
 - snNTPPollInterval
 - snNTPSync
 - All the objects in snNTPServerTable
- The web management support for SNTP is removed

The following optional features are not supported

- NTP version 4 Extension fields
- The NTP packets having control (6) or private (7) packet modes
- Autokey public key authentication
- NTP version 1 and 2
- Hostnames

Configuring NTP

Before you begin to configure NTP, you must use the **clock set** command to set the time on your device to within 1000 seconds of the coordinated Universal Time (UTC).

Changing to the NTP mode

Use the **ntp** command to enable the NTP client and server mode.

```
Brocade(config)# ntp
```

Syntax: ntp

Enabling NTP authentication

To enable Network Time Protocol (NTP) strict authentication, use the **authenticate** command. To disable the function, use the **no** form of this command.

```
Brocade(config-ntp)# authenticate
```

Syntax: [no] **authenticate**

Defining an authentication key

To define an authentication key for Network Time Protocol (NTP), use the **authentication-key** command. To remove the authentication key for NTP, use the **no** form of this command.

```
Brocade(config-ntp)# authentication-key key-id 1 md5 moof
```

Syntax: [no] **authentication-key** *key-id* **md5** *key string*

The valid *key-id* parameter is 1 to 65535.

MD5 is the message authentication support that is provided using the Message Digest 5 Algorithm. The key type md5 is currently the only key type supported.

The *key string* option is the value of the MD5 key. The maximum length of the key string may be defined up to 16 characters. Up to 32 keys may be defined.

Specifying a source interface

To use a particular source interface in Network Time Protocol (NTP) packets, use the **source-interface** command. To remove the specified source address, use the **no** form of this command.

NOTE

If the **source-interface** is not configured, then the lowest IP address in the outgoing interface will be used in the NTP packets.

```
Brocade(config-ntp)# source-interface ethernet 3/1
```

Syntax: [no] **source-interface** **ethernet** *slot/port* | **loopback** *num* | **ve** *num*

The **ethernet** *slot/port* parameter specifies the ethernet port number.

The **loopback** *num* parameter specifies the loopback interface number.

The **ve** *number* parameter specifies the virtual port number.

Enable or disable the VLAN containment for NTP

To enable or disable the VLAN containment for NTP, use the **access-control vlan** command. To remove the specified NTP VLAN configuration, use the **no** form of this command.

NOTE

The management interface is not part of any VLAN. When configuring the VLAN containment for NTP, it will not use the management interface to send or receive the NTP packets.

```
Brocade(config-ntp)# access-control vlan 100
```

Syntax: [no] **access-control vlan** *vlan-id*

The *vlan-id* parameter specifies the VLAN ID number.

Configuring the NTP client

To configure the device in client mode and specify the NTP servers to synchronize the system clock, use the **server** command. A maximum 8 NTP servers can be configured. To configured NTP server, use the **no** form of this command.

```
Brocade(config-ntp)#server 10.2.3.4 key 1234
```

Syntax: # [no] **server** *ipv4 address* | *ipv6 address* [**version** 3 | 4] [**key** *key id*] [**minpoll** *interval*] [**maxpoll** *interval*]

The *ipv4 address* | *ipv6 address* parameter is the IP address of the server providing the clock synchronization.

The **version** 3 | 4 option defines the Network Time Protocol (NTP) version number. If this option is not specified, then it defaults to 4.

The **key** *key id* option defines the authentication key. By default, no authentication key is configured.

The **minpoll** *interval* option is the shortest polling interval. The range is from 4 through 17. Default is 6. The interval argument is power of 2 (4=16s, 5=32s, 6=64s, 7=128s, 8=256s, 9=512s etc.).

The **maxpoll** *interval* option is the longest polling interval. The range is 4 through 17. Default is 10. The interval argument is calculated by the power of 2 (4=16s, 5=32s, 6=64s, 7=128s, 8=256s, 9=512s etc.).

Configuring the NTP peer

To configure the software clock to synchronize a peer or to be synchronized by a peer, use the **peer** command. A maximum of 8 NTP peers can be configured. To disable this capability, use the **no** form of this command.

```
Brocade(config-ntp)# peer 10.2.3.4 key 1234
```

Syntax: **[no] peer** *ipv4 address* | *ipv6 address* [**version** 3 | 4] [**key** *key id*] [**minpoll** *interval*] [**maxpoll** *interval*]

The *ipv4 address* | *ipv6 address* parameter is the IP address of the peer providing the clock synchronization.

The **version** 3 | 4 option defines the Network Time Protocol (NTP) version number. If this option is not specified, then it defaults to 4.

The **key** *key id* option defines the authentication key. By default, no authentication key is configured.

The **minpoll** *interval* option is the shortest polling interval. The range is from 4 through 17. Default is 6. The interval argument is power of 2 (4=16s, 5=32s, 6=64s, 7=128s, 8=256s, 9=512s etc.).

The **maxpoll** *interval* option is the longest polling interval. The range is 4 through 17. Default is 10. The interval argument is calculated by the power of 2 (4=16s, 5=32s, 6=64s, 7=128s, 8=256s, 9=512s etc.).

Configuring NTP on an interface

To configure the NTP interface context, use the **ntp-interface** command. The broadcast server or client is configured on selected interfaces. To remove the NTP broadcast configurations on the specified interface, use the **no** form of this command.

NOTE

The **ntp-interface** command is a mode change command, and will not be included in to the **show run output** unless there is configuration below that interface.

```
Brocade(config-ntp)# ntp-interface ethernet 2/13
Brocade(config-ntp-if-e1000-2/13)#
```

```
Brocade(config-ntp)# ntp-interface management 1
(config-ntp-mgmt-1)#
```

```
Brocade(config-ntp)# ntp-interface ve 100
Brocade(config-ntp-ve-100)#
```

Syntax: **[no] ntp-interface** [**management 1** | **ethernet** *slot/port* | **ve** *id*]

The **management 1** parameter is the management port 1

The **ethernet** *slot/port* parameter specifies the ethernet port number.

The **ve** *id* parameter specifies the virtual port number.

Configuring the broadcast client

To configure a device to receive Network Time Protocol (NTP) broadcast messages on a specified interface, use the **broadcast client** command. NTP broadcast client can be enabled on maximum of 16 ethernet interfaces. If the interface is operationally down or NTP is disabled, then NTP broadcast server packets are not received. To disable this capability, use the **no** form of this command.

```
Brocade(config-ntp mgmt-1)# broadcast client
```

Syntax: [no] broadcast client

Configuring the broadcast destination

To configure the options for broadcasting Network Time Protocol (NTP) traffic, use the **ntp broadcast destination** command. NTP broadcast server can be enabled on maximum 16 ethernet interfaces and four subnet addresses per interface. If the interface is operationally down or there is no ip address configured for the subnet address, then NTP broadcast server packets are not sent. To disable this capability, use the **no** form of this command.

NOTE

This command is not effective, if the NTP server is disabled.

```
Brocade(config)#int m1
Brocade(config-if-mgmt-1)#ip address 10.20.99.173/24
Brocade(config-if-mgmt-1)#ntp
Brocade(config-ntp)#ntp-interface m1
Brocade(config-ntp -mgmt-1)# broadcast destination 10.20.99.0 key 2
```

Syntax: [no] broadcast destination *ip-address* [key *key-id*] [version 3|4]

The *IP-address* parameter is the IPv4 subnet address of the device to send NTP broadcast messages to.

The **key** *key id* option defines the authentication key. By default, no authentication key is configured.

The **version** 3|4 option defines the Network Time Protocol (NTP) version number. If this option is not specified, then it defaults to 4.

Disabling NTP

To disable the NTP server and client, use the **disable** command. Disabling the NTP server or client mode will not remove the configurations. To enable receipt of NTP packets, use the **no** form of this command.

```
Brocade(config-ntp)# disable
```

Syntax: [no] disable [serve]

If the [serve] keyword is specified, then NTP will not serve the time to downstream devices. This keyword disables the NTP server mode functionalities.

If this keyword is not specified, then both NTP client mode and NTP server mode functionalities will be disabled

Configuring the master

To configure the Multi-Service IronWare as a Network Time Protocol (NTP) master clock to which peers synchronize themselves when an external NTP source is not available, use the **master** command. The master clock is disabled by default. To disable the master clock function, use the **no** form of this command.

NOTE

This command is not effective, if the NTP is enabled in client-only mode.

```
Brocade(config-ntp)# master stratum 5
```

Syntax: [no] master [stratum *number*]

Stratum *number* is the number from 2 to 15. It indicates the NTP stratum number that the system will claim.

Enable or disable NTP logging

To enable or disable Network Time Protocol (NTP) message logging, use the **logging enable ntp** command. By default, the logging is enabled for NTP. To disable NTP logging, use the **no** form of this command.

```
Brocade(config)# logging enable ntp
```

Syntax: [no] logging enable ntp

Show commands

Displaying NTP status

Use the **show ntp status** command to display the NTP status

```
Brocade#show ntp status
Clock is synchronized, stratum 4, reference clock is 10.20.99.174
precision is 2**-16
reference time is D281713A.80000000 (03:21:29.3653007907 GMT+00 Thu Dec 01 2011)
clock offset is -2.3307 msec, root delay is 24.6646 msec
root dispersion is 130.3376 msec, peer dispersion is 84.3335 msec
system poll interval is 64, last clock update was 26 sec ago
NTP server mode is enabled, NTP client mode is enabled
NTP master mode is disabled, NTP master stratum is 8
NTP is not in panic mode
```

TABLE 37 show ntp status command output descriptions

Field.....	Description.....
synchronized	Indicates the system clock is synchronized to NTP server or peer.
stratum	Indicates the stratum number that this system is operating. Range 2..15.
reference	IPv4 address or first 32 bits of the MD5 hash of the IPv6 address of the peer to which clock is synchronized.
precision	Precision of the clock of this system in Hz.
reference time	Reference time stamp.
clock offset	Offset of clock (in milliseconds) to synchronized peer.
root delay	Total delay (in milliseconds) along path to root clock.
root dispersion	Dispersion of root path.
peer dispersion	Dispersion of root path.
system poll interval	Poll interval of the local system.
last update	Time the router last updated its NTP information
server mode	Status of the NTP server mode for this device.
client mode	Status of the NTP client mode for this device.
master	Status of the master mode
master stratum	Stratum number that will be used by this device when master is enabled and no upstream time servers are accessible.
panic mode	Status of the panic mode. If the clock offset is more than 1000 seconds with the current time, then panic mode will be on.

Displaying NTP associations

Use the **show ntp associations** command to display detailed association information of the NTP server or peers.

```
Brocade# show ntp associations
address ref clock  st  when poll reach delay  offset  disp
*~172.19.69.1 172.24.114.33 3 25 64 3 2.89 0.234 39377
~2001:db8::234
      INIT 16 - 64 0 0.00 0.000 15937
```

* synced, # selected, + candidate, - outlayer, x falseticker, ~ configured

TABLE 38 show ntp associations command output descriptions

Field.....	Description.....
*	The peer has been declared the system peer and lends its variables to the system variables.
#	This peer is a survivor in the selection algorithm.
+	This peer is a candidate in the combine algorithm.
-	This peer is discarded as outlier in the clustering algorithm.
x	This peer is discarded as falseticker in the selection algorithm.
~	The server or peer is statically configured.
addressl	IPv4 or IPv6 address of the peer.
ref clock	IPv4 address or first 32 bits of the MD5 hash of the IPv6 address of the peer to which clock is synchronized.
St	Stratum setting for the peer.
when	Time, in seconds, since last NTP packet was received from peer.
poll	Polling interval (seconds).
reach	Peer reachability (bit string, in octal).
delay	Round-trip delay to peer, in milliseconds.
offset	Relative time difference between a peer clock and a local clock, in milliseconds.
disp	Dispersion.

Displaying NTP associations details

Use the **show ntp associations detail** command to display all the NTP servers and peers association information.

```
Brocade# show ntp association detail
2001:db8:99:30::1 configured server, sys peer, stratum 3
ref ID 10.235.61.9, time d288dc3b.f2a17891 (10:23:55.4070668433 Pacific Tue Dec 06 2011)
our mode client, peer mode server, our poll intvl 10, peer poll intvl 10,
root delay 0.08551025 msec, root disp 0.09309387, reach 17, root dist 0.17668502
delay 0.69961487 msec, offset -13.49459670 msec, dispersion 17.31550718,
precision 2**-16, version 4
org time d288df70.a91de561 (10:37:36.2837308769 Pacific Tue Dec 06 2011)
rcv time d288df70.a0c8d19e (10:37:36.2697515422 Pacific Tue Dec 06 2011)
xmt time d288df70.a086e4de (10:37:36.2693194974 Pacific Tue Dec 06 2011)
filter delay 1.7736 0.9933 0.8873 0.6699 0.7709 0.7712 0.7734 6.7741
filter offset -17.9936 33.0014 -13.6604 -13.4494 -14.4481 -16.4453 -18.4423 -22.0025
filter disp 15.6660 0.0030 17.7730 17.7700 17.6670 17.6640 17.6610 16.6635
filter epoch 55824 56866 55686 55688 55690 55692 55694 55759
```

Use the **show ntp associations detail IPv4 address | IPv6 address** command to display the NTP servers and peers association information for a specific ip address.

```
Brocade# show ntp association detail 10.99.40.1
10.99.40.1 configured server, candidate, stratum 3
```

```

ref ID 10.45.57.38, time d288de7d.690ca5c7 (10:33:33.1762436551 Pacific Tue Dec 06 2011)
our mode client, peer mode server, our poll intvl 10, peer poll intvl 10,
root delay 0.02618408 msec, root disp 0.10108947, reach 3, root dist 0.23610585
delay 0.92163588 msec, offset 60.77749188 msec, dispersion 70.33842156,
precision 2**-16, version 4
org time d288defa.b260a71f (10:35:38.2992678687 Pacific Tue Dec 06 2011)
rcv time d288defa.a2efbd41 (10:35:38.2733620545 Pacific Tue Dec 06 2011)
xmt time d288defa.a2ae54f8 (10:35:38.2729334008 Pacific Tue Dec 06 2011)
filter delay      0.000      6.7770      6.7773      6.7711      6.7720      6.7736      6.7700      0.9921
filter offset     0.000     19.0047     19.1145     19.2245     19.3313     17.4410     15.4463     60.7777
filter disp      16000.000    16.0005    15.9975    15.9945    15.9915    15.8885    15.8855     0.0030
filter epoch      55683      55683      55685      55687      55689      55691      55693      56748

```

Syntax: `show ntp association detail IPv4 address | IPv6 address`

The IPv4 or IPv6 address of the NTP peer

TABLE 39 `show ntp associations detail` command output descriptions

Field.....	Description.....
server	Indicates server is statically configured.
symmetric active peer	Indicates peer is statically configured.
symmetric passive peer	Indicates peer is dynamically configured.
sys_peer	This peer is the system peer
candidate	This peer is chosen as candidate in the combine algorithm.
reject	This peer is rejected by the selection algorithm
falsetick	This peer is dropped as falseticker by the selection algorithm
outlyer	This peer is dropped as outlyer by the clustering algorithm
Stratum	Stratum number
ref ID	IPv4 address or hash of IPv6 address of the upstream time server to which the peer is synchronized.
Time	Last time stamp that the peer received from its master.
our mode	This system's mode relative to peer (active /passive /client /server /bdcast /bdcast client).
peer mode	Mode of peer relative to this system
our poll intvl	This system's poll interval to this peer.
peer poll intvl	Poll interval of peer to this system
root delay	The delay along path to root (the final stratum 1 time source).
root disp	Dispersion of path to root.
root distance	The distance from the server or peer to the client
reach peer	The peer reachability (bit string in octal).
Delay	Round-trip delay to peer.
offset	Offset of peer clock relative to this clock
Dispersion	Dispersion of peer clock
precision	Precision of peer clock

TABLE 39 show ntp associations detail command output descriptions

Field.....	Description.....
version	Peer NTP version number
org time	Originate time stamp of the last packet.
rcv time	Receive time stamp of the last packet.
xmt time	Transmit time stamp of the last packet.
filter delay	Round-trip delay in milliseconds of last 8 samples.
filter offset	Clock offset in milliseconds of last 8 samples.
filter error	Approximate error of last 8 samples

Configuration Examples

The following sections lists configuration examples to configure the Brocade device.

NTP server and client mode configuration

Sample CLI commands to configure the NI device in NTP server and client modes.

```
Brocade(config-ntp)# server 10.1.2.3 minpoll 5 maxpoll 10
Brocade(config-ntp)# server 2001:db8::1/64
Brocade(config-ntp)# peer 10.100.12.18
Brocade(config-ntp)# peer 10.100.12.20
Brocade(config-ntp)# peer 10.100.12.67
Brocade(config-ntp)# peer 10.100.12.83
```

NTP client mode configuration

Sample CLI commands to configure the Brocade device in NTP client mode.

```
Brocade(config-ntp)# server 10.1.2.3 minpoll 5 maxpoll 10
Brocade(config-ntp)# server 2001:db8::1/24
Brocade(config-ntp)# peer 10.100.12.83
Brocade(config-ntp)# disable serve
```

NTP strict authentication configuration

Sample CLI commands to configure the Brocade device in strict authentication mode.

```
Brocade(config-ntp)# authenticate
Brocade(config-ntp)# authentication-key key-id1 md5 key123
Brocade(config-ntp)# server 10.1.2.4 key 1
```

NTP loose authentication configuration

Sample CLI commands to configure the NI device in loose authentication mode. This allows some of the servers or clients to use the authentication keys.

```
Brocade(config-ntp)# authentication-key-id key-id1 md5 key123
Brocade(config-ntp)# server 10.1.2.4 key 1
Brocade(config-ntp)# server 10.1.2.7
```

NTP interface context for broadcast server or client mode

Sample CLI command enter the NTP interface context.

```
Brocade(config)#int management 1
Brocade(config-if-mgmt-1)#ip address 10.20.99.173/24
Brocade(config-if-mgmt-1)#ntp
Brocade(config-ntp)# ntp-interface management 1
Brocade(config-ntp-mgmt-1)# broadcast destination 10.23.45.128

Brocade(config-ntp)# ntp-interface ethernet 1/3
Brocade(config-ntp-if-e1000-1/3)# broadcast destination 10.1.1.0 key 1

Brocade(config-ntp)# ntp-interface ve 100
Brocade(config-ntp-ve-100)# broadcast destination 10.2.2.0 key 23
```

NTP broadcast client configuration

Sample CLI commands to configure the NTP broadcast client

```
Brocade(config-ntp)# ntp-interface management 1
Brocade(config-ntp-mgmt-1)# broadcast client

Brocade(config-ntp)# ntp-interface ethernet 1/5
Brocade(config-ntp-if-e1000-1/5)# broadcast client

Brocade(config-ntp)# ntp-interface ve 100
Brocade(config-ntp-ve-100)# broadcast client
```

7 Show commands

Network Configuration Protocol

Table 40 displays the individual Brocade devices and the NETCONF features they support.

TABLE 40 Supported Brocade NETCONF features

Features supported	Brocade NetIron XMR Series	Brocade MLX Series	Brocade NetIron CES 2000 Series BASE package	Brocade NetIron CES 2000 Series ME_PREM package	Brocade NetIron CES 2000 Series L3_PREM package	Brocade NetIron CER 2000 Series Base package	Brocade NetIron CER 2000 Series Advanced Services package
NETCONF protocol	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Data models and mappings	Yes	Yes	Yes	Yes	Yes	Yes	Yes

NETCONF protocol introduction

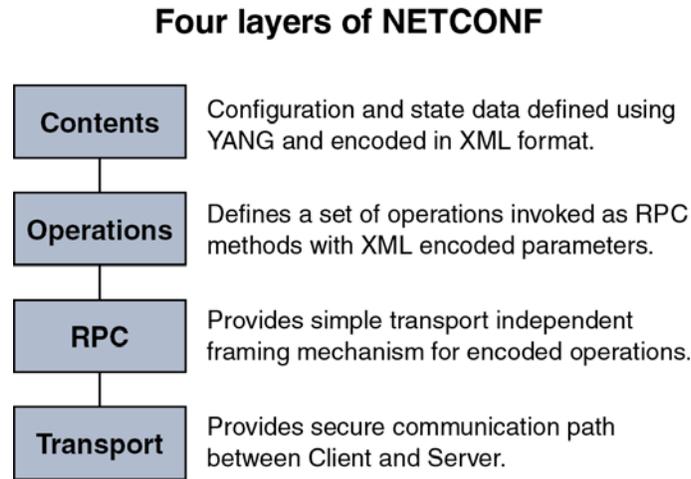
The Network Configuration protocol (NETCONF) uses Extensible Markup Language (XML) for automated configuration management. The NETCONF protocol runs on top of a secure transport, such as Secure Shell version 2 (SSHv2). Only one NETCONF session is supported at a time and any new NETCONF connection requests are rejected after the first session has been established.

NETCONF provides mechanisms through which you can do the following:

- Manage multiple network devices
- Retrieve full or partial configuration and state data
- Upload and manipulate new configurations

NETCONF can be conceptually partitioned into four layers, as shown in [Figure 16](#).

FIGURE 16 Four layers of NETCONF



Platforms

NETCONF is supported on the Brocade MLX series, Brocade NetIron XMR, Brocade NetIron CER, and Brocade NetIron CES devices.

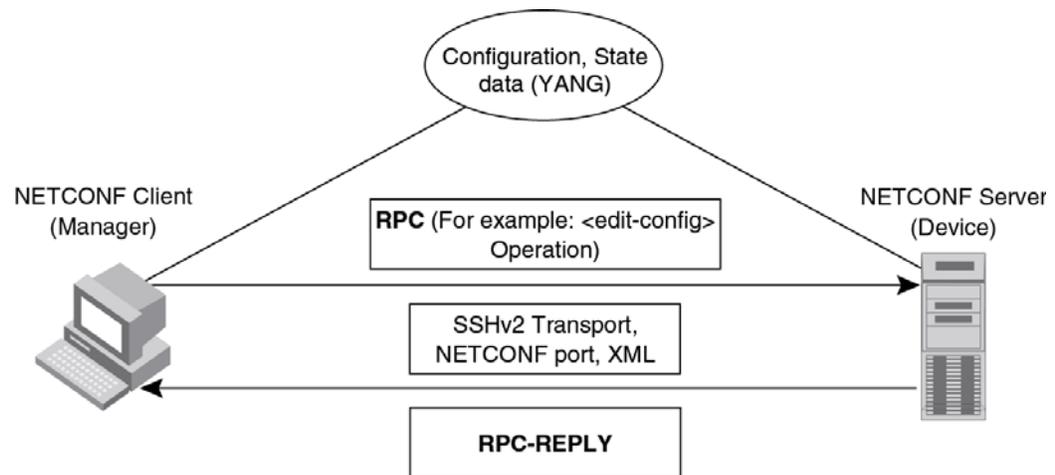
Related documentation

For detailed information about NETCONF, refer to RFC 4741.

For detailed information about using the NETCONF protocol over the Secure Shell (SSH), refer to RFC 4742.

NETCONF in client/server architecture

The NETCONF protocol uses a Remote Procedure Call (RPC) paradigm to facilitate communication between the client (NETCONF Manager or application) and the server (NETCONF Agent or device). A client encodes an RPC in XML and sends it to a server using a secure, connection-oriented session. The server responds with a reply encoded in XML, as shown in [Figure 17](#).

FIGURE 17 NETCONF communication

The communication between the client and server consists of a series of alternating request and reply messages. The NETCONF peers use `<rpc>` and `<rpc-reply>` elements to provide transport protocol-independent framing of NETCONF requests and responses. The NETCONF server processes the RPC requests sequentially in the order in which they are received.

RPC request

The `<rpc>` element is used for enclosing a NETCONF request sent from the client to the server. Every `<rpc>` element contains a mandatory attribute, the `message-id`. This attribute has a unique value for every RPC request, and is used to associate every RPC request with the corresponding response. The `message-id` value is a monotonically increasing integer string. The maximum length of the string is 4095 characters. If the `message-id` is not present in the RPC request, the server rejects the request by returning an `<rpc-error>` with the `error-tag` element set to the `missing-attribute`.

If there are any additional attributes present in the RPC request, the NETCONF server returns them unmodified in the corresponding RPC reply.

RPC reply

An `<rpc-reply>` element is sent in response to every RPC request. The `<rpc-reply>` element contains the mandatory attribute `message-id` copied from the corresponding RPC request, along with any additional attributes that are present in the RPC request.

For successfully processed `get` or `get-config` requests, the response data is encoded as the content of the `<rpc-reply>` element.

For successfully processed `edit-config` or `close-session` requests, the `<ok>` element is encoded as the content of the `<rpc-reply>` element.

For unsuccessful RPC requests, one or more `<rpc-error>` elements are encoded inside the `<rpc-reply>` element.

RPC and error handling

If the RPC request fails, an `<rpc-error>` element, the first detected error, is encoded inside the `<rpc-reply>` element and sent to the client. The server is not required to detect or report multiple errors. If the server detects multiple errors then the order of the error detection and reporting is at the discretion of the server.

CLI and SSH subsystem

The NETCONF client must use Secure Shell Version 2 (SSHv2) as the network transport to connect to the NETCONF server. Only the SSHv2 protocol is supported as the NETCONF transport protocol.

To run NETCONF over SSHv2, the client establishes an SSH transport connection using the SSH transport protocol to the NETCONF port. The default NETCONF port is 830. The underlying SSH client and server exchange keys for message integrity and encryption.

The SSHv2 client invokes the **ssh-userauth** service to authenticate the user. All currently supported SSH user authentication methods such as the **public-key**, **password**, and **keyboard-interactive** authentications are supported for a NETCONF session also. If the SSH user authentication is disabled, the user is allowed full access.

On successful user authentication, the client invokes the **ssh-connection** service, also known as the SSH connection protocol. After the SSH session is established, the NETCONF client invokes NETCONF as an SSH subsystem called **netconf**.

NETCONF user privileges

Every NETCONF session has a corresponding authentication, authorization, and accounting (AAA) session. The AAA attributes apply to the NETCONF session. Only authentication and EXEC authorization are supported. Other forms of accounting and command authorization are not supported.

The privilege level of the user (read-only(5), read-write(0)) is obtained from the AAA server, if it is provided. If the privilege level is not provided by the AAA server, the default privilege level applies for the NETCONF session.

Table 41 provides the mapping between the NETCONF privilege levels and the AAA privilege levels.

AAA privilege level	NETCONF privilege level
0	NETCONF_PRIVILEGE_LEVEL_0
1-5	NETCONF_PRIVILEGE_LEVEL_5

Table 42 provides the mapping between the NETCONF privilege levels and the supported NETCONF operations.

Operations	NETCONF_PRIVILEGE_LEVEL_0	NETCONF_PRIVILEGE_LEVEL_5
<get>	Yes	Yes
<get-config>	Yes	Yes

TABLE 42 NETCONF operations and privilege levels

Operations	NETCONF_PRIVILEGE_LEVEL_0	NETCONF_PRIVILEGE_LEVEL_5
<edit-config>	Yes	No
<close-session>	Yes	Yes

Recommendations for NETCONF

- Use an authentication method to secure the underlying SSH session and to prevent any unauthorized access.
- Use a NETCONF client to generate the RPCs. If you have manually written the XML requests, recycling the XML from a successful request is recommended.
- Refer to the *Brocade MLX Series and Brocade NetIron Family YANG guide* for XML and data verification.
- Plan the configuration or state information that must be sent or retrieved to avoid sending and receiving large RPC messages.
- Use of a scripting language or other custom interface is recommended.

Basic NETCONF operations

The NETCONF protocol provides a small set of low-level operations to manage device configurations and retrieve device state information. The base protocol provides operations to retrieve, configure, copy, and delete configuration data stores. Additional operations are provided based on the capabilities advertised by the device.

The following base protocol operations are supported:

- `get`
- `get-config`
- `edit-config`
- `close-session`

NOTE

Other operations, including `copy-config`, `delete-config`, `lock`, `unlock`, and `kill-session` are not supported.

Initial connection

Each NETCONF session begins with a handshake in which the NETCONF server and the client specify the NETCONF capabilities they support. The following sections describe how to start a NETCONF session.

Hello messages

After establishing a secure transport connection, both the NETCONF server and client send a `<hello>` element simultaneously to announce their capabilities and session identifier.

After sending the hello message, the server starts the hello timer (default is 600 seconds) and waits for the hello message from the client. If no hello message is received by the server before the hello timer expires, the server aborts the NETCONF session by closing the underlying SSH session.

The NETCONF server must include the `<session-id>` element, which contains the unique session value for the NETCONF session, in the `<hello>` element. If the client receives the `<hello>` element without the `<session-id>`, the client aborts the NETCONF session by closing the underlying SSH session.

The NETCONF client must not include the `<session-id>` element in the `<hello>` element. If the server receives the `<hello>` element with the `<session-id>`, the server aborts the NETCONF session by closing the underlying SSH session.

The NETCONF client must include a valid `xmlns` attribute in the `<hello>` element. If the server receives the `<hello>` element without a valid `xmlns` attribute, the server aborts the NETCONF session by closing the underlying SSH session.

The NETCONF client must include a base capability. The server receiving the `<hello>` element without a NETCONF base capability aborts the NETCONF session by closing the underlying SSH session.

The server receiving the `<rpc>` element without receiving the `<hello>` element aborts the NETCONF session by closing the underlying SSH session.

The NETCONF client may send arbitrary data before sending a valid hello message. The server discards the data until a valid `<hello>` element is received from the client.

The following is an example for a `<hello>` element from the NETCONF server.

```
<hello xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <capabilities>
    <capability>
      urn:ietf:params:netconf:base:1.0
    </capability>
    <capability>
      urn:ietf:params:netconf:capability:writable-running:1.0
    </capability>
  </capabilities>
  <session-id>4</session-id>
</hello>
```

Capabilities

A NETCONF capability is a set of protocol extensions that supplements the base NETCONF specification. A NETCONF capability is identified with a Uniform Resource Identifier (URI). Capabilities augment the base operations of the NETCONF server, describing both the additional operations and the contents allowed inside the operations. To support a capability, the NETCONF server must support all the dependent capabilities.

The following capabilities are supported on the NetIron platforms:

- **Base capability:** The base capability is the set of operations and contents that any NETCONF implementation must support. The URI for the base capability is `urn:ietf:params:xml:ns:netconf:base:1.0`. Both the NETCONF client and server must support the base capability.
- **Writable-running capability:** The writable-running capability indicates that the device supports `edit-config` and `copy-config` operations where the `<running>` configuration is the target. The URI is `urn:ietf:params:netconf:capability:writable-running:1.0`.

NOTE

Other capabilities, including Candidate Configuration Capability, Confirmed Commit Capability, and Validate Capability, are not supported.

<get> operation

The NETCONF <get> operation retrieves the devices and the state data, or a filtered subset of the data.

If the device can satisfy the request, the server sends an <rpc-reply> element containing a <data> element with the results of the query. If the request cannot be completed, an <rpc-error> element is included in the <rpc-reply> element.

Parameter

The <get> operation uses the `filter` parameter. The `filter` parameter specifies the portion of the system data to retrieve. If this parameter is not present, show version information is returned.

Examples

The following is an example of a <get> operation:

```
<nc:rpc message-id="1" xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0">
<nc:get>
</nc:get>
</nc:rpc>
]]>]]>
```

The following is an example of a <get> operation for MPLS state data.

```
<nc:rpc xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/"
message-id="25">
<nc:get>
<nc:filter >
<brcd:netiron-statedata>
<brcd:mpls-statedata/>
</brcd:netiron-statedata>
</nc:filter>
</nc:get>
</nc:rpc>
]]>]]>
```

The following is an example of a <get> operation for a specific LSP.

```
<nc:rpc xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/"
message-id="25">
<nc:get>
<nc:filter nc:type="subtree" >
<brcd:netiron-statedata>
<brcd:mpls-statedata>
<brcd:mpls-lsp-statedata>
<brcd:name>scriptlsp1001</brcd:name>
</brcd:mpls-lsp-statedata>
</brcd:mpls-statedata>
</nc:filter>
</nc:get>
</nc:rpc>
]]>]]>
```

```

</brcd:netiron-statedata>
</nc:filter>
</nc:get>
</nc:rpc>
]]>]]>
<nc:rpc-reply xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/"
message-id="25">
<nc:data>
<netiron-statedata xmlns="http://brocade.com/ns/netconf/config/netiron-config/">
<brcd:mpls-statedata>
<brcd:mpls-lsp-statedata>
  <brcd:name>scriptlsp1001</brcd:name>
  <brcd:to>10.0.0.1</brcd:to>
  <brcd:admin-state>
    <brcd:up></brcd:up>
  </brcd:admin-state>
  <brcd:oper-state>
    <brcd:down></brcd:down>
  </brcd:oper-state>
  <brcd:tunnel-intf>tnl0</brcd:tunnel-intf>
  <brcd:up-dn-times>0</brcd:up-dn-times>
  <brcd:retry-no>273</brcd:retry-no>
</brcd:mpls-lsp-statedata>
</brcd:mpls-statedata>
</netiron-statedata>
</nc:data>
</nc:rpc-reply>
]]>]]>

```

The following is an example of a <get> operation for VLAN state data.

```

<nc:rpc xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/"
message-id="25">
<nc:get>
<nc:filter >
<brcd:netiron-statedata>
<brcd:vlan-statedata/>
</brcd:netiron-statedata>
</nc:filter>
</nc:get>
</nc:rpc>
]]>]]>

```

The following is an example of a <get> operation for VLAN 1001.

```

<nc:rpc xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/"
message-id="25">
<nc:get>
<nc:filter nc:type="subtree" >
<brcd:netiron-statedata>
<brcd:vlan-statedata>
<brcd:vlan>
<brcd:vlan-id>1001</brcd:vlan-id>
</brcd:vlan>
</brcd:vlan-statedata>
</brcd:netiron-statedata>

```

```

</nc:filter>
</nc:get>
</nc:rpc>
]]>]]>
<nc:rpc-reply xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/"
message-id="25">
<nc:data>
<netiron-statedata xmlns="http://brocade.com/ns/netconf/config/netiron-config/">
<brcd:vlan-statedata>
<brcd:vlan>
  <brcd:vlan-id>1001</brcd:vlan-id>
  <brcd:vlan-name>scriptVlan1001</brcd:vlan-name>
  <brcd:topo-hw-idx>65535</brcd:topo-hw-idx>
  <brcd:topo-sw-idx>257</brcd:topo-sw-idx>
  <brcd:topo-next-vlan>0</brcd:topo-next-vlan>
  <brcd:port>
    <brcd:port-id>ethernet 1/20 </brcd:port-id>
    <brcd:tag-mode>TAGGED</brcd:tag-mode>
    <brcd:state>DISABLED</brcd:state>
  </brcd:port>
  <brcd:bytes-received>0</brcd:bytes-received>
</brcd:vlan>
<brcd:dual-mode>ethernet 1/1 to 1/20 </brcd:dual-mode>
<brcd:dual-mode>ethernet 2/1 to 2/8 </brcd:dual-mode>
<brcd:dual-mode>ethernet 5/1 to 5/8 </brcd:dual-mode>
<brcd:dual-mode>ethernet 7/1 to 7/8 </brcd:dual-mode>
<brcd:dual-mode>ethernet 8/1 to 8/48 </brcd:dual-mode>
<brcd:dual-mode>ethernet 9/1 to 9/8 </brcd:dual-mode>
<brcd:dual-mode>ethernet 10/1 to 10/8 </brcd:dual-mode>
<brcd:dual-mode>ethernet 11/1 to 11/8 </brcd:dual-mode>
<brcd:dual-mode>ethernet 12/1 to 12/8 </brcd:dual-mode>
<brcd:dual-mode>ethernet 13/1 to 13/8 </brcd:dual-mode>
<brcd:dual-mode>ethernet 14/1 to 14/8 </brcd:dual-mode>
<brcd:dual-mode>ethernet 15/1 to 15/2 </brcd:dual-mode>
<brcd:default-vlan-id>1</brcd:default-vlan-id>
<brcd:control-vlan-id>4095</brcd:control-vlan-id>
<brcd:maximum-port-vlan-entries>4095</brcd:maximum-port-vlan-entries>
</brcd:vlan-statedata>
</netiron-statedata>
</nc:data>
</nc:rpc-reply>
]]>]]>

```

The following is an example of a <get> operation for Interface state data.

```

<nc:rpc xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/"
message-id="25">
<nc:get>
<nc:filter >
<brcd:netiron-statedata>
<brcd:interface-statedata/>
</brcd:netiron-statedata>
</nc:filter>
</nc:get>
</nc:rpc>
]]>]]>

```

The following is an example of a <get> operation for a specific Interface state data.

```

<nc:rpc xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/"
message-id="25">
<nc:get>
<nc:filter nc:type="subtree" >
<brcd:netiron-statedata>
<brcd:interface-statedata>
<brcd:interface>
<brcd:interface-id>ethernet 1/20</brcd:interface-id>
</brcd:interface>
</brcd:interface-statedata>
</brcd:netiron-statedata>
</nc:filter>
</nc:get>
</nc:rpc>
]]>]]>
<nc:rpc-reply xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/"
message-id="25">
<nc:data>
<netiron-statedata xmlns="http://brocade.com/ns/netconf/config/netiron-config/">
<brcd:interface-statedata>
<brcd:interface>
  <brcd:interface-id>ethernet 1/20</brcd:interface-id>
  <brcd:link-state>
    <brcd:down></brcd:down>
  </brcd:link-state>
  <brcd:l2-state>
    <brcd:disabled></brcd:disabled>
  </brcd:l2-state>
  <brcd:duplex>
    <brcd:none></brcd:none>
  </brcd:duplex>
  <brcd:speed></brcd:speed>
  <brcd:tag-mode>
    <brcd:yes></brcd:yes>
  </brcd:tag-mode>
  <brcd:priority-level>
    <brcd:level0></brcd:level0>
  </brcd:priority-level>
  <brcd:mac-address>0000.0085.2d00</brcd:mac-address>
</brcd:interface>
</brcd:interface-statedata>
</netiron-statedata>
</nc:data>
</nc:rpc-reply>
]]>]]>

```

<get-config> operation

The NETCONF <get-config> operation retrieves all or part of a configuration from the source data store. The <get-config> operation is similar to the **show running-config** command.

If the device can satisfy the request, the server sends an <rpc-reply> element containing a <data> element with the results of the query. If the request cannot be completed, an <rpc-error> element is included in the <rpc-reply> element.

Parameters

The parameters used for `<get-config>` are as follows:

- `source`: Name of the configuration data store being queried, such as `<running/>`. Only running configuration data store is supported.
- `filter`: Specifies the portions of the device configuration to retrieve. If this parameter is not present, no configuration is returned. The `filter` parameter must contain a `type` attribute. This attribute indicates the type of filtering syntax used within the `filter` parameter. The subtree filtering is the default filtering mechanism used in NETCONF.

NOTE

xpath filtering is not supported.

Examples

The following is an example of a `<get-config>` operation for MPLS configuration.

```
<nc:rpc xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/"
message-id="1">
<nc:get-config>
<nc:source>
<nc:running/>
</nc:source>
<nc:filter nc:type="subtree">
<brcd:netiron-config>
<brcd:mpls-config/>
</brcd:netiron-config>
</nc:filter>
</nc:get-config>
</nc:rpc>
]]>]]>
<nc:rpc-reply xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/"
message-id="1">
<nc:data>
<brcd:netiron-config>
<brcd:mpls-config>
<brcd:path>
  <brcd:name>example</brcd:name>
  <brcd:strict>10.99.161.1</brcd:strict>
</brcd:path>
<brcd:path>
  <brcd:name>example2</brcd:name>
  <brcd:strict>10.99.145.1</brcd:strict>
</brcd:path>
<brcd:lsp>
  <brcd:name>examplelsp1</brcd:name>
  <brcd:adaptive></brcd:adaptive>
  <brcd:from>10.99.10.1</brcd:from>
  <brcd:to>10.99.161.1</brcd:to>
  <brcd:enable></brcd:enable>
  <brcd:hop-limit>10</brcd:hop-limit>
  <brcd:ipmtu>1526</brcd:ipmtu>
  <brcd:ldp-tunneling></brcd:ldp-tunneling>
  <brcd:metric>600</brcd:metric>
  <brcd:primary-path>example</brcd:primary-path>
```

```

    <brcd:record></brcd:record>
    <brcd:reoptimize-timer>3600</brcd:reoptimize-timer>
    <brcd:revert-timer>43200</brcd:revert-timer>
    <brcd:mpls-traffic-eng>
      <brcd:max-burst>44736</brcd:max-burst>
      <brcd:max-rate>6312</brcd:max-rate>
      <brcd:mean-rate>1544</brcd:mean-rate>
    </brcd:mpls-traffic-eng>
    <brcd:secondary-path>
      <brcd:name>example2</brcd:name>
    </brcd:secondary-path>
  </brcd:lsp>
</brcd:lsp>
  <brcd:name>examplelsp2</brcd:name>
  <brcd:adaptive></brcd:adaptive>
  <brcd:from>10.99.10.1</brcd:from>
  <brcd:to>10.99.145.1</brcd:to>
  <brcd:enable></brcd:enable>
  <brcd:hop-limit>10</brcd:hop-limit>
  <brcd:ipmtu>1526</brcd:ipmtu>
  <brcd:ldp-tunneling></brcd:ldp-tunneling>
  <brcd:metric>600</brcd:metric>
  <brcd:primary-path>example2</brcd:primary-path>
  <brcd:record></brcd:record>
  <brcd:reoptimize-timer>3600</brcd:reoptimize-timer>
  <brcd:revert-timer>43200</brcd:revert-timer>
  <brcd:mpls-traffic-eng>
    <brcd:max-burst>44736</brcd:max-burst>
    <brcd:max-rate>6312</brcd:max-rate>
    <brcd:mean-rate>1544</brcd:mean-rate>
  </brcd:mpls-traffic-eng>
  <brcd:secondary-path>
    <brcd:name>example</brcd:name>
  </brcd:secondary-path>
</brcd:lsp>
</brcd:lsp>
  <brcd:name>exmaplelsp</brcd:name>
  <brcd:from>10.99.10.1</brcd:from>
  <brcd:to>10.99.161.1</brcd:to>
  <brcd:disable></brcd:disable>
  <brcd:record></brcd:record>
</brcd:lsp>
</brcd:mpls-config>
</brcd:netiron-config>
</nc:data>
</nc:rpc-reply>
]]>]]>

```

The following is an example of a `<get-config>` operation for VLAN configuration.

```

<nc:rpc xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/"
message-id="1">
  <nc:get-config>
  <nc:source>
  <nc:running/>
  </nc:source>
  <nc:filter nc:type="subtree">
  <brcd:netiron-config>
  <brcd:vlan-config/>

```

```

</brcd:netiron-config>
</nc:filter>
</nc:get-config>
</nc:rpc>
]]>]]>
<nc:rpc-reply xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/"
message-id="1">
<nc:data>
<brcd:netiron-config>
<brcd:vlan-config>
<brcd:vlan>
  <brcd:vlan-id>1</brcd:vlan-id>
  <brcd:vlan-name>DEFAULT-VLAN</brcd:vlan-name>
  <brcd:untagged>ethernet 1/1 </brcd:untagged>
  <brcd:untagged>ethernet 1/3 to 1/24 </brcd:untagged>
</brcd:vlan>
<brcd:vlan>
  <brcd:vlan-id>100</brcd:vlan-id>
  <brcd:vlan-name></brcd:vlan-name>
  <brcd:untagged>ethernet 2/1 to 2/2 </brcd:untagged>
  <brcd:router-interface>ve 100</brcd:router-interface>
</brcd:vlan>
<brcd:vlan>
  <brcd:vlan-id>200</brcd:vlan-id>
  <brcd:vlan-name></brcd:vlan-name>
  <brcd:tagged>ethernet 1/1 </brcd:tagged>
  <brcd:tagged>ethernet 1/3 </brcd:tagged>
  <brcd:tagged>ethernet 1/5 </brcd:tagged>
</brcd:vlan>
<brcd:vlan>
  <brcd:vlan-id>300</brcd:vlan-id>
  <brcd:vlan-name></brcd:vlan-name>
  <brcd:untagged>ethernet 1/2 </brcd:untagged>
</brcd:vlan>
<brcd:vlan>
  <brcd:vlan-id>4095</brcd:vlan-id>
  <brcd:vlan-name>CONTROL-VLAN</brcd:vlan-name>
  <brcd:tagged>ethernet 1/1 to 1/24 </brcd:tagged>
  <brcd:tagged>ethernet 2/1 to 2/2 </brcd:tagged>
</brcd:vlan>
<brcd:default-vlan-id>1</brcd:default-vlan-id>
</brcd:vlan-config>
</brcd:netiron-config>
</nc:data>
</nc:rpc-reply>
]]>]]>

```

The following is an example of a `<get-config>` operation for Interface configuration.

```

<nc:rpc xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/"
message-id="1">
<nc:get-config>
<nc:source>
<nc:running/>
</nc:source>
<nc:filter nc:type="subtree">
<brcd:netiron-config>
<brcd:interface-config/>

```

```

</brcd:netiron-config>
</nc:filter>
</nc:get-config>
</nc:rpc>
]]>]]>
<nc:rpc-reply xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/"
message-id="1">
<nc:data>
<brcd:netiron-config>
<brcd:interface-config>
<brcd:interface>
  <brcd:interface-id>management 1</brcd:interface-id>
  <brcd:enable></brcd:enable>
  <brcd:ip>
    <brcd:address>10.20.99.187/20</brcd:address>
  </brcd:ip>
  <brcd:ipv6>
    <brcd:address>2001:db8::10:20:99:187/64</brcd:address>
  </brcd:ipv6>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/1</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/2</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/3</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/4</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>

```

```

    <brcd:priority>
  </brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/5</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/6</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/7</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/8</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/9</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/10</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>

```

```

    <brcd:priority>
  </brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/11</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/12</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/13</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/14</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/15</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/16</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>

```

```
<brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/17</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/18</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/19</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/20</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/21</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/22</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
```

```

    <brcd:priority>
  </brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/23</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/24</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 2/1</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 2/2</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ve 100</brcd:interface-id>
  <brcd:enable></brcd:enable>
  <brcd:ip>
    <brcd:address>10.2.2.2/24</brcd:address>
  </brcd:ip>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>loopback 1</brcd:interface-id>
  <brcd:enable></brcd:enable>
  <brcd:ip>
    <brcd:address>10.13.32.1/32</brcd:address>
  </brcd:ip>
</brcd:interface>
</brcd:interface-config>
</brcd:netiron-config>
</nc:data>
</nc:rpc-reply>
]]>]]>

```

<edit-config> operation

The NETCONF <edit-config> operation loads all the configurations into the specified target configuration.

Elements in the <config> subtree may contain an `operation` attribute. The attribute identifies the point in the configuration to perform the operation and might appear on multiple elements throughout the <config> subtree.

The operation attribute contains any one of the following values: `merge`, `replace`, `create`, `delete`.

The values `merge`, `replace`, or `create` is enforced by the behavior of the individual CLI, so these options are ignored. The `delete` operation alone is supported.

Parameters

The parameters used for <edit-config> are as follows:

- `target`: Name of the configuration data store being edited, such as `<nc:running/>`.
- `test-option`: This option is not supported.
- `default-operation`: Only the `none` value is supported. The other values such as `merge`, `replace`, `create`, and `delete` are ignored because the behaviors are enforced by the individual CLI.
- `error-option`: Only the `stop-on-error` option is supported. The other values such as `continue-on-error` and `rollback-on-error` are ignored.
- `config`: A hierarchy of configuration data as defined by the data models of the device. The new configuration must be inline configuration and other configuration options such as local file, remote file, and URL are not supported.

Examples

The following is an example for an <edit-config> operation for MPLS configuration.

```
<nc:rpc message-id="1" xmlns:nc="urn:iETF:params:xml:ns:netconf:base:1.0"
xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/">
<nc:edit-config>
<nc:target>
<nc:running/>
</nc:target>
<nc:default-operation>merge</nc:default-operation>
<nc:config>
<brcd:netiron-config>
<brcd:mpls-config>
<brcd:lsp nc:operation="delete">
<brcd:name>examplelsp2</brcd:name>
</brcd:lsp>
</brcd:mpls-config>
</brcd:netiron-config>
</nc:config>
</nc:edit-config>
</nc:rpc>
]]>]]>
<nc:rpc-reply xmlns:nc="urn:iETF:params:xml:ns:netconf:base:1.0"
xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/"
message-id="1">
<nc:ok></nc:ok>
```

```
</nc:rpc-reply>
]]>]]>
```

The following is an example for an `<edit-config>` operation for VLAN configuration.

```
<nc:rpc message-id="1" xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/">
<nc:edit-config>
<nc:target>
<nc:running/>
</nc:target>
<nc:default-operation>merge</nc:default-operation>
<nc:config>
<brcd:netiron-config>
<brcd:vlan-config>
<brcd:vlan nc:operation="delete">
<brcd:vlan-id>200</brcd:vlan-id>
</brcd:vlan>
</brcd:vlan-config>
</brcd:netiron-config>
</nc:config>
</nc:edit-config>
</nc:rpc>
]]>]]>
<nc:rpc-reply xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/"
message-id="1">
<nc:ok></nc:ok>
</nc:rpc-reply>
]]>]]>
```

The following is an example for an `<edit-config>` operation to configure interface ethernet 1/1 with the IP address of 10.1.1.1/24 and enable the interface.

```
<nc:rpc message-id="1" xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/">
<nc:edit-config>
<nc:target>
<nc:running/>
</nc:target>
<nc:config>
<brcd:netiron-config>
<brcd:interface-config>
<brcd:interface>
<brcd:interface-id>ethernet 1/1</brcd:interface-id>
<brcd:enable></brcd:enable>
<brcd:ip>
<brcd:address>10.1.1.1/24</brcd:address>
</brcd:ip>
</brcd:interface>
</brcd:interface-config>
</brcd:netiron-config>
</nc:config>
</nc:edit-config>
</nc:rpc>
]]>]]>
<nc:rpc-reply xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/"
message-id="1">
<nc:ok></nc:ok>
</nc:rpc-reply>
```

```
]]>]]>
```

Error handling

The `error-option` element contains the `stop-on-error` value. The `stop-on-error` value aborts the `edit-config` operation on the first error. All the configuration items before the error are already applied on the system. This is the default error option.

After receiving the complete `edit-config` RPC, the configuration items specified in the XML are applied sequentially as per the order specified in the YANG. If all the configuration items are successfully applied, an `<ok>` element is sent in the `<rpc-reply>` element. Otherwise, an `<rpc-error>` element with the details of the error is sent in the `<rpc-reply>` element.

Closing sessions

The NETCONF `<close-session>` operation is used for gracefully closing the current NETCONF session. The `<close-session>` operation uses no additional parameters.

When a NETCONF server receives a `<close-session>` request, the server releases resources associated with the session and closes the underlying SSH connection. Any NETCONF requests received after a `<close-session>` request are ignored. If the device is able to close the connection, an `<rpc-reply>` element is sent that includes an `<ok>` element. Otherwise, an `<rpc-error>` element with the details of the error is sent in the `<rpc-reply>` element.

NETCONF commands and specifications

The following sections describe the configuration of NETCONF using the CLI and the associated show commands, the syslog messages, and the system limitations of NETCONF.

Configuring NETCONF server

To enable the NETCONF server on a device, enter the following command.

```
Brocade(config)# netconf server
Brocade# netconf server?
server    Enable NETCONF server functionality
```

When no port number is specified, the command applies to the default port (830).

To enable the NETCONF server for a specific port, enter the following command.

```
Brocade(config)# netconf server port 2001
```

Syntax: `[no] netconf server [port port-number]`

The **port** option allows you to enable NETCONF on a non-default port.

The *port-number* variable specifies the port number of the device. The range is from 1 through 65535.

Both the SSH server and the NETCONF server must be enabled to establish a NETCONF session. The **netconf server** command displays the following warning message if the SSH server configuration is disabled.

```
Warning: SSH server is disabled. Please enable the SSH server.
```

Configuring session hello-timeout

A NETCONF session hello-timeout indicates the number of seconds a session waits before the hello message is received from the NETCONF client. A session is dropped if no hello message is received before the specified number of seconds elapses. If this parameter is set to zero, the server never drops a session.

NOTE

Setting the NETCONF session hello-timeout value to zero permits denial of service attacks.

To configure a NETCONF session hello-timeout, enter the following command.

```
Brocade(config)# netconf hello-timeout 300
```

Syntax: [no] netconf hello-timeout [seconds]

The *seconds* variable specifies the number of seconds the server waits to receive a hello message. The range is from 1 through 3600 seconds. The default value is 600 seconds.

Configuring session idle-timeout

A NETCONF session idle-timeout indicates the number of seconds that a session may remain idle without issuing any RPC requests. A session is dropped if it is idle for an interval longer than the specified number of seconds. If this parameter is set to zero, the server never drops a session because it is idle.

To configure a NETCONF session idle-timeout, enter the following command.

```
Brocade(config)# netconf idle-timeout 86400
```

Syntax: [no] netconf idle-timeout [seconds]

The *seconds* variable specifies the number of seconds a session remains idle. The range is from 1 through 360000 seconds. The default value is 3600 seconds.

Displaying NETCONF statistics

To display the NETCONF server level information and statistics, enter the following command.

```
Brocade# show netconf server
NETCONF server status: Enabled, Port: 830, Transport: SSH
  Start Time: Feb  4 19:20:31
  Max allowed sessions: 1, Active sessions: 1
  Hello timeout: 600 seconds, Idle timeout: 3600 seconds
  Server statistics:
    In sessions      : 1           In bad hellos    : 0
    Dropped sessions: 0           In too big rpcs : 0
    In rpcs          : 1           In bad rpcs     : 0
    Out rpcs         : 1           Out rpc errors  : 0
    Out too big rpcs: 0
```

Syntax: show netconf server

[Table 43](#) describes the output of the **show netconf server** command.

TABLE 43 NETCONF server parameters

Field	Description
server status	The admin status (enabled or disabled) of the NETCONF server. Also displays the SSH status, when SSH is not enabled.
Port	The NETCONF server port number.
Transport	The NETCONF transport (currently only SSH is supported).
Start Time	The time at which the NETCONF subsystem is started.
Max allowed sessions	The maximum number of simultaneous NETCONF sessions supported by the server.
Active sessions	The number of active NETCONF sessions.
Hello timeout	The NETCONF session hello message timeout in seconds.
Idle timeout	The NETCONF session idle message timeout in seconds.
In sessions	The number of sessions started.
In bad hellos	The number of sessions silently dropped because an invalid hello message was received.
Dropped sessions	The number of sessions that were abnormally terminated (for example, due to transport close).
In too big rpcs	The total number of RPC requests received by the server that are larger than the supported maximum RPC request size.
In rpcs	The total number of correct RPC requests received by the server.
In bad rpcs	The total number of incorrect RPC messages received by the server. This includes XML parse errors and errors on the RPC layer.
Out rpcs	The total number of RPC reply messages sent by the server containing an <code><rpc-ok></code> element or <code><data></code> element.
Out rpc errors	The total number of RPC reply messages sent by the server containing an <code><rpc-error></code> element.
Out too big rpcs	The total number of RPC reply messages sent by the server containing an <code><rpc-error></code> element with <code>too-big</code> as the error tag.

To display the NETCONF session level statistics, enter the following command.

```
Brocade# show netconf session
Session Id: 1  SSH session Id: 1
  Username: lab  Login time: Feb 7 21:28:47
  Client Ip Address: 10.120.73.112
  Privilege Level: <edit-config> <get-config> <get> <close-session>
  Session Statistics:
    In rpcs      : 1          In bad rpcs   : 0
    Out rpcs     : 1          Out rpc errors : 0
    Edit-Config : 0          Get-Config    : 0
    Get         : 1          Un-supported   : 0
```

Syntax: `show netconf sessions`

Table 44 describes the output of the `show netconf sessions` command.

TABLE 44 NETCONF session parameters

Field	Description
Session Id	The unique identification value for the NETCONF session.
SSH session Id	The unique identification value for the SSH session.
Username	The authenticated SSH user name. The value is <none> for public key authentication.
Login time	The time at which the session is established.
Client Ip Address	The IP address of the NETCONF client.
Privilege Level	The supported NETCONF privilege level operations for a session, where privilege is derived from the SSH user privilege.
In rpcs	The number of correct RPC requests received.
In bad rpcs	The total number of incorrect RPC messages received by the server. This includes XML parse errors and errors on the RPC layer.
Out rpcs	The total number of RPC reply messages sent by the server containing an <rpc-ok> element or <data> element.
Out rpc errors	The total number of RPC reply messages sent by the server containing an <rpc-error> element.
Edit-Config	The number of well-formed <edit-config> operations received.
Get-Config	The number of well-formed <get-config> operations received.
Get	The number of well-formed <get> operations received.
Unsupported	The number of unsupported operations received.

Syslog messages for NETCONF

The following syslog message is generated when the NETCONF session is established.

```
SYSLOG: <14>Feb 8 01:03:00 NETCONF session [1] from 10.20.99.130 user ncradsuper has been established.
```

Syntax: NETCONF session id from IP address user username has been established

The following syslog message is generated when the NETCONF session is disconnected.

```
SYSLOG: <14>Feb 8 01:03:00 NETCONF session [1] from 10.20.99.130 user ncradsuper has been disconnected.
```

Syntax: NETCONF session id from IP address user username has been disconnected

Clearing NETCONF statistics

To clear the NETCONF server level statistics, enter the following command.

```
Brocade# clear netconf server-stats
```

Syntax: clear netconf server-stats

To clear the NETCONF session level statistics, enter the following command.

```
Brocade# clear netconf session-stats
```

Syntax: clear netconf session-stats

System limitations for NETCONF

The following are the system limitations for NETCONF.

- Only one NETCONF session is supported at a time. Any new NETCONF connection requests are rejected after the first session is established.
- Only the <running> configuration data store is supported.
- The <running> configuration data store displays the commands that are currently supported by NETCONF.
- The NETCONF notifications are not supported.
- A partial set of configuration and state display commands are supported.
- The XPATH filtering is not supported.
- A 32K response size limit is supported. An error is returned, if the response size limit is exceeded.
- A 16K request buffer limit is supported. An error is returned, if the request size limit is exceeded.
- Only a subset of the subtree filtering is supported.

Data models and mapping

YANG is a data modeling language used to model configuration and state data manipulated by the NETCONF protocol, NETCONF remote procedure calls, and NETCONF notifications. In order for NETCONF to be an interoperable protocol, models must be defined in a vendor-neutral language. YANG provides the language and rules for defining models for use with NETCONF.

The YANG language is currently being developed by the IETF NETCONF Data Modeling Language Working Group (NETMOD) and is defined in RFC 6020.

Each block of YANG data is encapsulated as a module, containing a header statement, linkage information, meta information, and revision history. Modules can contain one or more submodules with the same structure.

The following code example shows the structure of a header statement, along with linkage and meta information, which contains contact information and a high-level description of the module.

```

module netiron-config
{
  namespace "http://brocade.com/ns/netconf/config/netiron-config/";
  prefix "brcd";

  include common-defs;
  include vlan-config;
  include interface-config;
  include mpls-config;

  organization
  "Brocade Communications Inc.";

  contact
    "Technical Support Center"+
    "130 Holger Way,"+
    "San Jose, CA 95134"+
    "Email: ipsupport@brocade.com"+
    "Phone: 1-800-752-8061"+
    "URL: www.brocade.com";

  description
  "NetIron Config module. VERSION: ";

  revision 2011-04-20
  {
    description "Initial revision";
  }
}

```

Example in YANG, XML, and CLI

Table 45 provides an example to describe the VLAN name in the YANG model and the equivalent XML and CLI.

TABLE 45 Example in YANG, XML, and CLI

YANG	XML	CLI
<pre> leaf vlan-name { type string { length "1..31"; } description "VLAN Name"; } </pre>	<pre> <brcd:vlan-name>example</brcd:vlan-name> </pre>	<pre> [no] vlan vlan-id [name vlan-name] </pre>

For further examples and information on the YANG model, refer to the *Brocade MLX Series and Brocade NetIron Family YANG Guide*.

Foundry Direct Routing and CAM Partition Profiles for the NetIron XMR and the Brocade MLX Series

Table 46 displays the individual devices and the Foundry Direct Routing (FDR) and CAM Partition features they support.

TABLE 46 Supported Foundry Direct Routing and CAM partition features

Features supported	Brocade NetIron XMR Series	Brocade MLX Series	Brocade NetIron CES 2000 Series BASE package	Brocade NetIron CES 2000 Series ME_PREM package	Brocade NetIron CES 2000 Series L3_PREM package	Brocade NetIron CER 2000 Series Base package	Brocade NetIron CER 2000 Series Advanced Services package
Foundry Direct Routing	Yes	Yes	Yes	Yes	Yes	Yes	Yes
CAM Partition Profiles	Yes	Yes	No	No	No	No	No
Supernet CAM Partition Sharing	Yes	Yes	No	No	No	No	No
CAM Overflow Logging	Yes	Yes	No	No	No	No	No
IPv6 Host CAM Enhancement	Yes	Yes	No	No	No	No	No

Configuring Density Mode for the 2x100G CAM

NOTE

There is no configuration support for **cam-mode** in NetIron CES and NetIron CER.

Setting the CAM to double density mode will automatically disable uRPF, while uRPF is allowed under single density mode. The default setting for the XMR is double density. The default setting for the MLX is single density. You can set the density mode using the following command:

```
Brocade(config)# cam-mode ip urpf-100g
```

You must reload the device for this command to take effect.

Syntax: [no] **cam-mode ip urpf-100g**

Configuring IPv6 host CAM mode

NOTE

There is no configuration support for **cam-mode** in NetIron CES and NetIron CER.

The CAM mode for IPv6 routes can be configured to host. You can set the CAM mode to **host** by using the following command.

```
Brocade(config)# cam-mode ipv6 host
```

You must reload the device for this command to take effect.

Syntax: [no] **cam-mode ipv6 host**

The **host** parameter programs the complete 128 bit IPv6 address into the CAM.

Configuring IPv6 host drop CAM limit

To limit the usage of CAM by IPv6 hosts with unresolved ND, enter the **ipv6 max-host-drop-cam** command.

```
Brocade(config)# ipv6 max-host-drop-cam 5
```

Syntax: [no] **ipv6 max-host-drop-cam** [*limit*]

The optional *limit* variable is the IPv6 drop CAM limit for a port per packet processor (PPCR). The limit value can be from 0 through 65535.

CAM partition profiles

CAM is partitioned on the device by a variety of profiles that you can select depending on your application. The available profiles are described in [Table 47](#) for Brocade NetIron XMR and [Table 48](#) for Brocade MLX series.

To implement a CAM partition profile, enter the following command.

```
Brocade(config) cam-partition profile ipv4
```

Syntax: **cam-partition profile** [**ipv4** | **ipv4-ipv6** | **ipv4-ipv6-2** | **ipv4-vpls** | **ipv4-vpn** | **ipv6** | **I2-metro** | **I2-metro-2** | **mpls-I3vpn** | **mpls-I3vpn-2** | **mpls-vpls** | **mpls-vpls-2** | **mpls-vpn-vpls** | **multi-service** | **multi-service-2** | **multi-service-3** | **multi-service-4**]

The **ipv4** parameter adjusts the CAM partitions, as described in [Table 47](#) for Brocade NetIron XMR and [Table 48](#) for Brocade MLX series, to optimize the device for IPv4 applications.

The **ipv4-ipv6** parameter adjusts the CAM partitions, as described in [Table 47](#) for Brocade NetIron XMR and [Table 48](#) for Brocade MLX series, to optimize the device for IPv4 and IPv6 dual stack applications

The **ipv4-ipv6-2** parameter that was introduced in version 03.7.00, adjusts the CAM partitions, as described in [Table 47](#) for Brocade NetIron XMR and [Table 48](#) for Brocade MLX series, to optimize the device for increased IPv4 routes with room or IPv6.

The **ipv4-vpls** parameter adjusts the CAM partitions, as described in [Table 47](#) for Brocade NetIron XMR and [Table 48](#) for Brocade MLX series, to optimize the device for IPv4 and MPLS VPLS applications

The **ipv4-vpn** parameter adjusts the CAM partitions, as described in [Table 47](#) for Brocade NetIron XMR and [Table 48](#) for Brocade MLX series, to optimize the device for IPv4 and MPLS Layer-3 VPN applications

The **ipv6** parameter adjusts the CAM partitions, as described in [Table 47](#) for Brocade NetIron XMR and [Table 48](#) for Brocade MLX series, to optimize the device for IPv6 applications.

The **I2-metro** parameter adjusts the CAM partitions, as described in [Table 47](#) for Brocade NetIron XMR and [Table 48](#) for Brocade MLX series routers, to optimize the device for Layer 2 Metro applications.

The **I2-metro-2** parameter provides another alternative to **I2-metro** to optimize the device for Layer 2 Metro applications. It adjusts the CAM partitions, as described in [Table 47](#) for Brocade NetIron XMR and [Table 48](#) for Brocade MLX series routers.

The **mpls-l3vpn** parameter adjusts the CAM partitions, as described in [Table 47](#) for Brocade NetIron XMR and [Table 48](#) for Brocade MLX series routers, to optimize the device for Layer 3, BGP or MPLS VPN applications.

The **mpls-l3vpn-2** parameter provides another alternative to **mpls-l3vpn** to optimize the device for Layer 3, BGP or MPLS VPN applications. It adjusts the CAM partitions, as described in [Table 47](#) for Brocade NetIron XMR and [Table 48](#) for Brocade MLX series routers.

The **mpls-vpls** parameter adjusts the CAM partitions, as described in [Table 47](#) for Brocade NetIron XMR and [Table 48](#) for Brocade MLX series routers, to optimize the device for MPLS VPLS applications.

The **mpls-vpls-2** parameter provides another alternative to **mpls-vpls** to optimize the device for MPLS VPLS applications. It adjusts the CAM partitions, as described in [Table 47](#) for Brocade NetIron XMR and [Table 48](#) for Brocade MLX series routers.

The **mpls-vpn-vpls** parameter adjusts the CAM partitions, as described in [Table 47](#) for Brocade NetIron XMR and [Table 48](#) for Brocade MLX series routers, to optimize the device for MPLS Layer-3 and Layer-2 VPN applications.

The **multi-service** parameter adjusts the CAM partitions, as described in [Table 47](#) for Brocade NetIron XMR and [Table 48](#) for Brocade MLX series routers, to optimize the device for Multi-Service applications.

The **multi-service-2** parameter provides another alternative to **multi-service** to optimize the device for Multi-Service applications. It adjusts the CAM partitions, as described in [Table 47](#) for Brocade NetIron XMR and [Table 48](#) for Brocade MLX series routers.

NOTE

You must reload your device for this command to take effect.

The **multi-service-3** parameter provides another alternative to **multi-service** to optimize the device for Multi-Service applications to support IPv6 VRF. It adjusts the CAM partitions, as described in [Table 47](#) for Brocade NetIron XMR and [Table 48](#) for Brocade MLX series routers.

The **multi-service-4** parameter provides another alternative to **multi-service** to optimize the device for Multi-Service applications to support IPv6 VRF. It adjusts the CAM partitions, as described in [Table 47](#) for Brocade NetIron XMR and [Table 48](#) for Brocade MLX series routers.

There are fourteen CAM partitioning profiles for Brocade NetIron XMR and [Table 48](#) for Brocade MLX series routers. The profiles for Brocade XMR routers are described in [Table 47](#) and the profiles for Brocade MLX routers are described in [Table 48](#).

TABLE 47 CAM partitioning profiles available for Brocade NetIron XMR routers

Profile	IPv4	IPv6	MAC or VPLS MAC	IPv4 VPN	IPv6 VPN	IPv4 or L2 Inbound ACL	IPv6 Inbound ACL	IPv4 or L2 Outbound ACL	IPv6 Outbound ACL
Default Profile	Logical size: 512K	Logical size: 64K	Logical size: 128K	Logical size: 128K	0	Logical size: 48K	Logical size: 4K	Logical size: 48K	Logical size: 4K
ipv4 Profile	Logical size: 1M	0	Logical size: 32K	0	0	Logical size: 112K	0	Logical size: 64K	0
ipv6 Profile	Logical size: 64K	Logical size: 240K	Logical size: 32K	0	0	Logical size: 16K	Logical size: 24K	Logical size: 16K	Logical size: 12K
I2-metro Profile	Logical size: 256K	0	Logical size: 512K	0	0	Logical size: 64K	0	Logical size: 64K	0
mpls-I3vpn Profile	Logical size: 256K	0	Logical size: 32K	Logical size: 480K	0	Logical size: 64K	0	Logical size: 64K	0
mpls-vpls Profile	Logical size: 256K	0	Logical size: 512K	0	0	Logical size: 64K	0	Logical size: 64K	0
multi-service Profile	Logical size: 256K	Logical size: 32K	Logical size: 192K	Logical size: 256K	0	Logical size: 32K	Logical size: 8K	Logical size: 32K	Logical size: 8K
multi-service-2 Profile	512K	64K	128K	128K	0	48K	4K	48K	4K
multi-service-3 Profile	256K	32K	128K	192K	32K	32K	8K	32K	8K
multi-service-4 Profile	768K	32K	64K	64K	8K	32K	8K	48K	4K
mpls-vpn-vpls Profile	Logical size: 128K	0	Logical size: 224K	Logical size: 384K	0	Logical size: 48K	0	Logical size: 64K	0
ipv4-vpn Profile	Logical size: 320K	0	Logical size: 32K	Logical size: 448K	0	Logical size: 64K	0	Logical size: 64K	0
I2-metro-2 Profile	Logical size: 64K	0	Logical size: 608K	0	0	Logical size: 64K	0	Logical size: 64K	0
mpls-I3vpn-2 Profile	Logical size: 128K	0	Logical size: 32K	Logical size: 544K	0	Logical size: 64K	0	Logical size: 64K	0
mpls-vpls-2 Profile	Logical size: 128K	0	Logical size: 576K	0	0	Logical size: 64K	0	Logical size: 64K	0

TABLE 47 CAM partitioning profiles available for Brocade NetIron XMR routers (Continued)

Profile	IPv4	IPv6	MAC or VPLS MAC	IPv4 VPN	IPv6 VPN	IPv4 or L2 Inbound ACL	IPv6 Inbound ACL	IPv4 or L2 Outbound ACL	IPv6 Outbound ACL
ipv4-ipv6 Profile	Logical size: 320K	Logical size: 160K	Logical size: 32K	0	0	Logical size: 48K	Logical size: 20K	Logical size: 32K	Logical size: 8K
ipv4-vpls Profile	Logical size: 320K	0	Logical size: 480K	0	0	Logical size: 64K	0	Logical size: 64K	0
ipv4-ipv6-2 Profile	Logical size: 768K	Logical size: 64K	Logical size: 64K	0	0	Logical size: 64K	Logical size: 8K	Logical size: 48K	Logical size: 4K

TABLE 48 CAM partitioning profiles available for Brocade MLX Series routers

Profile	IPv4	IPv6	MAC or VPLS MAC	IPv4 VPN	IPv6 VPN	IPv4 or L2 Inbound ACL	IPv6 Inbound ACL	IPv4 or L2 Outbound ACL	IPv6 Outbound ACL
Default Profile	Logical size: 256K	Logical size: 32K	Logical size: 64K	Logical size: 64K	0	Logical size: 24K	Logical size: 2K	Logical size: 48K	Logical size: 4K
ipv4 Profile	Logical size: 512K	0	Logical size: 16K	0	0	Logical size: 56K	0	Logical size: 64K	0
ipv6 Profile	Logical size: 64K	Logical size: 112K	Logical size: 16K	0	0	Logical size: 8K	Logical size: 12K	Logical size: 16K	Logical size: 12K
I2-metro Profile	Logical size: 128K	0	Logical size: 256K	0	0	Logical size: 32K	0	Logical size: 64K	0
mpls-l3vpn Profile	Logical size: 128K	0	Logical size: 16K	Logical size: 240K	0	Logical size: 32K	0	Logical size: 64K	0
mpls-vpls Profile	Logical size: 128K	0	Logical size: 256K	0	0	Logical size: 32K	0	Logical size: 64K	0
multi-service Profile	Logical size: 128K	Logical size: 16K	Logical size: 96K	Logical size: 128K	0	Logical size: 16K	Logical size: 4K	Logical size: 32K	Logical size: 8K
multi-service-2 Profile	448K	16K	32K	32K	0	24K	2K	48K	4K
multi-service-3 Profile	128K	16K	64K	96K	32K	16K	4K	32K	8K
multi-service-4 Profile	448K	8K	0	32K	8K	16K	4K	48K	4K

TABLE 48 CAM partitioning profiles available for Brocade MLX Series routers

Profile	IPv4	IPv6	MAC or VPLS MAC	IPv4 VPN	IPv6 VPN	IPv4 or L2 Inbound ACL	IPv6 Inbound ACL	IPv4 or L2 Outbound ACL	IPv6 Outbound ACL
mpls-vpn-vpls Profile	Logical size: 64K	0	Logical size: 112K	Logical size: 192K	0	Logical size: 24K	0	Logical size: 64K	0
ipv4-vpn Profile	Logical size: 160K	0	16K	Logical size: 224K	0	Logical size: 32K	0	Logical size: 64K	0
I2-metro-2 Profile	Logical size: 64K	0	Logical size: 288K	0	0	Logical size: 32K	0	Logical size: 64K	0
mpls-I3vpn-2 Profile	Logical size: 64K	0	Logical size: 16K	Logical size: 272K	0	Logical size: 32K	0	Logical size: 64K	0
mpls-vpls-2 Profile	Logical size: 64K	0	Logical size: 288K	0	0	Logical size: 32K	0	Logical size: 64K	0
ipv4-ipv6 Profile	Logical size: 160K	Logical size: 80K	Logical size: 16K	0	0	Logical size: 24K	Logical size: 10K	Logical size: 32K	Logical size: 8K
ipv4-vpls Profile	Logical size: 160K	0	Logical size: 240K	0	0	Logical size: 32K	0	Logical size: 64K	0
ipv4-ipv6-2 Profile	480K	8K	32K	0	0	32K	4K	48K	4K

TABLE 49 CAM partitioning profiles available for the NI-MLX-10Gx8-D modules

Profile	IPv4	IPv6	MAC or VPLS MAC	IPv4 VPN	IPv6 VPN	IPv4 or L2 Inbound ACL	IPv6 Inbound ACL	IPv4 or L2 Outbound ACL	IPv6 Outbound ACL
Default Profile	128K	16K	64K	32K	0	24K	2K	48K	4K
ipv4 Profile	256K	0	16K	0	0	56K	0	64K	0
ipv6 Profile	68K	48K	16K	0	0	8K	12K	16K	12K
I2-metro Profile	64K	0	160K	0	0	32K	0	64K	0
mpls-I3vpn Profile	64K	0	16K	144K	0	32K	0K	64K	0
mpls-vpls Profile	64K	0	160K	0	0	32K	0	64K	0

Profile	IPv4	IPv6	MAC or VPLS MAC	IPv4 VPN	IPv6 VPN	IPv4 or L2 Inbound ACL	IPv6 Inbound ACL	IPv4 or L2 Outbound ACL	IPv6 Outbound ACL
multi-service Profile	64K	8K	80K	64K	0	16K	4K	32K	8K
multi-service-2 Profile	192K	16K	32K	32K	0	24K	2K	48K	4K
multi-service-3 Profile	64K	8K	32K	48K	32K	16K	4K	32K	8K
multi-service-4 Profile	198K	8K	32K	48K	32K	16K	4K	32K	8K
mpls-vpn-vpls Profile	64K	0	80K	96K	0	24K	0	64K	0
ipv4-vpn Profile	96K	0K	16K	128K	0	32K	0	64K	0
l2-metro-2 Profile	64K	0	160K	0	0	32K	0	64K	0
mpls-l3vpn-2 Profile	64K	0	16K	144K	0	32K	0	64K	0
mpis-vpls-2 Profile	64K	0	160K	0	0	32K	0	64K	0
ipv4-ipv6 Profile	96K	32K	16K	0	0	24K	10K	32K	8K
ipv4-vpls Profile	64K	0	160K	0	0	32K	0	64K	0
ipv4-ipv6-2 Profile	192K	16K	32K	0	0	32K	4K	48K	4K

TABLE 50

CAM partitioning profiles available for the BR-MLX-100Gx2-X modules

Profile	IPv4 Note: All IPv4 values listed are when operating in double density mode	IPv6	MAC or VPLS MAC	IPv4 VPN	IPv6 VPN	L2 Inbound ACL	IPv4 Inbound ACL	IPv6 Inbound ACL	IPv4 or L2 Outbound ACL	IPv6 Outbound ACL
Default Profile	512K	64K	160K	256K	0	16K	96K	16K	48K	8K
ipv4 Profile	1024K	0	160K	0	0	16K	128K	0	64K	0
ipv6 Profile	64K	240K	160K	0	0	16K	32K	48K	16K	24K
l2-metro Profile	1024K	0	160K	0	0	16K	128K	0	64K	0
l2-metro-2 Profile	1024K	0	160K	0	0	16K	128K	0	64K	0
multi-service Profile	288K	40K	160K	576K	0	16K	64K	32K	32K	16K
multi-service-2 Profile	512K	64K	160K	128K	0	96K	96K	16K	48K	8K
multi-service-3 Profile	256K	32K	160K	384K	64K	16K	80K	24K	32K	16K
multi-service-4 Profile	384K	16K	160K	64K	32K	16K	64K	32K	48K	8K
mpls-vpn-vpls Profile	128K	0	160K	896K	0	16K	128K	0	64K	0
mpls-l3vpn Profile	128K	0	160K	896K	0	16K	128K	0	64K	0
mpls-l3vpn-2 Profile	1024K	0	160K	0	0	16K	128K	0	64K	0
mpls-vpls Profile	1024K	0	160K	0	0	16K	128K	0	64K	0
mpls-vpls-2 Profile	1024K	0	160K	0	0	16K	128K	0	64K	0
ipv4-vpn Profile	320K	0	160K	704K	0	16K	128K	0	64K	0
ipv4-ipv6 Profile	320K	176K	160K	0	0	16K	48K	40K	32K	16K
ipv4-vpls Profile	1024K	0	160K	0	0	16K	96K	16K	64K	0
ipv4-ipv6-2 Profile	768K	64K	160K	0	0	16K	96K	16K	48K	8K

TABLE 51

CAM partitioning profiles available for the BR-MLX-10GX24-DM modules

Profile	IPv4	IPv6	MAC	IPv4 VPN	IPv4 ACL/ MCAST VPLS	IPv6 VPN	IPv6 DAVC	IPv6 ACL	OUT ACL	OUT_ IPv6 ACL	Src_ Ingrs Chk	MCAST VPLS	OUT_ LBL ACL	SRVC LKUP	L2 ACL
Default Profile	128K	16K	128K	32K	48K	0	0	16K	48K	8K	0	NA	NA	64K	16K
ipv4 Profile	256K	0	128K	0	80K	0	0	0	64K	0	0	NA	NA	64K	16K
ipv6 Profile	32K	56K	128K	0	16K	0	0	32K	16K	24K	0	NA	NA	64K	16K
mpls-vpn Profile	64K	0	128K	96K	80K	0	0	0	64K	0	0	NA	NA	64K	16K
mpls-vpls Profile	256K	0	128K	0	80K	0	0	0	64K	0	0	NA	NA	64K	16K
I2-metro Profile	256K	0	128K	0	80K	0	0	0	64K	0	0	NA	NA	64K	16K
I2-metro-2 Profile	256K	0	128K	0	80K	0	0	0	64K	0	0	NA	NA	64K	16K
mpls-vpn-2 Profile	32K	0	128K	112K	80K	0	0	0	64K	0	0	NA	NA	64K	16K
mpls-vpls-2 Profile	256K	0	128K	0	80K	0	0	0	64K	0	0	NA	NA	64K	16K
multi-service Profile	96K	8	128K	64K	48K	0	0	16K	32K	16K	0	NA	NA	64K	16K
mpls- vpn-vpls Profile	64K	0	128K	96K	80K	0	0	0	64K	0	0	NA	NA	64K	16K
ipv4- vpn Profile	64K	0	128K	96K	80K	0	0	0	64K	0	0	NA	NA	64K	16K
ipv4- ipv6 Profile	96K	40	128K	0	64K	0	0	8K	32K	16K	0	NA	NA	64K	16K
ipv4- vpls Profile	256K	0	128K	0	80K	0	0	0	64K	0	0	NA	NA	64K	16K
ipv4- ipv6-2 Profile	224K	8K	128K	0	48K	0	0	16K	48K	8K	0	NA	NA	64K	16K
multi-service- 2 Profile	6K	8K	128K	16K	64K	0	0	8K	48K	8K	0	NA	NA	64K	16K
multi-service- 3 Profile	192K	8K	128K	48K	48K	16K	0	16K	32K	16K	0	NA	NA	64K	16K
multi-service- 4 Profile	128K	8K	128K	32K	48K	8K	0	16K	48K	8K	0	NA	NA	64K	16K

Supernet CAM partition sharing

As of 3.2.00 code and later, TCAM sharing within a particular CAM Section is now supported.

For example: In Multi-Service IronWare software prior to version 03.2.00, the TCAM resources could not be shared between the 32 levels of the IP Forwarding Information Base (FIB). Beginning with version 03.2.00, TCAM allocation is optimized to allow dynamic allocation of resources to each level within a particular resource pool. If one level runs out of TCAM resources, it can use resources that have been allocated to another level and remain unused. This feature is applicable to IPv4, IPv6, and L-3 VPN routes.

NOTE

CAM Sharing is not shared across resource pools, such as IPv4, IPv6 or L-3 VPN. Only shared between levels of each pool.

For example: IPv4 may not use CAM resources from the IPv6 resource pool.

Displaying CAM partition

The **show cam-partition** command provides information about available CAM in three formats: raw size, user size, and reserved size.

```
Brocade#show cam-partition
CAM partitioning profile: default
Slot 1 XPP20SP 0:
# of CAM device           = 4
Total CAM Size           = 917504 entries (63Mbits)
IP: Raw Size 524288, User Size 524288(0 reserved)
  Subpartition 0: Raw Size 12288, User Size 12288, (0 reserved)
  Subpartition 1: Raw Size 468107, User Size 468107, (0 reserved)
  Subpartition 2: Raw Size 37335, User Size 37335, (0 reserved)
  Subpartition 3: Raw Size 5140, User Size 5140, (0 reserved)
  Subpartition 4: Raw Size 778, User Size 778, (0 reserved)
IPv6: Raw Size 131072, User Size 65536(0 reserved)
  Subpartition 0: Raw Size 12288, User Size 6144, (0 reserved)
  Subpartition 1: Raw Size 107496, User Size 53748, (0 reserved)
  Subpartition 2: Raw Size 9332, User Size 4666, (0 reserved)
  Subpartition 3: Raw Size 1284, User Size 642, (0 reserved)
  Subpartition 4: Raw Size 384, User Size 192, (0 reserved)
IP VPN Raw Size 131072, User Size 131072(0 reserved)
  Subpartition 0: Raw Size 2048, User Size 2048, (0 reserved)
  Subpartition 1: Raw Size 116886, User Size 116886, (0 reserved)
  Subpartition 2: Raw Size 9333, User Size 9333, (0 reserved)
  Subpartition 3: Raw Size 1285, User Size 1285, (0 reserved)
  Subpartition 4: Raw Size 384, User Size 384, (0 reserved)
MAC: Raw Size 131072, User Size 131072(0 reserved)
  Subpartition 0: Raw Size 10, User Size 10, (0 reserved)
  Subpartition 1: Raw Size 32, User Size 32, (0 reserved)
  Subpartition 2: Raw Size 131030, User Size 131030, (0 reserved)
Session: Raw Size 98304, User Size 49152(0 reserved)
  Subpartition 0: Raw Size 79872, User Size 39936, (0 reserved)
  Subpartition 1: Raw Size 2048, User Size 1024, (0 reserved)
  Subpartition 2: Raw Size 16384, User Size 8192, (0 reserved)
IPv6 Session: Raw Size 32768, User Size 4096(0 reserved)
  Subpartition 0: Raw Size 15872, User Size 1984, (0 reserved)
  Subpartition 1: Raw Size 512, User Size 64, (0 reserved)
  Subpartition 2: Raw Size 16384, User Size 2048, (0 reserved)
Out Session: Raw Size 196608, User Size 98304(49152 reserved)
```

```

Out IPv6 Session: Raw Size 65536, User Size 8192(4096 reserved)
Slot 1 XPP20SP 0:
IP Section(Left):      0(000000) - 262143(03ffff)
IP Section(Right):    0 (000000) - 262143 (03ffff)
  IP SNet 0:(Left):    0(000000) - 12287(002fff)
  IP SNet 1:(Left): 12288(003000) - 262143(03ffff)
  IP SNet 1:(Right): 0 (000000) - 218250 (03548a)
  IP SNet 2:(Right): 218251 (03548b) - 255585 (03e661)
  IP SNet 3:(Right): 255586 (03e662) - 260725 (03fa75)
  IP SNet 4:(Right): 260726 (03fa76) - 261503 (03fd7f)
  IP SNet 5:(Right): 261504 (03fd80) - 261631 (03fdff)
  IP SNet 6:(Right): 261632 (03fe00) - 261695 (03fe3f)
  IP SNet 7:(Right): 261696 (03fe40) - 261727 (03fe5f)
  IP SNet 8:(Right): 261728 (03fe60) - 261759 (03fe7f)
  IP SNet 9:(Right): 261760 (03fe80) - 261791 (03fe9f)
  IP SNet 10:(Right): 261792 (03fea0) - 261807 (03feaf)
  IP SNet 11:(Right): 261808 (03feb0) - 261823 (03febf)
  IP SNet 12:(Right): 261824 (03fec0) - 261839 (03fecf)
  IP SNet 13:(Right): 261840 (03fed0) - 261855 (03fedf)
  IP SNet 14:(Right): 261856 (03fee0) - 261871 (03feef)
  IP SNet 15:(Right): 261872 (03fef0) - 261887 (03feff)
  IP SNet 16:(Right): 261888 (03ff00) - 261903 (03ff0f)
  IP SNet 17:(Right): 261904 (03ff10) - 261919 (03ff1f)
  IP SNet 18:(Right): 261920 (03ff20) - 261935 (03ff2f)
  IP SNet 19:(Right): 261936 (03ff30) - 261951 (03ff3f)
  IP SNet 20:(Right): 261952 (03ff40) - 261967 (03ff4f)
  IP SNet 21:(Right): 261968 (03ff50) - 261983 (03ff5f)
  IP SNet 22:(Right): 261984 (03ff60) - 261999 (03ff6f)
  IP SNet 23:(Right): 262000 (03ff70) - 262015 (03ff7f)
  IP SNet 24:(Right): 262016 (03ff80) - 262031 (03ff8f)
  IP SNet 25:(Right): 262032 (03ff90) - 262047 (03ff9f)
  IP SNet 26:(Right): 262048 (03ffa0) - 262063 (03ffaf)
  IP SNet 27:(Right): 262064 (03ffb0) - 262079 (03ffbf)
  IP SNet 28:(Right): 262080 (03ffc0) - 262095 (03ffcf)
  IP SNet 29:(Right): 262096 (03ffd0) - 262111 (03ffdf)
  IP SNet 30:(Right): 262112 (03ffe0) - 262127 (03ffef)
  IP SNet 31:(Right): 262128 (03fff0) - 262143 (03ffff)
IPV6 Section      : 262144 (040000) - 393215 (05ffff)
  IPV6 SNet 0: 262144 (040000) - 274431 (042fff)
  IPV6 SNet 1: 274432 (043000) - 381927 (05d3e7)
  IPV6 SNet 2: 381928 (05d3e8) - 391259 (05f85b)
  IPV6 SNet 3: 391260 (05f85c) - 392543 (05fd5f)
  IPV6 SNet 4: 392544 (05fd60) - 392927 (05fedf)
  IPV6 SNet 5: 392928 (05fee0) - 393055 (05ff5f)
  IPV6 SNet 6: 393056 (05ff60) - 393119 (05ff9f)
  IPV6 SNet 7: 393120 (05ffa0) - 393151 (05ffbf)
  IPV6 SNet 8: 393152 (05ffc0) - 393183 (05ffdf)
  IPV6 SNet 9: 393184 (05ffe0) - 393215 (05ffff)
IP VPN Section: 393216 (060000) - 524287 (07ffff)
  IP VPN SNet 0: 393216 (060000) - 395263 (0607ff)
  IP VPN SNet 1: 395264 (060800) - 512149 (07d095)
  IP VPN SNet 2: 512150 (07d096) - 521482 (07f50a)
  IP VPN SNet 3: 521483 (07f50b) - 522767 (07fa0f)
  IP VPN SNet 4: 522768 (07fa10) - 523151 (07fb8f)
  IP VPN SNet 5: 523152 (07fb90) - 523279 (07fc0f)
  IP VPN SNet 6: 523280 (07fc10) - 523343 (07fc4f)
  IP VPN SNet 7: 523344 (07fc50) - 523375 (07fc6f)
  IP VPN SNet 8: 523376 (07fc70) - 523407 (07fc8f)
  IP VPN SNet 9: 523408 (07fc90) - 523439 (07fcdf)
  IP VPN SNet 10: 523440 (07fcb0) - 523455 (07fcbf)

```

```

IP VPN SNet 11: 523456 (07fcc0) - 523471 (07cccf)
IP VPN SNet 12: 523472 (07fcd0) - 523487 (07cdf)
IP VPN SNet 13: 523488 (07fce0) - 523503 (07cef)
IP VPN SNet 14: 523504 (07fcf0) - 523519 (07ccff)
IP VPN SNet 15: 523520 (07fd00) - 523535 (07fd0f)
IP VPN SNet 16: 523536 (07fd10) - 523551 (07fd1f)
IP VPN SNet 17: 523552 (07fd20) - 523567 (07fd2f)
IP VPN SNet 18: 523568 (07fd30) - 523583 (07fd3f)
IP VPN SNet 19: 523584 (07fd40) - 523599 (07fd4f)
IP VPN SNet 20: 523600 (07fd50) - 523615 (07fd5f)
IP VPN SNet 21: 523616 (07fd60) - 523631 (07fd6f)
IP VPN SNet 22: 523632 (07fd70) - 523647 (07fd7f)
IP VPN SNet 23: 523648 (07fd80) - 523663 (07fd8f)
IP VPN SNet 24: 523664 (07fd90) - 523679 (07fd9f)
IP VPN SNet 25: 523680 (07fda0) - 523695 (07daf)
IP VPN SNet 26: 523696 (07fdb0) - 523711 (07dbf)
IP VPN SNet 27: 523712 (07fdc0) - 523727 (07dcf)
IP VPN SNet 28: 523728 (07fdd0) - 523743 (07ddf)
IP VPN SNet 29: 523744 (07fde0) - 523759 (07def)
IP VPN SNet 30: 523760 (07fdf0) - 523775 (07dff)
IP VPN SNet 31: 523776 (07fe00) - 524287 (07ffff)
MAC Section      : 524288 (080000) - 655359 (09ffff)
MAC Forwarding  : 524288 (080000) - 655317 (09ffd5)
MAC Flooding    : 655318 (09ffd6) - 655327 (09ffd6)
Misc Protocol   : 655350 (09fff6) - 655381 (0a0015)
Session Section : 655360 (0a0000) - 753663 (0b7fff)
IP Multicast    : 655360 (0a0000) - 671743 (0a3fff)
Broadcast ACL : 673792 (0a4800) - 675839 (0a4fff)
Receive ACL    : 671744 (0a4000) - 673791 (0a47ff)
Rule-based ACL : 673792 (0a4800) - 753663 (0b7fff)
IPv6 Session Sec: 753664 (0b8000) - 786431 (0bffff)
IP Multicast    : 753664 (0b8000) - 770047 (0bbfff)
Receive ACL    : 770048 (0bc000) - 770559 (0bc1ff)
Rule-based ACL : 770560 (0bc200) - 786431 (0bffff)
Out Session     : 786432 (0c0000) - 983039 (0effff)
Out IPv6 Session: 983040 (0f0000) - 104857 (0fffff)
...

```

Syntax: show cam-partition

The output displays the CAM partitioning profile name, slot number, number of CAM device, and total CAM size. It also displays the raw size, user size, and reserved size for each of the CAM sub-partitions.

The information related to the broadcast ACL CAM partition is shown in bold text in the previous output.

Table 52 describes the output parameters of the **show cam-partition** command.

TABLE 52 Output parameters of the **show cam-partition** command

Field	Description
CAM partitioning profile	Shows the CAM profile name.
Slot	Shows the slot number.
# of CAM device	Shows the number of the CAM device.
Total CAM Size	Shows the total available CAM size.

TABLE 52 Output parameters of the **show cam-partition** command (Continued)

Field	Description
Raw Size	Shows the value double that of the CAM partition standard entry count. A standard entry contains 64 bits for the data and 64 bits for the mask. The raw size may cover invalid entries.
User Size	Shows the actual number of entries that the application can use. For a 128-bit application, such as Layer 4 ACL and IPV6, two standard entries equal one user entry. The user size may also cover invalid entries.
reserved	Shows the number of entries not usable in a specific sub-partition.
IP	Shows the raw size, user size, and reserved size for the IP CAM partition and its sub-partitions.
IPv6	Shows the raw size, user size, and reserved size for the IPv6 CAM partition and its sub-partitions.
IP VPN	Shows the raw size, user size, and reserved size for the IP VPN CAM partition and its sub-partitions.
MAC	Shows the raw size, user size, and reserved size for the MAC CAM partition and its sub-partitions.
Session	Shows the raw size, user size, and reserved size for the session CAM partition and its sub-partitions.
IPv6 Session	Shows the raw size, user size, and reserved size for the IPv6 session CAM partition and its sub-partitions.
Out Session	Shows the raw size, user size, and reserved size for the out session CAM partition and its sub-partitions.
Out IPv6 Session	Shows the raw size, user size, and reserved size for the out IPv6 session CAM partition and its sub-partitions.
IP Section	Shows the CAM partition size of the IP section and its subnets.
IPv6 Section	Shows the CAM partition size of the IPv6 section and its subnets.
IP VPN Section	Shows the CAM partition size of the IP VPN section and its subnets.
MAC Section	Shows the CAM partition size of the MAC section.
MAC Forwarding	Shows the CAM partition size of the MAC forwarding section.
MAC Flooding	Shows the CAM partition size of the MAC flooding section.
Misc Protocol	Shows the CAM partition size of the miscellaneous protocol section.
Session Section	Shows the CAM partition size of the session section.
IP Multicast	Shows the CAM partition size of the IP multicast ACL.
Broadcast ACL	Shows the CAM partition size of the IP broadcast ACL.
Receive ACL	Shows the CAM partition size of the IP receive ACL.
Rule-based ACL	Shows the CAM partition size of the rule-based ACL.
IPv6 Session Sec	Shows the CAM partition size of the IPv6 session section.

Displaying CAM Partition for IPv6 VPN

The IPv6 VPN CAM partition is created when multi-service-3 or multi-service-4 CAM profile is configured. The IPv6 VPN CAM partition contains 10 sub partitions. The sub-partition is allocated with a fixed size, but can be dynamically changed. If the size of sub-partition is dynamically changed, the output from the **show cam-partition** command is affected. The following example displays information about IPv6 VPN CAM partition when the current CAM profile is multi-service-3:

```
Brocade# show cam-partition
CAM partitioning profile: multi-service-3
Slot 1 XPP20SP 0:
# of CAM device           = 4
Total CAM Size            = 917504 entries (63Mbits)
.....
IPv6 VPN: Raw Size 131072, User Size 65536(0 reserved)
  Subpartition 0: Raw Size 2048, User Size 1024, (0 reserved)
  Subpartition 1: Raw Size 117734, User Size 58867, (0 reserved)
  Subpartition 2: Raw Size 9333, User Size 4666, (0 reserved)
  Subpartition 3: Raw Size 1285, User Size 642, (0 reserved)
  Subpartition 4: Raw Size 384, User Size 192, (0 reserved)
.....
Slot 1 XPP20SP 0:
.....
IPv6 VPN Section: 524288 (080000) - 655359 (09ffff)
  IPV6 VPN SNet 0: 524288 (080000) - 526335 (0807ff)
  IPV6 VPN SNet 1: 526336 (080800) - 644069 (09d3e5)
  IPV6 VPN SNet 2: 644070 (09d3e6) - 653402 (09f85a)
  IPV6 VPN SNet 3: 653403 (09f85b) - 654687 (09fd5f)
  IPV6 VPN SNet 4: 654688 (09fd60) - 655071 (09fedf)
  IPV6 VPN SNet 5: 655072 (09fee0) - 655199 (09ff5f)
  IPV6 VPN SNet 6: 655200 (09ff60) - 655263 (09ff9f)
  IPV6 VPN SNet 7: 655264 (09ffa0) - 655295 (09ffbf)
  IPV6 VPN SNet 8: 655296 (09ffc0) - 655327 (09ffdf)
  IPV6 VPN SNet 9: 655328 (09ffe0) - 655359 (09ffff)
```

Syntax: show cam-partition

Output from show CAM partition usage command

The **show cam-partition usage** command shows the CAM size available per partition, the amount free, and the percent used. This information is shown here for slot 1.

```
Brocade# show cam-partition usage
CAM partitioning profile: multi-service-3

Slot 1 XPP20SP 0:
Slot 1 XPP20SP 0:
  [IP]262144(size), 262129(free), 00.00%(used)
  :SNet 0: 2048(size), 2036(free), 00.58%(used)
  :SNet 1:237830(size), 237828(free), 00.00%(used)
  :SNet 2: 18667(size), 18667(free), 00.00%(used)
  :SNet 3: 2570(size), 2570(free), 00.00%(used)
  :SNet 4: 389(size), 389(free), 00.00%(used)
  :SNet 5: 128(size), 128(free), 00.00%(used)
  :SNet 6: 64(size), 64(free), 00.00%(used)
  :SNet 7: 32(size), 32(free), 00.00%(used)
  :SNet 8: 32(size), 32(free), 00.00%(used)
  :SNet 9: 32(size), 32(free), 00.00%(used)
```

```

:SNet 10: 16(size), 16(free), 00.00%(used)
:SNet 11: 16(size), 16(free), 00.00%(used)
:SNet 12: 16(size), 16(free), 00.00%(used)
:SNet 13: 16(size), 16(free), 00.00%(used)
:SNet 14: 16(size), 16(free), 00.00%(used)
:SNet 15: 16(size), 16(free), 00.00%(used)
:SNet 16: 16(size), 16(free), 00.00%(used)
:SNet 17: 16(size), 16(free), 00.00%(used)
:SNet 18: 16(size), 16(free), 00.00%(used)
:SNet 19: 16(size), 16(free), 00.00%(used)
:SNet 20: 16(size), 16(free), 00.00%(used)
:SNet 21: 16(size), 16(free), 00.00%(used)
:SNet 22: 16(size), 16(free), 00.00%(used)
:SNet 23: 16(size), 16(free), 00.00%(used)
:SNet 24: 16(size), 16(free), 00.00%(used)
:SNet 25: 16(size), 16(free), 00.00%(used)
:SNet 26: 16(size), 16(free), 00.00%(used)
:SNet 27: 16(size), 16(free), 00.00%(used)
:SNet 28: 16(size), 16(free), 00.00%(used)
:SNet 29: 16(size), 16(free), 00.00%(used)
:SNet 30: 16(size), 16(free), 00.00%(used)
:SNet 31: 16(size), 15(free), 06.25%(used)

[IPV6] 32768(size), 32762(free), 00.01%(used)
:SNet 0: 1024(size), 1022(free), 00.19%(used)
:SNet 1: 28756(size), 28754(free), 00.00%(used)
:SNet 2: 2332(size), 2332(free), 00.00%(used)
:SNet 3: 320(size), 320(free), 00.00%(used)
:SNet 4: 192(size), 192(free), 00.00%(used)
:SNet 5: 64(size), 64(free), 00.00%(used)
:SNet 6: 32(size), 32(free), 00.00%(used)
:SNet 7: 16(size), 16(free), 00.00%(used)
:SNet 8: 16(size), 15(free), 06.25%(used)
:SNet 9: 16(size), 15(free), 06.25%(used)

[IP VPN] 196608(size), 196532(free), 00.03%(used)
:SNet 0: 2048(size), 1999(free), 02.39%(used)
:SNet 1: 177113(size), 177086(free), 00.01%(used)
:SNet 2: 14000(size), 14000(free), 00.00%(used)
:SNet 3: 1927(size), 1927(free), 00.00%(used)
:SNet 4: 384(size), 384(free), 00.00%(used)
:SNet 5: 128(size), 128(free), 00.00%(used)
:SNet 6: 64(size), 64(free), 00.00%(used)
:SNet 7: 32(size), 32(free), 00.00%(used)
:SNet 8: 32(size), 32(free), 00.00%(used)
:SNet 9: 32(size), 32(free), 00.00%(used)
:SNet 10: 16(size), 16(free), 00.00%(used)
:SNet 11: 16(size), 16(free), 00.00%(used)
:SNet 12: 16(size), 16(free), 00.00%(used)
:SNet 13: 16(size), 16(free), 00.00%(used)
:SNet 14: 16(size), 16(free), 00.00%(used)
:SNet 15: 16(size), 16(free), 00.00%(used)
:SNet 16: 16(size), 16(free), 00.00%(used)
:SNet 17: 16(size), 16(free), 00.00%(used)
:SNet 18: 16(size), 16(free), 00.00%(used)
:SNet 19: 16(size), 16(free), 00.00%(used)
:SNet 20: 16(size), 16(free), 00.00%(used)
:SNet 21: 16(size), 16(free), 00.00%(used)
:SNet 22: 16(size), 16(free), 00.00%(used)
:SNet 23: 16(size), 16(free), 00.00%(used)
:SNet 24: 16(size), 16(free), 00.00%(used)

```

9 CAM partition profiles

```
:SNet 25: 16(size), 16(free), 00.00%(used)
:SNet 26: 16(size), 16(free), 00.00%(used)
:SNet 27: 16(size), 16(free), 00.00%(used)
:SNet 28: 16(size), 16(free), 00.00%(used)
:SNet 29: 16(size), 16(free), 00.00%(used)
:SNet 30: 16(size), 16(free), 00.00%(used)
:SNet 31: 512(size), 512(free), 00.00%(used)

[MAC]131072(size), 131061(free), 00.00%(used)
:Protocol: 10(size), 6(free), 40.00%(used)
:Forwarding:131054(size), 131047(free), 00.00%(used)
:Flooding: 8(size), 8(free), 00.00%(used)

[IPV6 VPN] 65536(size), 15(free), 99.97%(used)
:SNet 0: 20(size), 0(free), 100.00%(used)
:SNet 1: 65500(size), 0(free), 100.00%(used)
:SNet 2: 2(size), 2(free), 00.00%(used)
:SNet 3: 2(size), 2(free), 00.00%(used)
:SNet 4: 2(size), 2(free), 00.00%(used)
:SNet 5: 2(size), 2(free), 00.00%(used)
:SNet 6: 2(size), 2(free), 00.00%(used)
:SNet 7: 2(size), 2(free), 00.00%(used)
:SNet 8: 2(size), 1(free), 50.00%(used)
:SNet 9: 2(size), 2(free), 00.00%(used)

[Session] 32768(size), 32767(free), 00.00%(used)
:IP Multicast: 8192(size), 8192(free), 00.00%(used)
:Receive ACL: 1024(size), 1023(free), 00.09%(used)
:Rule ACL: 23552(size), 23552(free), 00.00%(used)
:IP Source Guard Permit: 0(size), 0(free), 00.00%(used)
:IP Source Guard Denial: 0(size), 0(free), 00.00%(used)

[IPV6 Session] 8192(size), 8192(free), 00.00%(used)
:IP Multicast: 2048(size), 2048(free), 00.00%(used)
:Receive ACL: 0(size), 0(free), 00.00%(used)
:Rule ACL: 6144(size), 6144(free), 00.00%(used)

[Internal Forwarding Lookup] 8192(size), 8185(free), 00.08%(used)

[Out Session] 28672(size), 28672(free), 00.00%(used)

[Out V6 Session] 8192(size), 8192(free), 00.00%(used)
```

The type of CAM partitioning profile configured is displayed in the “CAM partitioning profile line. The “multi-service-3” or “multi-service-4” profile indicates that the system will allocate a partition for IPV6 VPN.

The output displays the size of the available CAM, amount of CAM currently free, and what percentage of the available CAM is used currently.

(size): The effective user size obtained by subtracting the reserved size from the user size.

(free): The amount of CAM currently available.

(used): The percentage of CAM currently being used.

Syntax: `show cam-partition usage slot slot-number`

Displaying CAM information

The following commands display CAM information.

Show cam l2vpn

To display all VLL or VPLS MAC entries, including local entries (Port or VLAN or MAC from end points) and remote entries (VC or MAC from VLL or VPLS peers) enter the following command.

```
Brocade# show cam l2vpn 2/1
Slot Index      MAC                Age  Port  IFL/  VC Label  Out Port Remote DA/ PRAM
      (Hex)                Dis  2/4  VLAN                Filter 0    SA  (Hex)
2     9fff6     0000.0034.5678  Dis  2/4  4096   74565    2/2    0    DA  8f
2     9fff7     0000.0034.5566  Dis  2/2   500    N/A      Filter 0    DA  8e
```

Syntax: `show cam l2vpn slot/port [MAC address]`

Show cam ipvpn

To display IPv4 VPN CAM entries, including local (Port+VLAN+IP) and remote (VC+IP) entries, enter the following command.

```
Brocade# show cam ipvpn 2/1
Slot Index  IP_Address      Port  Vlan  VC Lbl MAC                Age Out Vlan Out Port
2     0x60000  10.2.3.4/32    2/6   18   N/A   N/A                Dis 10      3/5
2     0x60001  224.7.8.9/32  N/A   N/A   4660  0000.0080.0600  Dis 20      3/5
```

Syntax: `show cam ipvpn slot/port [IP prefix]`

Show cam l4

To display all CAM entries on a Layer 4 interface, enter the following command.

```
Brocade# show cam l4 4/1
LP Index Src IP          SPort Pro Age IFL/  Out IF Group PRAM
      (Hex) (Dest IP      DPort)          VLAN Action      (Hex)
4     a4000  0.0.0.0      0     17  Dis 0    CPU   31    00084
      (10.0.0.0  3784 )
4     a4800  0.0.0.0      0     0   Dis 0    Pass  0     000c1
      (10.9.4.255  0   )
4     a4802  0.0.0.0      0     0   Dis 0    Pass  0     000c2
      (10.10.4.255 0   )
4     a4804  0.0.0.0      0     0   Dis 0    Pass  0     000c3
      (10.33.33.255 0   )
4     a4806  0.0.0.0      0     0   Dis 0    Pass  0     000c4
      (10.10.10.255 0   )
4     a4808  0.0.0.0      0     0   Dis 0    Pass  0     000c5
      (10.20.20.255 0   )
4     a480a  0.0.0.0      0     0   Dis 0    Pass  0     000c6
      (10.13.13.255 0   )
4     a480c  0.0.0.0      0     0   Dis 0    Pass  0     000c7
      (10.41.41.255 0   )
4     a480e  0.0.0.0      0     0   Dis 0    Pass  0     000c8
      (10.21.21.21  0   )
4     a4810  0.0.0.0      0     0   Dis 0    Pass  0     000c9
      (10.55.55.255 0   )
```

Syntax: `show cam I4 slot/port`

The *slot/port* parameter specifies the port for which you want to display the CAM entries.

Table 53 describes the output parameters of the `show cam I4 slot/port` command.

TABLE 53 Output parameters of the `show cam I4` command

Field	Description
LP	Shows the number of the interface module.
Index (Hex)	Shows the row number of this entry in the IP route table.
Src IP Dest IP	Shows the source IP address and the destination IP address.
SPort DPort	Shows the source port ID and the destination port ID.
Pro	Shows the type of the protocol (TCP, UDP) used.
Age	Shows whether the age is enabled or disabled.
IFL/ VLAN	Shows the VLAN to which the port belongs.
Out IF Action	Shows the state of outgoing interface action.
Group	Shows the group address.
PRAM (Hex)	Shows the ACL PRAM entries.

Show cam label-out

To display Outbound Label ACL CAMs, enter the following command.

```
Brocade# show cam label-out 2/1
Slot Index   Port  Outer Lbl  Inner Lbl  MAC           Action
2      0xc0000  2/1   1024      1025      0000.0034.5678 Drop
2      0xc0002  2/1   1027      1028      0000.0034.5577 Drop
```

Syntax: `show cam label-out slot/port`

Show IFL CAM partition

To display information about the IFL CAM partition, enter the following command.

```
Brocade#show cam ifl 2/1
Slot Index   Port  Outer VLAN  Inner VLAN  PRAM   IFL ID
      (Hex)
2      00c5fff  2/1   100         200         185fff  4096
```

Syntax: `show cam ifl slot/port`

Show CAM IP

To display IP CAM information, enter the following command.

```
Brocade# show cam ip 3/1
LP Index      IP Address      MAC              Age IFL/  Out IF  PRAM
  (Hex)                VLAN              (Hex)
3  02fef(L)  10.33.32.0/32   N/A              Dis N/A  Drop    00094
3  02ff0(L)  10.33.32.255/32 N/A              Dis N/A  Mgmt    0009d
3  02ff1(L)  10.33.32.1/32   N/A              Dis N/A  Mgmt    0009c
3  02ff2(L)  10.11.11.0/32   N/A              Dis N/A  Drop    00094
3  02ff3(L)  10.11.11.255/32 N/A              Dis N/A  Mgmt    0009b
3  02ff4(L)  10.11.11.3/32   N/A              Dis N/A  Mgmt    0009a
3  02ff5(L)  10.5.5.5/32     N/A              Dis N/A  Mgmt    00096
3  02ff6(L)  224.0.0.22/32   N/A              Dis N/A  Mgmt    00093
3  02ff7(L)  224.0.0.18/32   N/A              Dis N/A  Mgmt    00092
3  02ff8(L)  224.0.0.13/32   N/A              Dis N/A  Mgmt    00091
3  02ff9(L)  224.0.0.9/32    N/A              Dis N/A  Mgmt    00090
3  02ffa(L)  224.0.0.6/32    N/A              Dis N/A  Mgmt    0008f
3  02ffb(L)  224.0.0.5/32    N/A              Dis N/A  Mgmt    0008e
3  02ffc(L)  224.0.0.4/32    N/A              Dis N/A  Mgmt    0008d
3  02ffd(L)  224.0.0.2/32    N/A              Dis N/A  Mgmt    0008c
3  02ffe(L)  224.0.0.1/32    N/A              Dis N/A  Mgmt    0008b
3  02fff(L)  10.255.255.255/32 N/A              Dis N/A  Mgmt    0008a
3  35488(R)  10.33.32.0/24   N/A              Dis N/A  CPU     0009f
3  35489(R)  10.11.11.0/24   N/A              Dis N/A  CPU     0009e
3  3548a(R)  10.5.5.5/32     N/A              Dis N/A  Drop    00094
3  3ffff(R)  0.0.0.0/0       N/A              Dis N/A  Drop    00094
```

Syntax: `show cam ip slot/port`

The `slot/port` parameter specifies the port for which you want to display the CAM entries.

[Table 54](#) describes the output parameters of the `show cam ip slot/port` command.

TABLE 54 Output parameters of the `show cam ip` command

Field	Description
LP	Shows the number of the interface module.
Index (Hex)	Shows the row number of this entry in the IP route table.
IP Address	Shows the IP address of the interface.
MAC	Shows the MAC address of the interface.
Age	Shows whether the age is enabled or disabled.
IFL/ VLAN	Shows the VLAN to which the port belongs.
Out IF	Shows the state of outgoing interface action.
PRAM (Hex)	Shows the ACL PRAM entries.

Show CAM IPv6

To display IPv6 CAM information, enter the following command

```
Brocade# show cam ipv6 3/20
LP Index      IPV6 Address      MAC              Age IFL/ Out IF PRAM
  (Hex)                               (Hex)
 3 22ffc 2001:db8::/128   N/A              Dis N/A  Mgmt  000dc
 3 22ffe 2001:db8::1/128  N/A              Dis N/A  Mgmt  000db
 3 2e8a6 2001:db8::/64    N/A              Dis N/A  CPU   000dd
 3 2ffde fe80::/10        N/A              Dis N/A  CPU   00086
 3 2fffe ::/0             N/A              Dis N/A  Drop  00085
```

Syntax: `show cam ipv6 slot/port`

The `slot/port` parameter specifies the port for which you want to display the CAM entries.

Table 55 describes the output parameters of the `show cam ipv6 slot/port` command.

TABLE 55 Output parameters of the `show cam ipv6` command

Field	Description
LP	Shows the number of the interface module.
Index (Hex)	Shows the row number of this entry in the IP route table.
IPV6 Address	Shows the IPv6 address of the interface.
MAC	Shows the MAC address of the interface.
Age	Shows whether the age is enabled or disabled.
IFL/ VLAN	Shows the VLAN to which the port belongs.
Out IF	Shows the state of outgoing interface action.
PRAM (Hex)	Shows the ACL PRAM entries.

Displaying IPv6 VPN CAM information

The `show cam ipv6-vpn` command displays CAM information for an IPv6 VPN CAM entry on a single port, or for all ports on a device. IPv6 VPN CAM contains the destination IPv6 VPN address and layer 3 VPN ID. To display information for an IPv6 VPN CAM entry, enter the following command:

```
Brocade# show cam ipv6-vpn 1/1
LP Index IPV6 VPN Address      MAC              Age
  (Hex)                               IFL ID    Out IF    PRAM
 1 407f0 2001:db8:1::/128   N/A              Dis
                               (21847    Filter    1d615)
 1 407f2 2001:db8:2::/128   N/A              Dis
                               (21846    Drop     5af6d)
```

Syntax: `show cam ipv6-vpn slot/port`

Show cam v6acl

The **show cam v6acl** command displays IPv6 ACL CAM sessions configured on the device. The VLAN column is expanded to display either VLAN or IFL ID as shown in the example below:

```
Brocade# show cam v6acl 1/1
LP Index Src IP Addr          SPort IFL/VLAN ID
        Dst IP Addr          DPort Pro Age Out IF PRAM
1  74000 2001:db8:1::/64             0      536977
        2001:db8:2::/64             0      6   Dis Pass  000a4
1  74008 2001:db8:1::/64             0      536977
        2001:db8:2::/64             0      6   Dis Pass  000a5
1  74010 2001:db8:1::/64             0      536977
        2001:db8:2::/64             0      6   Dis Pass  000a6
1  74018 2001:db8:1::/64             0      536977
        2001:db8:2::/64             0      6   Dis Pass  000a7
1  74020 2001:db8:1::/64             0      536977
        2001:db8:2::/64             0      6   Dis Pass  000a8
1  74028 2001:db8:1::/64             0      536977
```

Syntax: **show cam ipv6-vpn slot/port**

Displaying IPv6 host drop CAM limit

Run the **show ipv6** command to display information about the IPv6 host drop CAM limit.

```
Brocade# show ipv6
Global Settings
  IPv6 Router-Id: 10.23.23.1  load-sharing path: 4
  unicast-routing enabled, ipv6 allowed to run, hop-limit 64
  reverse-path-check disabled
  host drop cam limit 5
  urpf-exclude-default disabled
  session-logging-age 5
  No Inbound Access List Set
  No Outbound Access List Set
  source-route disabled, forward-source-route disabled, icmp-redirect disabled
Configured Static Routes: 2
```

Syntax: **show ipv6**

Show IFL CAM ISID partition

To display information about 802.1AH for ISID, enter the following command:

```
Brocade#show cam ifl-isid 1/1
Slot Index   Port   Outer VLAN Itag ISID   PRAM   IFL ID   IPV4/V6
  (Hex)                               (Hex)   Routing
1     0085fe8 1/14    27     37    185fe8 1       0/0
1     0085fe9 1/13    26     36    185fe9 1       0/1
1     0085fea 1/16    25     35    185fea 1       1/0
1     0085feb 1/15    24     34    185feb 1       1/1
```

Syntax: `show cam ifl-isid slot/port`

This output includes an IPv4/ IPv6 Routing column. The IPv4/IPv6 Routing column indicates whether IPv4 or IPv6 is enabled or disabled on the interface. The number 1 represents enabled, and the number 0 represents disabled. For example, if 0/0 is displayed, then IPv4/IPv6 is disabled. If 0/1 is displayed, then IPv4 is disabled/ IPv6 is enabled. The IPv4/IPv6 Routing column is also displayed in the output of the `show cam ifl` command and `show cam ifl-mpls` command.

Configuring CAM partition size

When you configure a tftp file size into the device, the device can only perform a parameter check based on the default CAM profile configured. In this situation, it is possible that you have configured a CAM partition size that conflicts with the physical CAM size. The following `system-max` commands may cause a conflict with the physical CAM size:

- `system-max`
 - `ifl-cam`
 - `ip-source-guard-cam`
 - `ipv4-mcast-cam`
 - `ipv6-mcast-cam`
 - `lsp-out-acl-cam`
 - `subnet-broadcast-acl-cam`
 - `receive cam`

When you have configured a CAM size that conflicts with the physical CAM size, a partition is created with the maximum possible CAM indices assigned to it. The following Syslog message is generated:

```
Brocade# show log
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

Dynamic Log Buffer (50 lines):
Sep  9 18:48:23:A:CAM IPv6 VPN SNet 9 partition warning: request 32, actual 0,
s1
ot 1, ppcr 0
```

CAM overflow logging

At system initialization, a threshold value is calculated for each sub-partition. If a partition does not have any sub-partitions, the value is based on the entire partition size. If a partition has movable sub-partition boundaries, the threshold value is also based on the entire partition size. By default, the threshold value is 5% of the total entry count. A minimum logging interval (default of 5 minutes) is also set for each partition to check usage. For example, let us say CAM overflow logging duration was set to 5 minutes and the overflow log is generated during a CAM write at 2:00 pm, then any further CAM writes will not cause an overflow log until 5 minutes have elapsed. So the next CAM overflow logging would occur on a CAM write after 2:05 pm. When the interval elapses, if the number of unused CAM entries drops below the threshold percentage value, a log message is generated during a CAM write.

```
CAM partition <partition name including sub-partition ID if applicable> warning:  
total <total count>, free <current free count>, slot <1 based slot number>, ppcr  
<0 based ppcr id>
```

After the log message is generated, the sub-partition time stamp is updated to the current time.

Configuring minimum logging interval and threshold value

You can configure a minimum logging interval and threshold value for CAM partition logging using the following command.

```
Brocade(config)# cam-partition logging 10% 5
```

Syntax: [no] **cam-partition logging** *threshold percentage %* | *interval in minutes*

You can configure the *threshold percentage %* variable to change the threshold value from the default 95%.

The *interval in minutes* variable allows you to set the minimum logging interval. Default 5 minutes.

NOTE

Because IP and IPv6 sub-partitions can dynamically grow and shrink, for these partitions, logging is implemented at the entire partition level. An SNMP trap is generated with the logging message.

9 CAM overflow logging

Using Syslog

Table 56 displays the individual Brocade devices and the Syslog features they support.

TABLE 56 Supported Brocade Syslog features

Features supported	Brocade Netron XMR Series	Brocade MLX Series	Brocade Netron CES 2000 Series BASE package	Brocade Netron CES 2000 Series ME_PREM package	Brocade Netron CES 2000 Series L3_PREM package	Brocade Netron CER 2000 Series Base package	Brocade Netron CER 2000 Series Advanced Services package
Syslog Messages	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Real-Time Display of Syslog Messages	Yes	Yes	Yes	Yes	Yes	Yes	Yes
BFD Syslog	Yes	Yes	No	Yes	Yes	Yes	Yes
Increased Syslog buffer	Yes	Yes	No	No	No	No	No
Time Stamps	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Option for Show Log Command	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Disabling Logging of a Message Level	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Interface Name in Syslog Messages	Yes	Yes	Yes	Yes	Yes	Yes	Yes
TCP or UDP Port Numbers in Syslog Messages	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Logging all CLI Commands	Yes	Yes	Yes	Yes	Yes	Yes	Yes
BFD Logging	Yes	Yes	No	Yes	Yes	Yes	Yes

TABLE 56 Supported Brocade Syslog features

Features supported	Brocade Netron XMR Series	Brocade MLX Series	Brocade Netron CES 2000 Series BASE package	Brocade Netron CES 2000 Series ME_PREM package	Brocade Netron CES 2000 Series L3_PREM package	Brocade Netron CER 2000 Series Base package	Brocade Netron CER 2000 Series Advanced Services package
Number of Entries the Local Buffer Can Hold	1-5000	1-5000	1-5000	1-5000	1-5000	1-5000	1-5000
Disabling Syslog for an event	Yes	Yes	Yes	Yes	Yes	Yes	Yes

This appendix describes how to display Syslog messages and how to configure the Syslog facility, and lists the Syslog messages that a Brocade 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 Brocade 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.

```
Brocade# 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 330.

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.

```
Brocade(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@Brocade# 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.

A Configuring the Syslog service

```
telnet@Brocade# terminal monitor
Syslog trace was turned OFF
```

Here is an example of how the Syslog messages are displayed.

```
telnet@Brocade# terminal monitor
Syslog trace was turned ON
SYSLOG: <9>Brocade, Power supply 2, power supply on left connector, failed

SYSLOG: <14>Brocade, Interface ethernet 1/6, state down

SYSLOG: <14>Brocade, 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.

```
Brocade> 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 57 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 337. 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 340.
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.

A Configuring the Syslog service

```
Brocade(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.

```
Brocade# 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.

```
Brocade(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.

A Configuring the Syslog service

```
Brocade(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:

1. 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:

```
Brocade# copy tftp flash 10.25.101.121 cert.p12 client-certificate
Brocade# 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_ip> <priv_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

1. 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:

```
Brocade (config)# logging host 10.25.105.201 ssl-port 60514
```

Syntax: logging host [ipv6] <ip-address> | <ipv6-address> ssl-port <port>

The **ip-address** keyword specifies the syslog server. The **ssl-port** keyword 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:

```
Brocade# 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.

```
Brocade# show logging
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
```

A Configuring the Syslog service

```
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

```
Brocade# 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.

```
Brocade(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.

```
Brocade(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 particular event. In the following example, the **nologging enable** command disables the syslog for SNMP authentication failure.

```
Brocade (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 **vrp-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

```
Brocade(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-addr* | *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

```
Brocade(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-addr* | *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

```
Brocade(config)# no logging buffered debugging
```

A Configuring the Syslog service

```
Brocade(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.

Example

```
Brocade(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.

```
Brocade(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.

```
Brocade(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.

```
Brocade# 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
```

A Syslog messages

```
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.

```
Brocade# 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.

```
Brocade(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@ Brocade-2(config)# logging cli-command  
telnet@ Brocade-2(config)# router bgp  
telnet@ Brocade-2(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.

```
Brocade-2(config)# router bgp  
Brocade-2(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.

```
Brocade(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

TABLE 58 Syslog messages system

Message level	Message	Explanation
Alert	ISIS Memory Limit Exceeded	IS-IS is requesting more memory than is available.
Alert	System Power supply <i>num</i> , <i>location</i> , failed	A power supply has failed. The <i>num</i> is the power supply number. The <i>location</i> describes where the failed power supply is in the chassis.
Alert	System <i>power type</i> Power Supply <i>num</i> , <i>location</i> , <i>state</i>	The <i>power type</i> refers to AC or DC power supply. The <i>num</i> is the power supply number as positioned in the chassis. The <i>location</i> describes where the power supply is in the chassis in relation to its state. The <i>state</i> refers to how the power supply is functioning in the chassis. The <i>state</i> can be one of the following: <ul style="list-style-type: none"> • Installed (OK) - The power supply is installed and ok. • 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.
Alert	System Fan <i>num</i> , <i>location</i> , failed	A fan has failed. The <i>num</i> is the power supply number. The <i>location</i> describes where the failed power supply is in the chassis. The location can be one of the following

TABLE 58 Syslog messages system (Continued)

Message level	Message	Explanation
Alert	System Management module at slot <i>slot-num</i> state changed from <i>module-state</i> to <i>module-state</i> due to <i>reason</i> .	<p>Indicates a state change in a management module.</p> <p>The <i>slot-num</i> indicates the chassis slot containing the module.</p> <p>The <i>module-state</i> can be one of the following:</p> <ul style="list-style-type: none"> • active • standby • crashed • coming-up • unknown <p>A due to clause has been added to this message. The <i>reason</i> variable can be either or the following:</p> <ul style="list-style-type: none"> • MP upgrade to ver <i>version number</i> where <i>version number</i> is the version number of the Multi-Service IronWare software that the management module was upgraded to. • Active Reboot
Alert	System Temperature <i>degrees</i> C degrees, warning level <i>warn-degrees</i> C degrees, shutdown level <i>shutdown-degrees</i> C degrees	<p>Indicates an overtemperature condition on the active module.</p> <p>The <i>degrees</i> value indicates the temperature of the module.</p> <p>The <i>warn-degrees</i> value is the warning threshold temperature configured for the module.</p> <p>The <i>shutdown-degrees</i> value is the shutdown temperature configured for the module.</p>
Alert	System <i>num-modules</i> modules and 1 power supply, need more power supply!	<p>Indicates that the chassis needs more power supplies to run the modules in the chassis.</p> <p>The <i>num-modules</i> parameter indicates the number of modules in the chassis.</p>
Alert	System: Health Monitoring: FE access failure detected on SFM <i>num</i> /FE <i>num</i> (NetIron XMR and NetIron MLX only)	<p>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.</p> <p>The SFM and FE <i>num</i> parameters indicate the number of the switch fabric module and fabric element that could not be accessed</p>
Alert	System: Health Monitoring: TM Egress data errors detected on LP <i>num</i> /TM <i>num</i>	<p>The system has detected egress data errors on the specified line processor and traffic manager.</p> <p>The LP and TM <i>num</i> parameters indicate the number of the line processor and traffic manager on which the errors were detected.</p>
Error	Error Failed to shutdown Power Supply <i>PS-Num</i> . Write Failed (offset 0x2, value 44, size 2).(Brocade NetIron XMR and Brocade MLX only).	<p>A power supply failed to shutdown because of its failure to access its registers.</p>

TABLE 58 Syslog messages system (Continued)

Message level	Message	Explanation
Error	Error Module down in slot 3, reason CARD_DOWN_REASON_BOOT_FAILED.Err or Code (1).	<ul style="list-style-type: none"> 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. <p>The error code is 0 when there is no error code reported from the Interface Module.</p>
Warning	CAM partition <i>partition name</i> warning total <i>total-count</i> , free current <i>free-count</i> , slot <i>slot-number</i> , ppcr <i>ppcr-id</i>	<p>Indicates that the CAM partition specified by the <i>partition name</i> 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 <i>free-count</i> specifies the amount of free space still available in the partition. The <i>slot-number</i> and ppcr <i>ppcr-id</i> indicate where the overflow is occurring. The <i>partition-name</i> includes the sub-partition ID if applicable.</p>
Warning	System Not enough power to power on module in slot <i>num</i>	<p>There is not enough power available in the chassis to power on the module in the specific slot number.</p> <p>The slot <i>num</i> refers to the slot number in the chassis.</p>
Notification	System Module up in slot <i>n</i>	
Notification	System Module down in slot <i>n</i> reason	<p>Indicates that the module in the slot specified by the <i>n</i> variable is down for one of the following reasons as specified by the <i>reason</i> variable:</p> <ul style="list-style-type: none"> 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
Notification	System Module <i>n</i> powered on	
Notification	System Module <i>n</i> powered off	
Notification	System Switch fabric <i>n</i> powered on	
Notification	System Switch fabric <i>n</i> powered off	
Notification	System Enough power available to power on module in slot <i>num</i> .	<p>There is enough power available in the chassis to power on the module in the specific slot number. The slot <i>num</i> refers to the slot number in the chassis.</p>

TABLE 58 Syslog messages system (Continued)

Message level	Message	Explanation
Notification	System Module was inserted to slot <i>slot-num</i>	Indicates that a module was inserted into a chassis slot. The <i>slot-num</i> is the number of the chassis slot into which the module was inserted.
Notification	System Module was removed from slot <i>slot-num</i>	Indicates that a module was removed from a chassis slot. The <i>slot-num</i> is the number of the chassis slot from which the module was removed.
Notification	System Set fan speed to <i>speed percentage</i>	Indicates that the fan speed has been changed to the value described in the <i>speed</i> variable and that the fan is now operating at the <i>percentage</i> of capacity described. The possible <i>speed percentage</i> values are: <ul style="list-style-type: none"> • LOW (50%) • MEDIUM (75%) • MEDIUM-HIGH (90%) • HIGH (100%)
Informational	System Cold start	The device has been powered on.
Informational	System Warm start	The system software (flash code) has been reloaded.
Informational	System Interface <i>portnum</i> , state up	A port has come up. The <i>portnum</i> is the port number.
Informational	System Interface <i>portnum</i> , state down	A port has gone down. The <i>portnum</i> is the port number.
Informational	System Interface <i>portnum</i> , line protocol up	The line protocol on a port has come up. The <i>portnum</i> is the port number.
Informational	System Interface <i>portnum</i> , line protocol down	The line protocol on a port has gone down. The <i>portnum</i> is the port number.
Informational	System Interface <i>portnum</i> is down (remote fault)	The interface is down due to Remote Fault. This is indicated as "(remote fault)". The <i>portnum</i> is the port number of the interface.
Informational	System <i>portnum</i> is down (local fault)	The port is down due to Local Fault. This is indicated as "(local fault)". The <i>portnum</i> is the port number of the interface.
Informational	System Syslog server <i>IP-address</i> deleted added modified from console telnet ssh web snmp OR Syslog operation enabled disabled from console telnet ssh web snmp	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.
Informational	System SSH telnet server enabled disabled from console telnet ssh web snmp session [by user <i>username</i>]	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.
Informational	System Module <i>n</i> CPU <i>m</i> crashed	
Informational	System IfIndex assignment was changed.)	The maximum number of ifIndex per module has been changed.

TABLE 58 Syslog messages system (Continued)

Message level	Message	Explanation
Informational	System Power Supply <i>PS-Num</i> will be shutdown due to flapping next time it becomes available. (Brocade NetIronXMR and Brocade MLX only).	A power supply will shutdown because of flapping the next time it is available. The <i>PS-Num</i> is the power supply number.
Informational	System Power Supply <i>PS-Num</i> is shutdown due to flapping.(Brocade XMR and Brocade MLX only).	A power supply is shut down because of flapping. The <i>PS-Num</i> is the power supply number.

TABLE 59 Syslog messages security

Message level	Message	Explanation
Warning	Security Port security violation at interface <i>portnum</i> , address <i>mac</i> , <i>vlan id</i>	
Warning	Security Interface <i>portnum</i> was shut down due to port security violation	
Informational	Security console login {by <i>userInull</i> } to USER EXEC mode Security {telnet ssh} login {by <i>userInull</i> } from src {IP <i>ip</i> IPv6 <i>ipv6-addr</i> } to USER EXEC mode	A user has logged into the USER EXEC mode of the CLI. The <i>user</i> is the user name.
Informational	Security console logout {by <i>userInull</i> } from USER EXEC mode Security {telnet ssh} logout {by <i>userInull</i> } from src {IP <i>ip</i> IPv6 <i>ipv6-addr</i> } from USER EXEC mode	A user has logged out of the USER EXEC mode of the CLI. The <i>user</i> is the user name.
Informational	Security console login {by <i>userInull</i> } to Privileged EXEC mode Security {telnet ssh} login {by <i>userInull</i> } from src {IP <i>ip</i> IPv6 <i>ipv6-addr</i> } to Privileged EXEC mode	A user has logged into the Privileged EXEC mode of the CLI. The <i>user</i> is the user name.
Informational	Security console logout {by <i>userInull</i> } from Privileged EXEC mode Security {telnet ssh} logout {by <i>userInull</i> } from src {IP <i>ip</i> IPv6 <i>ipv6-addr</i> } from Privileged EXEC mode	A user has logged out of Privileged EXEC mode of the CLI. The <i>user</i> is the user name.
Informational	Security outbound telnet <i>session number</i> login to server IP <i>ip</i> from SSH session <i>session number</i>	A user has initiated an outbound Telnet session from an inbound SSH session. The first <i>session number</i> is the number of the outbound Telnet session. The <i>ip</i> is the IP address to which the Telnet session is connected. The second <i>sessions number</i> is the number of the inbound SSH session.

TABLE 59 Syslog messages security (Continued)

Message level	Message	Explanation
Informational	Security outbound telnet <i>session number</i> logout from server IP <i>ip</i> from SSH session <i>session number</i>	A user has terminated an outbound Telnet session initiated from an inbound SSH session. The first <i>session number</i> is the number of the outbound Telnet session. The <i>ip</i> is the IP address from which the Telnet session has disconnected. The second <i>sessions number</i> is the number of the inbound SSH session.
Informational	Security startup-config was changed {by <i>user</i> } from {web management snmp management ssh client <i>ipl</i> telnet client <i>ip</i> }	A configuration change was saved to the startup configuration file. The <i>user</i> is the user's ID, if they entered a user ID to log in.
Informational	Security running-config was changed {by <i>user</i> } from {web management snmp management ssh client <i>ip</i> telnet client <i>ip</i> }	A configuration change was saved to the running configuration file. The <i>user</i> is the user's ID, if they entered a user ID to log in.
Informational	Security telnet SSH web access [by <i>username</i>] from src IP <i>source ip address</i> , src MAC <i>source MAC address</i> rejected, <i>n</i> attempts	There were failed web, SSH, or Telnet login access attempts from the specified source IP and MAC address. <ul style="list-style-type: none"> [by <i>user username</i>] does not appear if telnet or SSH clients are specified. <i>n</i> is the number of times this SNMP trap occurred in the last five minutes, or other configured number of minutes.
Informational	Security user <i>username</i> added deleted modified from console telnet ssh web snmp	A user created, modified, or deleted a local user account through the Web, SNMP, console, SSH, or Telnet session.
Informational	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	A user created, re-configured, or deleted an Enable or Line password through the Web, SNMP, console, SSH, or Telnet session.

TABLE 60 Syslog messages VLAN

Message level	Message	Explanation
Informational	VLAN Id <i>vlan-id</i> added deleted modified from console telnet ssh web snmp session	A user created, modified, or deleted a VLAN through the Web, SNMP, console, SSH, or Telnet session.

TABLE 61 Syslog messages STP

Message level	Message	Explanation
Informational	STP VLAN <i>id</i> - New RootBridge <i>string</i> RootPort <i>portnum</i> (<i>reason</i>)	A Spanning Tree Protocol (STP) topology change has occurred. The <i>id</i> is the ID of the VLAN in which the STP topology change occurred. The <i>portnum</i> is the number of the port connected to the new root bridge.
Informational	STP VLAN <i>id</i> - Bridge is RootBridge <i>string</i> (<i>reason</i>)	A Spanning Tree Protocol (STP) topology change has occurred, resulting in the device becoming the root bridge. The <i>id</i> is the ID of the VLAN in which the STP topology change occurred.
Informational	STP VLAN <i>id</i> Port <i>portnum</i> - Bridge TC Event (<i>reason</i>)	A Spanning Tree Protocol (STP) topology change has occurred on a port. The <i>id</i> is the ID of the VLAN in which the STP topology change occurred. The <i>portnum</i> is the port number.
Informational	STP VLAN <i>vlanid</i> Port <i>portnum</i> - State <i>state</i> (<i>reason</i>)	
Informational	STP Root Guard Port <i>portnum</i> , VLAN <i>vlan-id</i> inconsistent (Received superior BPDU)	The specified port was blocked because it has Root Guard enabled and received a superior BPDU.
Informational	STP Root Guard Port <i>portnum</i> , VLAN <i>vlan-id</i> consistent (Timeout)	The specified block Root Guard-protected port was unblocked.
Informational	STP BPDU Guard port <i>portnum</i> disable System Interface ethernet <i>portnum</i> , state down - disabled	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.
Informational	STP BPDU Guard re-enabled on ports ethe <i>portnum</i> System Interface ethernet <i>portnum</i> , state up	The spanning-tree protect re-enable was issued to re-enable the specified port

TABLE 62 Syslog messages RSTP

Message level	Message	Explanation
Informational	RSTP VLAN <i>id</i> Port <i>portnum</i> - Bridge TC Event (<i>reason</i>)	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.
Informational	RSTP VLAN <i>id</i> Port <i>portnum</i> - STP State <i>state</i> (<i>reason</i>)	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.
Informational	RSTP VLAN <i>id</i> - New RootPort <i>portnum</i> (<i>reason</i>)	802.1W changed the port's role to Root port, using the root selection computation.
Informational	RSTP VLAN <i>id</i> - New RootBridge <i>string</i> RootPort <i>portnum</i> (<i>reason</i>)	802.1W selected a new root bridge as a result of the BPDUs received on a bridge port.

TABLE 62 Syslog messages RSTP (Continued)

Message level	Message	Explanation
Informational	RSTP VLAN <i>id</i> - Bridge is RootBridge <i>string</i> (reason)	802.1W changed the current bridge to be the root bridge of the given topology due to administrative change in bridge priority.
Informational	vlan <i>vlan-id</i> Bridge is RootBridge <i>mac-address</i> (MsgAgeExpiry)	The message age expired on the Root port so 802.1W changed the current bridge to be the root bridge of the topology.

TABLE 63 Syslog messages LAG

Message level	Message	Explanation
Informational	LAG group (<i>ports</i>) created by 802.3ad link-aggregation module.	802.3ad link aggregation is configured on the device, and the feature has dynamically created a LAG group (aggregate link). The <i>ports</i> is a list of the ports that were aggregated to make the LAG group.

TABLE 64 Syslog messages MRP

Message level	Message	Explanation
Informational	MRP interface ethernet <i>portnum</i> vlan <i>vlan-master</i> , changing to <i>state-string</i>	
Informational	MRP metro ring <i>ring-id</i> cannot be enabled. No free CAM entries	

TABLE 65 Syslog messages UDLD

Message level	Message	Explanation
Informational	UDLD Logical link on interface ethernet <i>portnum</i> is up	
Informational	UDLD Logical link on interface ethernet <i>portnum</i> is down	

TABLE 66 Syslog messages VSRP

Message level	Message	Explanation
Informational	VSRP VLAN <i>vlanid</i> VRID <i>id</i> - transition to <i>state-string</i>	
Informational	VSRP VLAN <i>vlanid</i> VRID <i>id</i> - aware change <i>old-portnum</i> -> <i>new-portnum</i> \n	
Informational	VSRP VLAN <i>vlanid</i> VRID <i>id</i> - aware learn <i>portnum</i>	

TABLE 67 Syslog messages VRRP

Message level	Message	Explanation
Notification	VRRP intf state changed, intf <i>portnum</i> , vrid <i>virtual-router-id</i> , state <i>vrrp-state</i>	A state change has occurred in a Virtual Router Redundancy Protocol (VRRP) interface. The <i>portnum</i> is the port. The <i>virtual-router-id</i> is the virtual router ID (VRID) configured on the interface. The <i>vrrp-state</i> can be one of the following: <ul style="list-style-type: none"> • init • master • backup • unknown

TABLE 68 Syslog messages IP

Message level	Message	Explanation
Warning	IP Dup IP <i>ip-addr</i> detected, sent from MAC <i>mac-addr</i> interface <i>portnum</i>	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 <i>ip-addr</i> is the duplicate IP address. The <i>mac-addr</i> is the MAC address of the device with the duplicate IP address. The <i>portnum</i> is the port that received the packet with the duplicate IP address. The address is the packet's source IP address.

TABLE 69 Syslog messages ICMP

Message level	Message	Explanation
Notification	ICMP Local ICMP exceeds <i>burst-max</i> burst packets, stopping for <i>lockup</i> seconds!	The number of ICMP packets exceeds the <i>burst-max</i> 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 <i>lockup</i> value. When the lockup period expires, the packet counter is reset and measurement is restarted.
Notification	ICMP Transit ICMP in interface <i>portnum</i> exceeds <i>num</i> burst packets, stopping for <i>num</i> seconds!	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 <i>portnum</i> is the port number. The first <i>num</i> is the maximum burst size (maximum number of packets allowed). The second <i>num</i> 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.

TABLE 70 Syslog messages ACL

Message level	Message	Explanation
Warning	ACL list <i>acl-num</i> denied <i>ip-PROTO</i> <i>src-ip-addr</i> (<i>src-tcp/udp-port</i>) (Ethernet <i>portnum mac-addr</i>) -> <i>dst-ip-addr</i> (<i>dst-tcp/udp-port</i>), 1 events	<p>Indicates that an Access Control List (ACL) denied (dropped) packets.</p> <p>The <i>acl-num</i> indicates the ACL number. Numbers 1 – 99 indicate standard ACLs. Numbers 100 – 199 indicate extended ACLs.</p> <p>The <i>ip-PROTO</i> indicates the IP protocol of the denied packets.</p> <p>The <i>src-ip-addr</i> is the source IP address of the denied packets.</p> <p>The <i>src-tcp/udp-port</i> is the source TCP or UDP port, if applicable, of the denied packets.</p> <p>The <i>portnum</i> indicates the port number on which the packet was denied.</p> <p>The <i>mac-addr</i> indicates the source MAC address of the denied packets.</p> <p>The <i>dst-ip-addr</i> indicates the destination IP address of the denied packets.</p> <p>The <i>dst-tcp/udp-port</i> indicates the destination TCP or UDP port number, if applicable, of the denied packets.</p>
Warning	ACL:rip filter list <i>list-num</i> <i>direction</i> V1 V2 denied <i>ip-addr</i> , <i>num</i> packets	<p>Indicates that a RIP route filter denied (dropped) packets.</p> <p>The <i>list-num</i> is the ID of the filter list.</p> <p>The <i>direction</i> indicates whether the filter was applied to incoming packets or outgoing packets. The value can be one of the following:</p> <ul style="list-style-type: none"> • in • out <p>The V1 or V2 value specifies the RIP version (RIPv1 or RIPv2).</p> <p>The <i>ip-addr</i> indicates the network number in the denied updates.</p> <p>The <i>num</i> indicates how many packets matching the values above were dropped during the five-minute interval represented by the log entry.</p>
Notification	ACL insufficient L4 session resource, using flow based ACL instead	<p>The device does not have enough Layer 4 session entries.</p> <p>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 <i>num</i></p>
Notification	ACL system fragment packet inspect rate <i>rate</i> exceeded	<p>The fragment rate allowed on the device has been exceeded.</p> <p>The <i>rate</i> indicates the maximum rate allowed.</p> <p>This message can occur if fragment throttling is enabled.</p>

TABLE 70 Syslog messages ACL (Continued)

Message level	Message	Explanation
Notification	AC port fragment packet inspect rate <i>rate</i> exceeded on port <i>portnum</i>	The fragment rate allowed on an individual interface has been exceeded. The <i>rate</i> indicates the maximum rate allowed. The <i>portnum</i> indicates the port. This message can occur if fragment throttling is enabled.
Notification	ACL Port <i>portnum</i> , exceed configured L4 rule-based CAM size, larger L4 partition size required	
Notification	ACL Port <i>portnum</i> , exceed configured L2 ACL rule-based CAM size, larger partition size is required	
Notification	ACL Port <i>portnum</i> , exceed configured outbound L4 rule-based CAM size, larger outbound L4 partition size required	
Notification	ACL Port <i>portnum</i> , exceed configured IPv6 L4 rule-based CAM size, larger IPv6 L4 partition size required	
Notification	ACL Port <i>portnum</i> , exceed configured IPv6 outbound L4 rule-based CAM size, larger IPv6 outbound L4 partition size required	
Notification	ACL Port <i>portnum</i> , error in allocating inbound L4 rule-based ACL CAM entry	
Notification	ACL Port <i>portnum</i> , error in allocating outbound L4 rule-based ACL CAM entry	
Notification	ACL Port <i>portnum</i> , inbound ACL CAM programming incomplete	
Notification	ACL Port <i>portnum</i> , outbound ACL CAM programming incomplete	
Informational	ACL <i>aclid</i> added deleted modified from console telnet ssh web snmp session	A user created, modified, deleted, or applied an ACL through the Web, SNMP, console, SSH, or Telnet session.

TABLE 71 Syslog messages RACL

Message level	Message	Explanation
Notification	RACL Port <i>portnum</i> , IP Receive ACL exceed configured CAM size, larger partition size required	
Notification	RACL Port <i>portnum</i> , IP Receive ACL exceed configured RL class limit	
Notification	RACL Port <i>portnum</i> , IP Receive ACL CAM malloc error	

TABLE 72 Syslog messages OSPF

Message level	Message	Explanation
Alert	OSPF Memory Overflow	OSPF has run out of memory.
Alert	OSPF LSA Overflow, LSA Type = <i>lsa-type</i>	<p>Indicates an LSA database overflow. The <i>lsa-type</i> parameter indicates the type of LSA that experienced the overflow condition. The LSA type is one of the following:</p> <ul style="list-style-type: none"> • 1 – Router • 2 – Network • 3 – Summary • 4 – Summary • 5 – External
Notification	OSPF interface state changed, rid <i>router-id</i> , intf addr <i>ip-addr</i> , state <i>ospf-state</i>	<p>Indicates that the state of an OSPF interface has changed. The <i>router-id</i> is the router ID of the device. The <i>ip-addr</i> is the interface's IP address. The <i>ospf-state</i> indicates the state to which the interface has changed and can be one of the following:</p> <ul style="list-style-type: none"> • down • loopback • waiting • point-to-point • designated router • backup designated router • other designated router • unknown
Notification	OSPF virtual intf state changed, rid <i>router-id</i> , area <i>area-id</i> , nbr <i>ip-addr</i> , state <i>ospf-state</i>	<p>Indicates that the state of an OSPF virtual routing interface has changed. The <i>router-id</i> is the router ID of the router the interface is on. The <i>area-id</i> is the area the interface is in. The <i>ip-addr</i> is the IP address of the OSPF neighbor. The <i>ospf-state</i> indicates the state to which the interface has changed and can be one of the following:</p> <ul style="list-style-type: none"> • down • loopback • waiting • point-to-point • designated router • backup designated router • other designated router • unknown

TABLE 72 Syslog messages OSPF (Continued)

Message level	Message	Explanation
Notification	OSPF nbr state changed, rid <i>router-id</i> , nbr addr <i>ip-addr</i> , nbr rid <i>nbr-router-id</i> , state <i>ospf-state</i>	<p>Indicates that the state of an OSPF neighbor has changed.</p> <p>The <i>router-id</i> is the router ID of the device.</p> <p>The <i>ip-addr</i> is the IP address of the neighbor.</p> <p>The <i>nbr-router-id</i> is the router ID of the neighbor.</p> <p>The <i>ospf-state</i> indicates the state to which the interface has changed and can be one of the following:</p> <ul style="list-style-type: none"> • down • attempt • initializing • 2-way • exchange start • exchange • loading • full • unknown
Notification	OSPF virtual nbr state changed, rid <i>router-id</i> , nbr addr <i>ip-addr</i> , nbr rid <i>nbr-router-id</i> , state <i>ospf-state</i>	<p>Indicates that the state of an OSPF virtual neighbor has changed.</p> <p>The <i>router-id</i> is the router ID of the device.</p> <p>The <i>ip-addr</i> is the IP address of the neighbor.</p> <p>The <i>nbr-router-id</i> is the router ID of the neighbor.</p> <p>The <i>ospf-state</i> indicates the state to which the interface has changed and can be one of the following:</p> <ul style="list-style-type: none"> • down • attempt • initializing • 2-way • exchange start • exchange • loading • full • unknown

TABLE 72 Syslog messages OSPF (Continued)

Message level	Message	Explanation
Notification	OSPF intf config error, rid <i>router-id</i> , intf addr <i>ip-addr</i> , pkt src addr <i>src-ip-addr</i> , error type <i>error-type</i> , pkt type <i>pkt-type</i>	<p>Indicates that an OSPF interface configuration error has occurred. The <i>router-id</i> is the router ID of the device. The <i>ip-addr</i> is the IP address of the interface on the device. The <i>src-ip-addr</i> is the IP address of the interface from which the device received the error packet. The <i>error-type</i> can be one of the following:</p> <ul style="list-style-type: none"> • 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 <p>The <i>packet-type</i> can be one of the following:</p> <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state ack • unknown

TABLE 72 Syslog messages OSPF (Continued)

Message level	Message	Explanation
Notification	OSPF virtual intf config error, rid <i>router-id</i> , intf addr <i>ip-addr</i> , pkt src addr <i>src-ip-addr</i> , error type <i>error-type</i> , pkt type <i>pkt-type</i>	<p>Indicates that an OSPF virtual routing interface configuration error has occurred. The <i>router-id</i> is the router ID of the device. The <i>ip-addr</i> is the IP address of the interface on the device. The <i>src-ip-addr</i> is the IP address of the interface from which the device received the error packet.</p> <p>The <i>error-type</i> can be one of the following:</p> <ul style="list-style-type: none"> • 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 <p>The <i>packet-type</i> can be one of the following:</p> <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state ack • unknown

TABLE 72 Syslog messages OSPF (Continued)

Message level	Message	Explanation
Notification	OSPF intf authen failure, rid <i>router-id</i> , intf addr <i>ip-addr</i> , pkt src addr <i>src-ip-addr</i> , error type <i>error-type</i> , pkt type <i>pkt-type</i>	<p>Indicates that an OSPF interface authentication failure has occurred. The <i>router-id</i> is the router ID of the device. The <i>ip-addr</i> is the IP address of the interface on the device. The <i>src-ip-addr</i> is the IP address of the interface from which the device received the authentication failure. The <i>error-type</i> can be one of the following:</p> <ul style="list-style-type: none"> • 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 <p>The <i>packet-type</i> can be one of the following:</p> <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state ack • unknown

TABLE 72 Syslog messages OSPF (Continued)

Message level	Message	Explanation
Notification	OSPF virtual intf authen failure, rid <i>router-id</i> , intf addr <i>ip-addr</i> , pkt src addr <i>src-ip-addr</i> , error type <i>error-type</i> , pkt type <i>pkt-type</i>	<p>Indicates that an OSPF virtual routing interface authentication failure has occurred.</p> <p>The <i>router-id</i> is the router ID of the device.</p> <p>The <i>ip-addr</i> is the IP address of the interface on the device.</p> <p>The <i>src-ip-addr</i> is the IP address of the interface from which the device received the authentication failure.</p> <p>The <i>error-type</i> can be one of the following:</p> <ul style="list-style-type: none"> • 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 <p>The <i>packet-type</i> can be one of the following:</p> <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state ack • unknown

TABLE 72 Syslog messages OSPF (Continued)

Message level	Message	Explanation
Notification	OSPF intf rcvd bad pkt, rid <i>router-id</i> , intf addr <i>ip-addr</i> , pkt src addr <i>src-ip-addr</i> , pkt type <i>pkt-type</i>	<p>Indicates that an OSPF interface received a bad packet.</p> <p>The <i>router-id</i> is the router ID of the device.</p> <p>The <i>ip-addr</i> is the IP address of the interface on the device.</p> <p>The <i>src-ip-addr</i> is the IP address of the interface from which the device received the authentication failure.</p> <p>The <i>packet-type</i> can be one of the following:</p> <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state ack • unknown <p>NOTE: This message is typically generated during BFD or OSPF reconverge within the following scenarios:</p> <ul style="list-style-type: none"> • 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. <p>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.</p>
Notification	OSPF virtual intf rcvd bad pkt, rid <i>router-id</i> , intf addr <i>ip-addr</i> , pkt src addr <i>src-ip-addr</i> , pkt type <i>pkt-type</i>	<p>Indicates that an OSPF interface received a bad packet.</p> <p>The <i>router-id</i> is the router ID of the device.</p> <p>The <i>ip-addr</i> is the IP address of the interface on the device.</p> <p>The <i>src-ip-addr</i> is the IP address of the interface from which the device received the authentication failure.</p> <p>The <i>packet-type</i> can be one of the following:</p> <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state ack • unknown

TABLE 72 Syslog messages OSPF (Continued)

Message level	Message	Explanation
Notification	OSPF intf retransmit, rid <i>router-id</i> , intf addr <i>ip-addr</i> , nbr rid <i>nbr-router-id</i> , pkt type is <i>pkt-type</i> , LSA type <i>lsa-type</i> , LSA id <i>lsa-id</i> , LSA rid <i>lsa-router-id</i>	An OSPF interface on the device has retransmitted a Link State Advertisement (LSA). The <i>router-id</i> is the router ID of the device. The <i>ip-addr</i> is the IP address of the interface on the device. The <i>nbr-router-id</i> is the router ID of the neighbor router. The <i>packet-type</i> can be one of the following: <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state ack • unknown The <i>lsa-type</i> is the type of LSA. The <i>lsa-id</i> is the LSA ID. The <i>lsa-router-id</i> is the LSA router ID.
Notification	OSPF virtual intf retransmit, rid <i>router-id</i> , intf addr <i>ip-addr</i> , nbr rid <i>nbr-router-id</i> , pkt type is <i>pkt-type</i> , LSA type <i>lsa-type</i> , LSA id <i>lsa-id</i> , LSA rid <i>lsa-router-id</i>	An OSPF interface on the device has retransmitted a Link State Advertisement (LSA). The <i>router-id</i> is the router ID of the device. The <i>ip-addr</i> is the IP address of the interface on the device. The <i>nbr-router-id</i> is the router ID of the neighbor router. The <i>packet-type</i> can be one of the following: <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state ack • unknown The <i>lsa-type</i> is the type of LSA. The <i>lsa-id</i> is the LSA ID. The <i>lsa-router-id</i> is the LSA router ID.
Notification	OSPF originate LSA, rid <i>router-id</i> , area <i>area-id</i> , LSA type <i>lsa-type</i> , LSA id <i>lsa-id</i> , LSA router id <i>lsa-router-id</i>	An OSPF interface has originated an LSA. The <i>router-id</i> is the router ID of the device. The <i>area-id</i> is the OSPF area. The <i>lsa-type</i> is the type of LSA. The <i>lsa-id</i> is the LSA ID. The <i>lsa-router-id</i> is the LSA router ID.
Notification	OSPF max age LSA, rid <i>router-id</i> , area <i>area-id</i> , LSA type <i>lsa-type</i> , LSA id <i>lsa-id</i> , LSA rid <i>lsa-router-id</i>	An LSA has reached its maximum age. The <i>router-id</i> is the router ID of the device. The <i>area-id</i> is the OSPF area. The <i>lsa-type</i> is the type of LSA. The <i>lsa-id</i> is the LSA ID. The <i>lsa-router-id</i> is the LSA router ID.

TABLE 72 Syslog messages OSPF (Continued)

Message level	Message	Explanation
Notification	OSPF LSDB overflow, rid <i>router-id</i> , limit <i>num</i>	A Link State Database Overflow (LSDB) condition has occurred. The <i>router-id</i> is the router ID of the device. The <i>num</i> is the number of LSAs.
Notification	OSPF LSDB approaching overflow, rid <i>router-id</i> , limit <i>num</i>	The software is close to an LSDB condition. The <i>router-id</i> is the router ID of the device. The <i>num</i> is the number of LSAs.
Notification	OSPF intf rcvd bad pkt Bad Checksum, rid <i>ip-addr</i> , intf addr <i>ip-addr</i> , pkt size <i>num</i> , checksum <i>num</i> , pkt src addr <i>ip-addr</i> , pkt type <i>type</i>	The device received an OSPF packet that had an invalid checksum. The rid <i>ip-addr</i> is device's router ID. The intf addr <i>ip-addr</i> is the IP address of the interface that received the packet. The pkt size <i>num</i> is the number of bytes in the packet. The checksum <i>num</i> is the checksum value for the packet. The pkt src addr <i>ip-addr</i> is the IP address of the neighbor that sent the packet. The pkt type <i>type</i> is the OSPF packet type and can be one of the following: <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state acknowledgement • unknown (indicates an invalid packet type)
Notification	OSPF intf rcvd bad pkt Bad Packet type, rid <i>ip-addr</i> , intf addr <i>ip-addr</i> , pkt size <i>num</i> , checksum <i>num</i> , pkt src addr <i>ip-addr</i> , pkt type <i>type</i>	The device received an OSPF packet with an invalid type. The parameters are the same as for the Bad Checksum message. The pkt type <i>type</i> value is "unknown", indicating that the packet type is invalid.
Notification	OSPF intf rcvd bad pkt Unable to find associated neighbor, rid <i>ip-addr</i> , intf addr <i>ip-addr</i> , pkt size <i>num</i> , checksum <i>num</i> , pkt src addr <i>ip-addr</i> , pkt type <i>type</i>	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.
Notification	OSPF intf rcvd bad pkt Invalid packet size, rid <i>ip-addr</i> , intf addr <i>ip-addr</i> , pkt size <i>num</i> , checksum <i>num</i> , pkt src addr <i>ip-addr</i> , pkt type <i>type</i>	The device received an OSPF packet with an invalid packet size. The parameters are the same as for the Bad Checksum message.

TABLE 73 Syslog messages OSPFv3

Message level	Message	Explanation
Alert	OSPFv3 Memory Overflow	OSPF has run out of memory.
Alert	OSPFv3 LSA Overflow, LSA Type = <i>lsa-type</i>	<p>Indicates an LSA database overflow.</p> <p>The <i>lsa-type</i> parameter indicates the type of LSA that experienced the overflow condition. The LSA type is one of the following:</p> <ul style="list-style-type: none"> • 1 – Router • 2 – Network • 3 – Summary • 4 – Summary • 5 – External
Notification	OSPFv3 interface state changed, rid <i>router-id</i> , intf addr <i>ip-addr</i> , state <i>ospf-state</i>	<p>Indicates that the state of an OSPF interface has changed.</p> <p>The <i>router-id</i> is the router ID of the device.</p> <p>The <i>ip-addr</i> is the interface's IP address.</p> <p>The <i>ospf-state</i> indicates the state to which the interface has changed and can be one of the following:</p> <ul style="list-style-type: none"> • down • loopback • waiting • point-to-point • designated router • backup designated router • other designated router • unknown
Notification	OSPFv3 virtual intf state changed, rid <i>router-id</i> , area <i>area-id</i> , nbr <i>ip-addr</i> , state <i>ospf-state</i>	<p>Indicates that the state of an OSPF virtual routing interface has changed.</p> <p>The <i>router-id</i> is the router ID of the router the interface is on.</p> <p>The <i>area-id</i> is the area the interface is in.</p> <p>The <i>ip-addr</i> is the IP address of the OSPF neighbor.</p> <p>The <i>ospf-state</i> indicates the state to which the interface has changed and can be one of the following:</p> <ul style="list-style-type: none"> • down • loopback • waiting • point-to-point • designated router • backup designated router • other designated router • unknown

TABLE 73 Syslog messages OSPFv3

Message level	Message	Explanation
Notification	OSPFv3 nbr state changed, rid <i>router-id</i> , nbr addr <i>ip-addr</i> , nbr rid <i>nbr-router-id</i> , state <i>ospf-state</i>	<p>Indicates that the state of an OSPF neighbor has changed.</p> <p>The <i>router-id</i> is the router ID of the device.</p> <p>The <i>ip-addr</i> is the IP address of the neighbor.</p> <p>The <i>nbr-router-id</i> is the router ID of the neighbor.</p> <p>The <i>ospf-state</i> indicates the state to which the interface has changed and can be one of the following:</p> <ul style="list-style-type: none"> • down • attempt • initializing • 2-way • exchange start • exchange • loading • full • unknown
Notification	OSPFv3 virtual nbr state changed, rid <i>router-id</i> , nbr addr <i>ip-addr</i> , nbr rid <i>nbr-router-id</i> , state <i>ospf-state</i>	<p>Indicates that the state of an OSPF virtual neighbor has changed.</p> <p>The <i>router-id</i> is the router ID of the device.</p> <p>The <i>ip-addr</i> is the IP address of the neighbor.</p> <p>The <i>nbr-router-id</i> is the router ID of the neighbor.</p> <p>The <i>ospf-state</i> indicates the state to which the interface has changed and can be one of the following:</p> <ul style="list-style-type: none"> • down • attempt • initializing • 2-way • exchange start • exchange • loading • full • unknown

TABLE 73 Syslog messages OSPFv3

Message level	Message	Explanation
Notification	OSPFv3 intf config error, rid <i>router-id</i> , intf addr <i>ip-addr</i> , pkt src addr <i>src-ip-addr</i> , error type <i>error-type</i> , pkt type <i>pkt-type</i>	<p>Indicates that an OSPF interface configuration error has occurred.</p> <p>The <i>router-id</i> is the router ID of the device.</p> <p>The <i>ip-addr</i> is the IP address of the interface on the device.</p> <p>The <i>src-ip-addr</i> is the IP address of the interface from which the device received the error packet.</p> <p>The <i>error-type</i> can be one of the following:</p> <ul style="list-style-type: none"> • 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 <p>The <i>packet-type</i> can be one of the following:</p> <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state ack • unknown
Notification	OSPFv3 virtual intf config error, rid <i>router-id</i> , intf addr <i>ip-addr</i> , pkt src addr <i>src-ip-addr</i> , error type <i>error-type</i> , pkt type <i>pkt-type</i>	<p>Indicates that an OSPF virtual routing interface configuration error has occurred.</p> <p>The <i>router-id</i> is the router ID of the device.</p> <p>The <i>ip-addr</i> is the IP address of the interface on the device.</p> <p>The <i>src-ip-addr</i> is the IP address of the interface from which the device received the error packet.</p> <p>The <i>error-type</i> can be one of the following:</p> <ul style="list-style-type: none"> • 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 <p>The <i>packet-type</i> can be one of the following:</p> <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state ack • unknown

TABLE 73 Syslog messages OSPFv3

Message level	Message	Explanation
Notification	OSPFv3 intf authen failure, rid <i>router-id</i> , intf addr <i>ip-addr</i> , pkt src addr <i>src-ip-addr</i> , error type <i>error-type</i> , pkt type <i>pkt-type</i>	<p>Indicates that an OSPF interface authentication failure has occurred.</p> <p>The <i>router-id</i> is the router ID of the device.</p> <p>The <i>ip-addr</i> is the IP address of the interface on the device.</p> <p>The <i>src-ip-addr</i> is the IP address of the interface from which the device received the authentication failure.</p> <p>The <i>error-type</i> can be one of the following:</p> <ul style="list-style-type: none"> • 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 <p>The <i>packet-type</i> can be one of the following:</p> <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state ack • unknown
Notification	OSPFv3 virtual intf authen failure, rid <i>router-id</i> , intf addr <i>ip-addr</i> , pkt src addr <i>src-ip-addr</i> , error type <i>error-type</i> , pkt type <i>pkt-type</i>	<p>Indicates that an OSPF virtual routing interface authentication failure has occurred.</p> <p>The <i>router-id</i> is the router ID of the device.</p> <p>The <i>ip-addr</i> is the IP address of the interface on the device.</p> <p>The <i>src-ip-addr</i> is the IP address of the interface from which the device received the authentication failure.</p> <p>The <i>error-type</i> can be one of the following:</p> <ul style="list-style-type: none"> • 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 <p>The <i>packet-type</i> can be one of the following:</p> <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state ack • unknown

TABLE 73 Syslog messages OSPFv3

Message level	Message	Explanation
Notification	OSPFv3 intf rcvd bad pkt, rid <i>router-id</i> , intf addr <i>ip-addr</i> , pkt src addr <i>src-ip-addr</i> , pkt type <i>pkt-type</i>	Indicates that an OSPF interface received a bad packet. The <i>router-id</i> is the router ID of the device. The <i>ip-addr</i> is the IP address of the interface on the device. The <i>src-ip-addr</i> is the IP address of the interface from which the device received the authentication failure. The <i>packet-type</i> can be one of the following: <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state ack • unknown
Notification	OSPFv3 virtual intf rcvd bad pkt, rid <i>router-id</i> , intf addr <i>ip-addr</i> , pkt src addr <i>src-ip-addr</i> , pkt type <i>pkt-type</i>	Indicates that an OSPF interface received a bad packet. The <i>router-id</i> is the router ID of the device. The <i>ip-addr</i> is the IP address of the interface on the device. The <i>src-ip-addr</i> is the IP address of the interface from which the device received the authentication failure. The <i>packet-type</i> can be one of the following: <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state ack • unknown
Notification	OSPFv3 intf retransmit, rid <i>router-id</i> , intf addr <i>ip-addr</i> , nbr rid <i>nbr-router-id</i> , pkt type is <i>pkt-type</i> , LSA type <i>lsa-type</i> , LSA id <i>lsa-id</i> , LSA rid <i>lsa-router-id</i>	An OSPF interface on the device has retransmitted a Link State Advertisement (LSA). The <i>router-id</i> is the router ID of the device. The <i>ip-addr</i> is the IP address of the interface on the device. The <i>nbr-router-id</i> is the router ID of the neighbor router. The <i>packet-type</i> can be one of the following: <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state ack • unknown The <i>lsa-type</i> is the type of LSA. The <i>lsa-id</i> is the LSA ID. The <i>lsa-router-id</i> is the LSA router ID.

TABLE 73 Syslog messages OSPFv3

Message level	Message	Explanation
Notification	OSPFv3 virtual intf retransmit, rid <i>router-id</i> , intf addr <i>ip-addr</i> , nbr rid <i>nbr-router-id</i> , pkt type is <i>pkt-type</i> , LSA type <i>lsa-type</i> , LSA id <i>lsa-id</i> , LSA rid <i>lsa-router-id</i>	An OSPF interface on the device has retransmitted a Link State Advertisement (LSA). The <i>router-id</i> is the router ID of the device. The <i>ip-addr</i> is the IP address of the interface on the device. The <i>nbr-router-id</i> is the router ID of the neighbor router. The <i>packet-type</i> can be one of the following: <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state ack • unknown The <i>lsa-type</i> is the type of LSA. The <i>lsa-id</i> is the LSA ID. The <i>lsa-router-id</i> is the LSA router ID.
Notification	OSPFv3 originate LSA, rid <i>router-id</i> , area <i>area-id</i> , LSA type <i>lsa-type</i> , LSA id <i>lsa-id</i> , LSA router id <i>lsa-router-id</i>	An OSPF interface has originated an LSA. The <i>router-id</i> is the router ID of the device. The <i>area-id</i> is the OSPF area. The <i>lsa-type</i> is the type of LSA. The <i>lsa-id</i> is the LSA ID. The <i>lsa-router-id</i> is the LSA router ID.
Notification	OSPFv3 max age LSA, rid <i>router-id</i> , area <i>area-id</i> , LSA type <i>lsa-type</i> , LSA id <i>lsa-id</i> , LSA rid <i>lsa-router-id</i>	An LSA has reached its maximum age. The <i>router-id</i> is the router ID of the device. The <i>area-id</i> is the OSPF area. The <i>lsa-type</i> is the type of LSA. The <i>lsa-id</i> is the LSA ID. The <i>lsa-router-id</i> is the LSA router ID.
Notification	OSPFv3 LSDB overflow, rid <i>router-id</i> , limit <i>num</i>	A Link State Database Overflow (LSDB) condition has occurred. The <i>router-id</i> is the router ID of the device. The <i>num</i> is the number of LSAs.
Notification	OSPFv3 LSDB approaching overflow, rid <i>router-id</i> , limit <i>num</i>	The software is close to an LSDB condition. The <i>router-id</i> is the router ID of the device. The <i>num</i> is the number of LSAs.

TABLE 73 Syslog messages OSPFv3

Message level	Message	Explanation
Notification	OSPFv3 intf rcvd bad pkt Bad Checksum, rid <i>ip-addr</i> , intf addr <i>ip-addr</i> , pkt size <i>num</i> , checksum <i>num</i> , pkt src addr <i>ip-addr</i> , pkt type <i>type</i>	<p>The device received an OSPF packet that had an invalid checksum.</p> <p>The rid <i>ip-addr</i> is device's device ID.</p> <p>The intf addr <i>ip-addr</i> is the IP address of the interface that received the packet.</p> <p>The pkt size <i>num</i> is the number of bytes in the packet.</p> <p>The checksum <i>num</i> is the checksum value for the packet.</p> <p>The pkt src addr <i>ip-addr</i> is the IP address of the neighbor that sent the packet.</p> <p>The pkt type <i>type</i> is the OSPF packet type and can be one of the following:</p> <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state acknowledgement • unknown (indicates an invalid packet type)
Notification	OSPFv3 intf rcvd bad pkt Bad Packet type, rid <i>ip-addr</i> , intf addr <i>ip-addr</i> , pkt size <i>num</i> , checksum <i>num</i> , pkt src addr <i>ip-addr</i> , pkt type <i>type</i>	<p>The device received an OSPF packet with an invalid type.</p> <p>The parameters are the same as for the Bad Checksum message. The pkt type <i>type</i> value is "unknown", indicating that the packet type is invalid.</p>
Notification	OSPFv3 intf rcvd bad pkt Unable to find associated neighbor, rid <i>ip-addr</i> , intf addr <i>ip-addr</i> , pkt size <i>num</i> , checksum <i>num</i> , pkt src addr <i>ip-addr</i> , pkt type <i>type</i>	<p>The neighbor IP address in the packet is not on the device's list of OSPF neighbors.</p> <p>The parameters are the same as for the Bad Checksum message.</p>
Notification	OSPFv3 intf rcvd bad pkt Invalid packet size, rid <i>ip-addr</i> , intf addr <i>ip-addr</i> , pkt size <i>num</i> , checksum <i>num</i> , pkt src addr <i>ip-addr</i> , pkt type <i>type</i>	<p>The device received an OSPF packet with an invalid packet size.</p> <p>The parameters are the same as for the Bad Checksum message.</p>

TABLE 74 Syslog messages IS-IS

Message level	Message	Explanation
Alert	ISIS Memory Limit Exceeded	IS-IS is requesting more memory than is available.
Notification	ISIS ENTERED INTO OVERLOAD STATE	The device has set the overload bit to on (1), indicating that the device's IS-IS resources are overloaded.

TABLE 74 Syslog messages IS-IS (Continued)

Message level	Message	Explanation
Notification	ISIS Entered Overload State Due to <i>overload-reason</i>	<p>The device has set the overload bit to on (1), indicating that the device's IS-IS resources are Overloaded.</p> <p>Reasons for the overload as expressed in the <i>overload-reason</i> variable are:</p> <ul style="list-style-type: none"> • 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 <p>Definitions of the <i>overload-reason</i> values are described in Table 75.</p>
Notification	ISIS Exited Overload State	The device has set the overload bit to off (0), indicating that the device's IS-IS resources are no longer overloaded.
Notification	ISIS L1 ADJACENCY DOWN <i>system-id</i> on circuit <i>circuit-id</i>	<p>The device's adjacency with this Level-1 IS has gone down.</p> <p>The <i>system-id</i> is the system ID of the IS. The <i>circuit-id</i> is the ID of the circuit over which the adjacency was established.</p>
Notification	ISIS L1 ADJACENCY UP <i>system-id</i> on circuit <i>circuit-id</i>	<p>The device's adjacency with this Level-1 IS has come up.</p> <p>The <i>system-id</i> is the system ID of the IS. The <i>circuit-id</i> is the ID of the circuit over which the adjacency was established.</p>
Notification	ISIS L2 ADJACENCY DOWN <i>system-id</i> on circuit <i>circuit-id</i>	<p>The device's adjacency with this Level-2 IS has gone down.</p> <p>The <i>system-id</i> is the system ID of the IS. The <i>circuit-id</i> is the ID of the circuit over which the adjacency was established.</p>
Notification	ISIS L2 ADJACENCY UP <i>system-id</i> on circuit <i>circuit-id</i>	<p>The device's adjacency with this Level-2 IS has come up.</p> <p>The <i>system-id</i> is the system ID of the IS. The <i>circuit-id</i> is the ID of the circuit over which the adjacency was established.</p>

TABLE 74 Syslog messages IS-IS (Continued)

Message level	Message	Explanation
Notification	ISIS <i>LSP-type</i> LSP <i>LSP-ID</i> Seq <i>sequence-number</i> Len <i>length</i> LifeTime <i>lifetime</i> on <i>interface-name</i> dropped due to <i>LSP-drop-reason</i>	<p>The device has dropped the received LSP. The <i>LSP-Type</i> can be one of the following:</p> <ul style="list-style-type: none"> • L1 • L2 <p>The <i>LSP-ID</i> variable is in the 8 byte LSP ID value.</p> <p>The <i>sequence-number</i> is a 4 byte value that is associated with each LSP ID.</p> <p>The <i>length</i> is the length of the LSP PDU.</p> <p>The <i>lifetime</i> is the life period of the LSP.</p> <p>The <i>interface-name</i> is the name of the interface and is displayed in the following form "Ethernet 1/1".</p> <p>The <i>LSP-drop-reason</i> variable describes the following reasons that the LSP was dropped:</p> <ul style="list-style-type: none"> • 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 <p>Definitions of the <i>LSP-drop-reason</i> values are described in Table 75.</p>
Notification	ISIS <i>NbrType</i> Neighbor <i>Hostname/systemID</i> DOWN on <i>interface-name</i> due to <i>neighbor-down-reason</i>	<p>The device's Neighbor has gone down.</p> <p>The <i>NbrType</i> can be one of the following:</p> <ul style="list-style-type: none"> • L1 • L2 • PTPT <p>The <i>interface-name</i> is the name of the interface and is displayed in the following form "Ethernet 1/1".</p> <p>The <i>neighbor-down-reason</i> variable can be any one of the following reasons that the Neighbor is Down:</p> <ul style="list-style-type: none"> • BFD Trigger • Maximum Adjacencies • User Trigger • Hold Timer Expiry • Adjacency ID Mismatch • Adjacency Type Mismatch • Interface Down • Interface State Change <p>Definitions of the <i>neighbor-down-reason</i> values are described in Table 75.</p>

TABLE 74 Syslog messages IS-IS (Continued)

Message level	Message	Explanation
Notification	ISIS <i>NbrType</i> neighbor <i>Hostname/systemID</i> UP on <i>interface-name</i>	The device's Neighbor has come up. The <i>NbrType</i> can be one of the following: <ul style="list-style-type: none"> • L1 • L2 • PTPT The <i>interface-name</i> is the name of the interface and is displayed in the following form "Ethernet 1/1".
Notification	ISIS PTP ADJACENCY DOWN <i>mac</i> on interface <i>portnum</i>	
Notification	ISIS PTP ADJACENCY UP <i>mac</i> on interface <i>portnum</i>	

TABLE 75 Definition of IS-IS variables

Variable	Value	Definition
<i>neighbor-down-reason</i>	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.
<i>overload-reason</i>	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.

TABLE 75 Definition of IS-IS variables (Continued)

Variable	Value	Definition
	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.
<i>LSP-drop-reason</i>	Adjacency not found	The LSP was dropped because there 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.
	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.

TABLE 76 Syslog messages BGP

Message level	Message	Explanation
Debug	BGP4 Not enough memory available to run BGP4	The device could not start the BGP4 routing protocol because there is not enough memory available.
Error	BGP No of prefixes received from BGP peer <i>ip-addr</i> exceeds maximum prefix-limit...shutdown	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.
Error	BGP received invalid AS4_PATH attribute length (3) - entire AS4_PATH ignored	Possible attribute length can be only even number and cannot be odd. If an attribute with odd length is received, this error is displayed.

TABLE 76 Syslog messages BGP (Continued)

Message level	Message	Explanation
Error	BGP received invalid AS4_PATH attribute flag (0x40) - entire AS4_PATH ignored	If the flag that describes the attribute has unacceptable values then this error is displayed.
Error	BGP received invalid Confed info in AS4_PATH (@byte 43) - entire AS4_PATH ignored	Confederation segments(AS_CONFED_SEQ/SET) must precede the (AS_SEQ/SET), if not, this error is displayed.
Error	BGP received incorrect Seq type/len in AS4_PATH (@byte 41) - entire AS4_PATH ignored	Valid segment types are (AS_SEQ/SET, AS_CONFED_SEQ/SET), any other values results in an error being displayed.
Error	BGP received multiple AS4_PATH attributes - used first AS4_PATH attribute only	When AS4_PATH is received more than one time in the update message, this error is displayed.
Warning	BGP No of prefixes received from BGP peer <i>ip-addr</i> exceeds warning limit <i>num</i>	The device has received more than the allowed percentage of prefixes from the neighbor. The <i>ip-addr</i> is the IP address of the neighbor. The <i>num</i> 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.
Notification	BGP Peer <i>ip-addr</i> UP (ESTABLISHED)	Indicates that a BGP4 neighbor has come up. The <i>ip-addr</i> is the IP address of the neighbor's BGP4 interface with the device.
Notification	BGP Peer <i>ip-addr</i> DOWN (IDLE)	Indicates that a BGP4 neighbor has gone down. The <i>ip-addr</i> is the IP address of the neighbor's BGP4 interface with the device.
Notification	BGP Peer <i>ip</i> DOWN (<i>reasonrecv notif</i>)	
Notification	Configuration (Wait for BGP)	IS-IS is waiting for BGP convergence to complete.

TABLE 77 Syslog messages NTP

Message level	Message	Explanation
Warning	NTP server <i>ip-addr</i> failed to respond	Indicates that a Network Time Protocol (NTP) server did not respond to the device's query for the current time. The <i>ip-addr</i> indicates the IP address of the NTP server.

TABLE 78 Syslog messages TCP

Message level	Message	Explanation
Notification	TCP Local TCP exceeds <i>burst-max</i> burst packets, stopping for <i>lockup</i> seconds!	The number of TCP SYN packets exceeds the <i>burst-max</i> 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 <i>lockup</i> value. When the lockup period expires, the packet counter is reset and measurement is restarted.
Notification	TCP Transit TCP in interface <i>portnum</i> exceeds <i>num</i> burst packets, stopping for <i>num</i> seconds!	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 <i>portnum</i> is the port number. The first <i>num</i> is the maximum burst size (maximum number of packets allowed). The second <i>num</i> 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.

TABLE 79 Syslog messages DOT1X

Message level	Message	Explanation
Warning	DOT1X security violation at port <i>portnum</i> , malicious mac address detected <i>mac-address</i>	A security violation was encountered at the specified port number.
Warning	DOT1X Port <i>portnum</i> , AuthControlledPortStatus change restricted	
Notification	DOT1X Port <i>portnum</i> port default vlan-id changes to <i>vlan-id</i>	
Notification	DOT1X Port <i>portnum</i> currently used vlan-id changes to <i>vlan-id</i> due to move to restricted vlan	
Notification	DOT1X issues software port up indication of Port <i>portnum</i> to other software applications	The device has indicated that the specified port has been authenticated, but the actual port may not be active.
Notification	DOT1X issues software port down indication of Port <i>portnum</i> to other software applications	The device has indicated that the specified is no longer authorized, but the actual port may still be active.
Informational	DOT1X Port <i>portnum</i> , AuthControlledPortStatus change authorized	The status of the interface's controlled port has changed from unauthorized to authorized.
Informational	DOT1X Port <i>portnum</i> , AuthControlledPortStatus change unauthorized	The status of the interface's controlled port has changed from authorized to unauthorized.

TABLE 79 Syslog messages DOT1X (Continued)

Message level	Message	Explanation
Informational	DOT1X Port <i>portnum</i> currently used <i>vlan-id</i> changes to <i>vlan-id</i> due to dot1x-RADIUS <i>vlan</i> assignment	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 <i>vlan-id</i> .
Informational	DOT1X Port <i>portnum</i> currently used <i>vlan-id</i> is set back to port default <i>vlan-id</i> <i>vlan-id</i>	The user connected to <i>portnum</i> has disconnected, causing the port to be moved back into its default VLAN, <i>vlan-id</i> .
Informational	DOT1X Port <i>portnum</i> is unauthorized because system resource is not enough or the invalid information to set the dynamic assigned IP ACLs or MAC address filters	802.1X authentication could not take place on the port. This happened because strict security mode was enabled and one of the following occurred: <ul style="list-style-type: none"> • 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)
Debug	DOT1X Not enough memory	There is not enough system memory for 802.1X authentication to take place. Contact device Technical Support.

TABLE 80 Syslog messages SNMP

Message level	Message	Explanation
Informational	SNMP Auth. failure, intruder IP <i>ip-addr</i>	A user has tried to open a management session with the device using an invalid SNMP community string. The <i>ip-addr</i> is the IP address of the host that sent the invalid community string.
Informational	SNMP read-only community read-write community contact location user group view engineId trap [host] [value-str] deleted added modified from console telnet ssh web snmp session	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 engineId is specified.

TABLE 81 Syslog messages MPLS

Message level	Message	Explanation
Informational	Deleting VLL <i>name</i> (ID <i>number</i>) at <i>string</i> port <i>slot/port</i> with peer IPv4 address <i>ip-address</i>	Sent when PW traps are generated if the PW has been deleted, i.e. 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.
Informational	MPLS Deleting VLL <i>vll-name</i> (ID <i>vll-id</i>)	Sent when the specified VLL is being deleted.

TABLE 81 Syslog messages MPLS (Continued)

Message level	Message	Explanation
Informational	MPLS Deleting VLL <i>vll-name</i> (ID <i>vll-id</i>) at {tagged untagged} port <i>slot/port</i>	Sent when the specified VLL with the at the specified tagged or untagged port is being deleted.
Informational	MPLS Deleting VLL <i>vll-name</i> (ID <i>vll-id</i>) with peer IPv4 address <i>ip</i>	Sent when the specified VLL with the specified IPv4 peer is being deleted.
Informational	VLL is down for table index <i>number</i>	Sent when PW traps are generated if the VLL is down for one index
Informational	VLL is up for table index <i>number</i>	Sent when PW traps are generated if VLL is up for one index
Informational	VLLs are down for table indexes <i>number</i> through <i>number</i>	Sent when PW traps are generated if the VLLs represented by sequential entries in the database are down
Informational	VLLs are up for table indexes <i>number</i> through <i>number</i>	Sent when PW traps are generate dif VLLs represented by sequential entries in the database are up
Informational	VRF Port <i>slot-port</i> added to VRF <i>name</i> with updated port count <i>number</i>	Sent when an MPLS L3 VPN trap is generated if the state of an interface within the VRF changed from down to up.
Informational	VRF Port <i>slot-port</i> deleted from VRF <i>name</i> with updated port count <i>number</i>	Sent when an MPLS L3 VPN trap is generated if the state of an interface within the VRF changed from down to up.
Notification	MPLS Deleting VLL <i>name</i> (ID <i>vc-id</i>) at {tagged untagged} port <i>portnum</i> with peer IPv4 address <i>ip</i>	
Notification	MPLS LSP <i>Ispname</i> switches to new active path <i>pathame</i>	
Notification	MPLS LSP <i>Ispname</i> using path <i>pathname</i> is down	
Notification	MPLS LSP <i>Ispname</i> using path <i>pathname</i> is up	
Notification	MPLS VLL is down for table index <i>n</i>	
Notification	MPLS VLL is up for table index <i>n</i>	
Notification	MPLS VLLs are down for table indexes <i>n</i> through <i>m</i>	
Notification	MPLS VLLs are up for table indexes <i>n</i> through <i>m</i>	
Notification	MPLS VPLS [ID <i>id</i>] peer <i>ip</i> is down	Sent when a single VPLS peer is transitioning to a down state
Notification	MPLS VPLS [ID <i>id</i>] peer <i>ip</i> is up	Sent when a single VPLS peer is transitioning to an up state
Notification	MPLS VPLS <i>name</i> (ID <i>id</i>) endpoint <i>ip-address</i> is down	Sent when a single VPLS endpoint is transitioning to a down state.
Notification	MPLS VPLS <i>name</i> (ID <i>id</i>) endpoint <i>ip-address</i> is up	Sent when a single VPLS endpoint is transitioning to an up state.

TABLE 81 Syslog messages MPLS (Continued)

Message level	Message	Explanation
Notification	MPLS VPLS for instance indices <i>list n</i> through <i>m</i> are up	Sent when multiple VPLS instances are transitioning to an up state.
Notification	MPLS VPLS for instance indices <i>list n</i> through <i>m</i> are down	Sent when multiple VPLS instances are transitioning to a down state.
Notification	MPLS VPLS peer <i>ip</i> associated with VC ID <i>n</i> is up	Sent when a single VPLS peer is transitioning to an up state.
Notification	MPLS VPLS peer <i>ip</i> associated with VC ID <i>n</i> is down	Sent when a single VPLS peer is transitioning to a down state.
Notification	MPLS VPLS peer <i>ip</i> associated with instances <i>n-m list</i> is down	Sent when multiple VPLS instances associated with a peer are transitioning to a down state.
Notification	MPLS VPLS peer <i>ip</i> associated with instances <i>n-m list</i> is up	Sent when multiple VPLS instances associated with a peer are transitioning to an up state.
Notification	MPLS VPL endpoint <i>slot/port</i> associated with instance indices <i>list</i> is down	Sent when multiple VPLS instances associated with an endpoint is transitioning to a down state.
Notification	MPLS VPL endpoint <i>slot/port</i> associated with instance indices <i>list</i> is up	Sent when multiple VPLS instances associated with an endpoint is transitioning to an up state.
Notification	MPLS VLL-Local <i>name</i> is down	Sent when a single VLL-Local instance is transitioning to a down state.
Notification	MPLS VLL-Local <i>name</i> is up	Sent when a single VLL-Local instance is transitioning to an up state.
Notification	MPLS VLL-Local for instance indices <i>list n</i> through <i>m</i> are up	Sent when multiple VLL-Local instances are transitioning to an up state.
Notification	MPLS VLL-Local for instance indices <i>list n</i> through <i>m</i> are down	Sent when multiple VLL-Local instances are transitioning to a down state.
Notification	MPLS VLL <i>name</i> (ID <i>id</i> is down	Sent when a single VLL peer is transitioning to a down state.
Notification	MPLS VLL <i>name</i> (ID <i>id</i> is up	Sent when a single VLL peer is transitioning to an up state.
Notification	MPLS VLL for instance indices <i>list n</i> through <i>m</i> are up	Sent when multiple VLL instances are transitioning to an up state.
Notification	MPLS VLL for instance indices <i>list n</i> through <i>m</i> are down	Sent when multiple VLL instances are transitioning to a down state.

TABLE 82 Syslog messages VRF

Message level	Message	Explanation
Notification	VRF Port <i>portnum</i> added to VRF <i>name</i> with updated port count <i>n</i>	
Notification	VRF Port <i>portnum</i> deleted from VRF <i>name</i> with updated port count <i>n</i>	

TABLE 82 Syslog messages VRF (Continued)

Message level	Message	Explanation
Informational	VRF <i>vrf_name</i> has been configured as management VRF.	Indicates that the specified VRF has been configured as a management VRF.
Informational	VRF <i>vrf_name</i> has been un-configured as management VRF.	Indicates that the specified VRF has been removed as a management VRF.

TABLE 83 Syslog messages

Message level	Message	Explanation
Notification	Authentication Enabled on <i>portnum</i>	The multi-device port authentication feature was enabled on the on the specified <i>portnum</i> .
Notification	Authenticaiion Disabled on <i>portnum</i>	The multi-device port authentication feature was disabled on the on the specified <i>portnum</i> .

TABLE 84 Syslog messages BFD

Message level	Message	Explanation
Notification	BFD Session UP for NBR <i>neighbor-ID</i> on <i>port</i>	The BFD session is UP with the neighbor specified by the <i>neighbor-ID</i> on the port specified by the <i>port</i> variable.
Notification	BFD Session DOWN for NBR <i>neighbor-ID</i> on <i>port</i> Reason Neighbor Signaled Session Down	The BFD session with the neighbor specified by the <i>neighbor-ID</i> on the port specified by the <i>port</i> variable is Down because the BFD neighbor has signaled the session to be down.
Notification	BFD Session DOWN for NBR <i>neighbor-ID</i> on <i>port</i> Reason Administratively Down	The BFD session with the neighbor specified by the <i>neighbor-ID</i> on the port specified by the <i>port</i> variable is Down for Administrative reasons.

TABLE 85 Syslog messages Optics

Message level	Message	Explanation
Notification	Transceiver type checking has been disabled!	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.
Notification	Session DOWN for LSP <i>lsp-name</i> Reason Administratively Down	The BFD session for the LSP specified by the <i>lsp-name</i> is Down for Administrative reasons.
Notification	Session Up for LSP <i>lsp-name</i>	The BFD session for the LSP specified by the <i>lsp-name</i> is Up.

TABLE 85 Syslog messages Optics (Continued)

Message level	Message	Explanation
Notification	Session DOWN for RSVP session <i>session-id</i> Reason Administratively Down	The BFD session for the RSVP session specified by the <i>session-id</i> is Down for Administrative reasons. The form of the <i>session-id</i> displayed is IPv4 tunnel endpoint or tunnel ID or extended tunnel ID. For example 10.22.22.2/3/11/11/11/1
Notification	Session UP for RSVP session <i>session-id</i>	The BFD session for the RSVP session specified by the <i>session-id</i> is Up. The form of the <i>session-id</i> displayed is IPv4 tunnel endpoint or tunnel ID or extended tunnel ID. For example 10.22.22.2/3/11/11/11/1
Notification	Transceiver type checking has been enabled!	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.
Warning	Optic is not Brocade qualified (<i>port</i>) Type <i>type-description</i> Vendor <i>vendor-name</i> , Version <i>version-num</i> Part# <i>part-no</i> , Serial# <i>serial-no</i>	The optic module installed in the Interface module at the port specified by the <i>port</i> variable is not Brocade qualified although the port is still operational. The Type, Vendor, Version, Part#, and Serial # of the optic module is provided.
Alert	Optic is not Brocade qualified, optical monitoring is not supported (<i>port</i>) Type <i>type-description</i> Vendor <i>vendor-name</i> , Version <i>version-num</i> Part# <i>part-no</i> , Serial# <i>serial-no</i>	The optic module installed in the Interface module at the port specified by the <i>port</i> variable is not Brocade 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.
Alert	Optic is not capable of optical monitoring (<i>port</i>) Type <i>type-description</i> Vendor <i>vendor-name</i> , Version <i>version-num</i> Part# <i>part-no</i> , Serial# <i>serial-no</i>	The optic module installed in the Interface module at the port specified by the <i>port</i> 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.
Alert	Incompatible optical trans-receiver detected on port <i>n</i>	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.

TABLE 86 Syslog messages LDP

Message level	Message	Explanation
Notification	MPLS LDP path vector limit mismatch for session <i>lsrld labelSpaceId</i> (value <i>local vector limit</i>) with peer <i>lsrld labelSpaceId</i> (value <i>peer vector limit</i>)	This notification is generated when the value of the LDP path vector limit value from the peer does not match that of the entity.
Notification	MPLS LDP entity session <i>lsrld labelSpaceId</i> with peer <i>lsrld labelSpaceId</i> is up	This notification is sent when the value of 'mplsLdpSessionState' enters the 'operational(5)' state.
Notification	MPL LDP entity session <i>lsrld labelSpaceId</i> with peer <i>lsrld labelSpaceId</i> is down	This notification is sent when the value of 'mplsLdpSessionState' leaves the 'operational(5)' state.

TABLE 87 Syslog messages DHCP

Message level	Message	Explanation
Warning	DHCPC: No DHCP service available on the network.	The DHCP OFFER message is not received within 16 seconds of starting the DHCP address configuration phase.
Warning	DHCPC: Failed to renew DHCP lease on port 1/1 with IP address 10.1.1.1 mask 255.255.255.0	The DHCP lease cannot be renewed.
Warning	DHCPC: Failed to configure IP address on port 1/1; with IP address 10.1.1.1,mask 255.255.255.0	The IP address cannot be configured without a reason.
Warning	Failed to download image file <i>image name</i>	The image file cannot be downloaded.
Warning	DHCPC: Failed to download configuration file <i>image name</i>	The configuration files cannot be downloaded.

TABLE 88 Syslog messages DHCPv6

Message level	Message	Explanation
Warning	DHCPv6: Maximum allowed 60000 delegated prefixes learned.	The delegated prefixes' limit has reached the maximum value at the system level.
Warning	DHCPv6: Write to flash file to save delegated prefixes information failed.	Saving the delegated prefixes to flash file failed.
Warning	DHCPv6: Maximum allowed 20000 delegated prefixes learned on interface ve 100.	The delegated prefixes' limit has reached the maximum value at the interface level.

TABLE 89 Syslog messages NSR

Message level	Message	Explanation
Notification	NSR: Successfully notified RTM6 that OSPF6 switchover complete	OSPFv3 completes the restart process after switching over to the new master MP.

A Syslog messages

Global and Address Family Configuration Levels

Table 90 displays the individual devices and the Global and Address Family Configuration Level features they support.

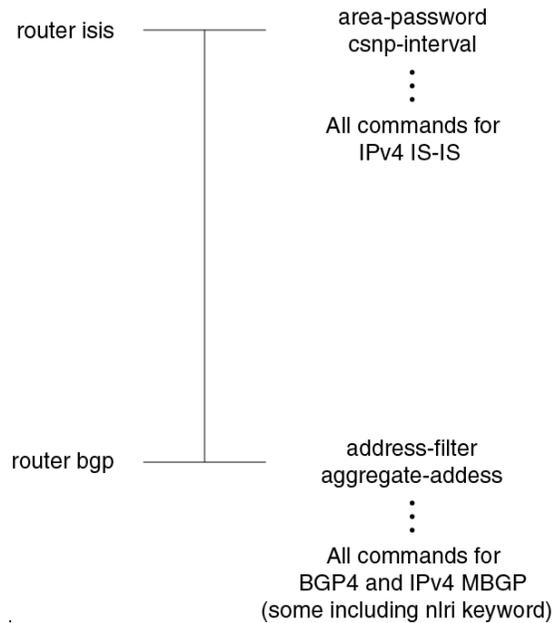
TABLE 90 Supported global and address family configuration level features

Features supported	Brocade Netron XMR Series	Brocade MLX Series	Brocade Netron CES 2000 Series BASE package	Brocade Netron CES 2000 Series ME_PREM package	Brocade Netron CES 2000 Series L3_PREM package	Brocade Netron CER 2000 Series Base package	Brocade Netron CER 2000 Series Advanced Services package
Accessing the Address Family Configuration Level Under BGP	Yes	Yes	No	No	Yes	Yes	Yes
Backward Compatibility for Existing BGP4 and IPv4 IS-IS	Yes	Yes	No	No	Yes	Yes	Yes
Global BGP4 Commands and BGP4 Unicast Route Commands	Yes	Yes	No	No	Yes	Yes	Yes

This appendix describes the restructured CLI for BGP and IS-IS on devices that support IPv6. In earlier versions of software, the CLI for BGP4 and IPv4 IS-IS is structured as shown in Figure 18.

FIGURE 18 Earlier structure of BGP4 and IPv4 IS-IS CLI

B Global and Address Family Configuration Levels



To configure BGP4 and IPv4 MBGP, enter the **router bgp** command, which takes you to the BGP router configuration level. At this level, you can access commands to configure all aspects of BGP4 and IPv4 MBGP, including commands that configure the protocol, and commands that configure unicast routes and multicast routes. (To configure aspects of multicast routes, specify the nri keyword with a command.)

To configure IPv4 IS-IS, enter the **router isis** command, which takes you to the IS-IS device configuration level. At this level, you can access commands that allow you to configure all aspects of IPv4 IS-IS, including commands that configure the protocol, and commands that configure unicast routes.

In both cases, the device determines, for example, whether commands entered at the BGP device configuration level apply to BGP4, to BGP4 unicast routes, or to IPv4 MBGP routes.

The introduction of IPv6 required the restructuring of existing BGP4 and IPv4 IS-IS CLI for the following reasons:

- To accommodate the IPv6-related CLI.
- To simplify the configuration of BGP4 unicast and IPv4 MBGP routes.

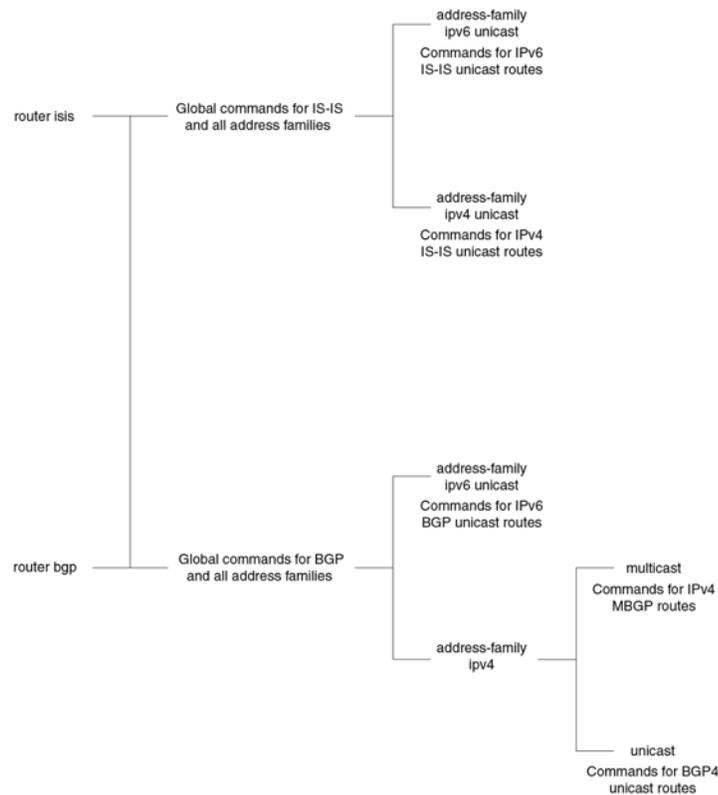
The CLI includes two layers of CLI for BGP and IS-IS (refer to [Figure 19](#)):

- A global layer to configure BGP and IS-IS protocols.
- Address families that separate the configuration of:
 - IPv6 and IPv4
 - Routing protocol

Sub-address families separate the configuration of:

- Unicast routes
- Multicast routes

FIGURE 19 IPv4, BGP4+ and IS-IS CLI structures



Accessing the address family configuration level

For example, to access the BGP4 multicast address family configuration level, enter the following command while at the global BGP configuration level.

```
Brocade(config-bgp)# address-family ipv4 multicast
Brocade(config-bgp-ipv4m)#
```

Syntax: `address-family ipv4 unicast | ipv4 multicast | ipv6 unicast`

The (config-bgp-ipv4m)# prompt indicates that you are at the IPv4 multicast address family configuration level. At this level, you can access commands that allow you to configure BGP4 multicast routes. The commands that you enter at this level apply to BGP4 multicast routes only. You can generate a configuration for BGP4 multicast routes that is separate and distinct from configurations for BGP4 unicast routes and BGP4+ unicast routes.

NOTE

Each address family configuration level gives you access to commands that apply to that address family only. To enable a feature in a particular address family, you must specify any associated commands for that feature in that address family.

To exit the BGP4 multicast address family configuration level, enter this command.

```
Brocade(config-bgp-ipv4m)# exit-address-family
Brocade(config-bgp)#
```

Syntax: `exit-address-family`

B Backward compatibility for existing BGP4 and IPv4 IS-IS configurations

When you enter the **exit-address-family** command at an address family configuration level you return to the global IS-IS configuration level, or the BGP4 unicast address family configuration level, (the default BGP4 level). For backward compatibility, you can currently access commands related to BGP4 unicast routes at both global BGP4 configuration and BGP4 unicast address family configuration levels. Both levels are indicated by the (config-bgp)# prompt.

Backward compatibility for existing BGP4 and IPv4 IS-IS configurations

When a device is upgraded to the current software version, the software automatically converts the existing BGP4 unicast and all IPv4 IS-IS configurations into the new address families. The software also automatically converts some of the IPv4 MBGP configuration into the new address family. Software conversion actions include:

- Leaves the global BGP4 and IPv4 IS-IS configurations as is.
- Converts the configuration for BGP4 unicast neighbors and routes into the BGP4 unicast address family.
- Converts the configuration for IPv4 IS-IS unicast routes into the IPv4 IS-IS unicast address family.
- Converts the configuration for IPv4 MBGP neighbors into IPv4 MBGP address family.

NOTE

The software does not convert all aspects of the IPv4 MBGP configuration. You must reconfigure the network routes, aggregate routes, redistribution of routes into IPv4 MBGP, and route map filters. Use the **show run** and **show ip bgp config** commands to check your IPv4 MBGP configuration.

Previously, IPv4 MBGP routes were configured using commands that included the **nlri** keyword. The current software version does not support the **nlri** keyword with IPv4 and IPv6 MBGP commands. You must now use the address families to configure all versions of BGP, IS-IS, and MBGP.

Global BGP4 commands and BGP4 unicast route commands

A global BGP command configures the BGP routing protocol and applies to all IPv4 and IPv6 address families. You can access the global commands while at the global BGP configuration level.

A BGP4 unicast route command configures a BGP4 unicast route. For backward compatibility, you can access BGP4 unicast route commands at both global BGP4 configuration and BGP4 unicast address family configuration levels. To help you distinguish the global BGP4 commands from the BGP4 unicast route commands at the global BGP4 configuration level, this section lists global BGP commands:

- **address-filter**
- **always-compare-med**
- **as-path-filter**
- **as-path-ignore**
- **bgp-redistribute-internal**

- **cluster-id**
- **community-filter**
- **compare-routerid**
- **confederation identifier**
- **confederation peers**
- **default-local-preference**
- **distance**
- **enforce-first-as**
- **fast-external-falover**
- **ignore-invalid-confed-as-path**
- **local-as**
- **med-missing-as-worst**
- **timers**

The following global BGP commands are used to configure peer groups and neighbors:

- **neighbor *ipv4-address* | *ipv6-address* | *peer-group-name* description**
- **neighbor *ipv4-address* | *ipv6-address* | *peer-group-name* distribute-list *acl_name* in**
- **neighbor *ipv4-address* | *ipv6-address* | *peer-group-name* distribute-list *acl_name* out**
- **neighbor *ipv4-address* | *ipv6-address* | *peer-group-name* distribute-list in**
- **neighbor *ipv4-address* | *ipv6-address* | *peer-group-name* distribute-list out**
- **neighbor *ipv4-address* | *ipv6-address* | *peer-group-name* ebgp-multihop**
- **neighbor *ipv4-address* | *ipv6-address* | *peer-group-name* filter-list in**
- **neighbor *ipv4-address* | *ipv6-address* | *peer-group-name* filter-list out**
- **neighbor *ipv4-address* | *ipv6-address* | *peer-group-name* next-hop-self**
- **neighbor *ipv4-address* | *ipv6-address* | *peer-group-name* password**
- **neighbor *ipv4-address* | *ipv6-address* | *peer-group-name* peer-group**
- **neighbor *ipv4-address* | *ipv6-address* | *peer-group-name* remote-as**
- **neighbor *ipv4-address* | *ipv6-address* | *peer-group-name* remove-private-as**
- **neighbor *ipv4-address* | *ipv6-address* | *peer-group-name* shut_down**
- **neighbor *ipv4-address* | *ipv6-address* | *peer-group-name* soft-reconfiguration inbound**
- **neighbor *ipv4-address* | *ipv6-address* | *peer-group-name* timers**
- **neighbor *ipv4-address* | *ipv6-address* | *peer-group-name* update-source**

The following address family commands modify the behavior of BGP for a specific address family:

- **aggregate-address**
- **client-to-client-reflection**
- **dampening**
- **default-information-originate**
- **default-metric**
- **maximum-paths**
- **multipath**

B Global BGP4 commands and BGP4 unicast route commands

- **neighbor** *ipv4-address* | *ipv6-address* | *peer-group-name* **filter-list in** (applies to the IPv4 unicast address family only)
- **network**
- **next-hop-enable-default**
- **next-hop-recursion** (applies to the IPv4 unicast address family only)
- **readvertise** (applies to the IPv4 unicast address family only)
- **redistribute**
- **table-map**
- **update-time**

The following commands configure policies for neighbors or peer groups for a specific address family:

- **neighbor** *ipv4-address* | *ipv6-address* | *peer-group-name* **activate**
- **neighbor** *ipv4-address* | *ipv6-address* | *peer-group-name* **capability orf prefixlist**
- **neighbor** *ipv4-address* | *ipv6-address* | *peer-group-name* **default-originate**
- **neighbor** *ipv4-address* | *ipv6-address* | *peer-group-name* **filter-list as-path-access-list in**
- **neighbor** *ipv4-address* | *ipv6-address* | *peer-group-name* **filter-list as-path-access-lis out**
- **neighbor** *ipv4-address* | *ipv6-address* | *peer-group-name* **maximum-prefix**
- **neighbor** *ipv4-address* | *ipv6-address* | *peer-group-name* **prefix-list prefix_list_name in**
- **neighbor** *ipv4-address* | *ipv6-address* | *peer-group-name* **prefix-list prefix_list_name out**
- **neighbor** *ipv4-address* | *ipv6-address* | *peer-group-name* **route-map in**
- **neighbor** *ipv4-address* | *ipv6-address* | *peer-group-name* **route-map out**
- **neighbor** *ipv4-address* | *ipv6-address* | *peer-group-name* **route-reflector-client**
- **neighbor** *ipv4-address* | *ipv6-address* | *peer-group-name* **send-community**
- **neighbor** *ipv4-address* | *ipv6-address* | *peer-group-name* **unsuppress-map**
- **neighbor** *ipv4-address* | *ipv6-address* | *peer-group-name* **weight**
- **neighbor** *ipv4-address* | *ipv6-address* | *peer-group-name* **route-map**
- **neighbor** *ipv4-address* | *ipv6-address* | *peer-group-name* **allowas-in**
- **neighbor** *ipv4-address* | *ipv6-address* | *peer-group-name* **send-label**

Currently, you can create a neighbor with an IPv4 or IPv6 address at the global BGP configuration or IPv4 unicast address family configuration level. For example, if you create a neighbor with an IPv4 address at this level, by default, the neighbor is enabled to exchange IPv4 unicast prefixes. However, this neighbor cannot exchange IPv4 multicast prefixes until you explicitly enable it to do so by entering the **neighbor** *ipv4-address* | *peer-group-name* **activate** command at the IPv4 multicast address family configuration level. Likewise, if you create a neighbor with an IPv6 address at this level, the neighbor will not exchange IPv6 unicast prefixes until you explicitly enable it to do so by entering the **neighbor** *ipv6-address* | *peer-group-name* **activate** command at the IPv6 unicast address family configuration level.

If you create a neighbor at the IPv4 multicast address family configuration or IPv6 unicast address family configuration levels, by default, the neighbor is enabled to exchange IPv4 multicast prefixes or IPv6 unicast prefixes, respectively. You do not need to explicitly enable the neighbor at either level.

Commands That Require a Reload

Table 91 displays the individual Brocade devices and the commands that require a reload features they support.

TABLE 91 Supported Brocade commands that require a reload features

Features supported	Brocade Netron XMR Series	Brocade MLX Series	Brocade Netron CES 2000 Series BASE package	Brocade Netron CES 2000 Series ME_PREM package	Brocade Netron CES 2000 Series L3_PREM package	Brocade Netron CER 2000 Series Base package	Brocade Netron CER 2000 Series Advanced Services package
cam-mode ip	Yes	Yes	No	No	No	No	No
cam-mode ipvpn	Yes	Yes	No	No	No	No	No
default-max-frame-size	Yes	Yes	Yes	Yes	Yes	Yes	Yes
multicast-flooding	Yes	Yes	Yes	Yes	Yes	Yes	Yes
port-priority	Yes	Yes	Yes	Yes	Yes	Yes	Yes
system-max	Yes	Yes	Yes	Yes	Yes	Yes	Yes
virtual-interface-mac	Yes	Yes	Yes	Yes	Yes	Yes	Yes
vll-mtu-enforcement	Yes	Yes	No	No	No	No	No

Most CLI commands take effect as soon as you enter them. However, a small number of commands require a software reload to take effect. Table 92 lists these commands.

To place a configuration change made by one of these commands into effect, you must save the change to the startup-config file, then reload the software. If you reload the software without saving the change to the startup-config file, the device does not make the change.

To reload the software, you must perform a warm start. To perform a warm start, do one of the following:

- Enter the **reload** command at the Privileged EXEC level of the CLI.
- Enter the **boot system** command at the Privileged EXEC level of the CLI.

TABLE 92 Commands that require a software reload

Command	See ...
cam-mode ip	Foundry Direct Routing and CAM Partition Profiles for the NetIron XMR and the Brocade MLX Series chapter
cam-mode ipvpn	Foundry Direct Routing and CAM Partition Profiles for the NetIron XMR and the Brocade MLX Series chapter

C Commands That Require a Reload

TABLE 92 Commands that require a software reload (Continued)

Command	See ...
default-max-frame-size	Foundry Direct Routing and CAM Partition Profiles for the NetIron XMR and the Brocade MLX Series chapter
multicast-flooding	Multi-Service IronWare Switching Configuration Guide
system-max	Foundry Direct Routing and CAM Partition Profiles for the NetIron XMR and the Brocade MLX Series chapter
virtual-interface-mac	Multi-Service IronWare Switching Configuration Guide
vll-mtu-enforcement	Multi-Service IronWare MPLS Configuration Guide

NIAP-CCEVS

Some devices have passed the Common Criteria (CC) certification testing. This testing is sponsored by the National Information Assurance Partnership (NIAP) - Common Criteria Evaluation and Validation Scheme (CCEVS). For more information regarding the NIAP-CCEVS certification process refer to the following link: <http://www.niap-ccevs.org/>.

In an effort to maintain a proper level of security as it relates to access to network infrastructure resources, Brocade recommends that all Brocade hardware be installed within a secure location that is accessible by approved personnel only.

NIAP-CCEVS certified Brocade equipment and Ironware releases

The following devices have been NIAP-CCEVS certified. The following IronWare software release must be used to remain compliant with this certification:

TABLE 93 NIAP-CCEVS certified equipment and IronWare software releases

Brocade product	Brocade IronWare software version	Discussed in
Brocade XMR Family	3.8.00a	<i>Multi-Service IronWare Administration Configuration Guide</i>
Brocade MLX Family	3.8.00a	<i>Multi-Service IronWare Administration Configuration Guide</i>
BigIron RX Family	2.5.00b	<i>BigIron RX Series Configuration Guide</i>
FastIron SuperX/SX Family	4.1.00	<i>FastIron and Turbolron Configuration Guide</i>
FastIron Edge X Family	4.1.00	<i>FastIron and Turbolron Configuration Guide</i>
FastIron GS/LS Family	4.2.00a	<i>FastIron and Turbolron Configuration Guide</i>
FastIron Edge Switch Family	4.0.00a	<i>Security Guide</i>

D Web management access to NIAP-CCEVS certified Brocade equipment

TABLE 93 NIAP-CCEVS certified equipment and IronWare software releases

Brocade product	Brocade IronWare software version	Discussed in
ServerIron JetCore Family	11.0.00a	<i>ServerIron TrafficWorks Graphical User Interface</i> <i>ServerIron TrafficWorks Server Load Balancing Guide</i> <i>ServerIron TrafficWorks Advanced Server Load Balancing Guide</i> <i>ServerIron TrafficWorks Global Server Load Balancing Guide</i> <i>ServerIron TrafficWorks Security Guide</i> <i>ServerIron TrafficWorks Administration Guide</i> <i>ServerIron TrafficWorks Switching and Routing Guide</i> <i>ServerIron Firewall Load Balancing Guide</i>
ServerIron ADX Family	12.0.00	<i>ServerIron ADX TrafficWorks Graphical User Interface</i> <i>ServerIron ADX TrafficWorks Server Load Balancing Guide</i> <i>ServerIron ADX TrafficWorks Advanced Server Load Balancing Guide</i> <i>ServerIron ADX TrafficWorks Global Server Load Balancing Guide</i> <i>ServerIron ADX TrafficWorks Security Guide</i> <i>ServerIron ADX TrafficWorks Administration Guide</i> <i>ServerIron ADX TrafficWorks Switching and Routing Guide</i> <i>ServerIron ADX Firewall Load Balancing Guide</i>

Web management access to NIAP-CCEVS certified Brocade equipment

All devices that are to remain in compliancy with the NIAP-CCEVS certification must disable all remote access through the integrated Web management graphical user interface (GUI). In accordance with NIAP-CCEVS this functionality is considered a security risk and must be disabled.

Please refer to the Brocade Configuration Guides associated with each product in the table [“NIAP-CCEVS certified equipment and IronWare software releases”](#) for detailed instructions on how to disable the Web Management Interface feature.

Warning: local user password changes

Please note that if existing usernames and passwords have been configured on a device with specific privilege levels (super-user, read-only, port-config) and if you attempt to change a user's password by executing the following command.

```
Brocade(config)# user fdryreadonly password <value>
```

The privilege level of this particular user will be changed from its current value to “super-user”. The “super-user” level username and password combination provides full access to the Brocade command line interface (CLI). To prevent this from occurring, use the following command.

```
Brocade(config)# user fdryreadonly privilege <value> password <value>
```

D Warning: local user password changes

Acknowledgements

This appendix presents the acknowledgements for portions of code from various vendors that are included in the Brocade devices covered in this manual.

Cryptographic software

MPL 1.1

The contents of this file are subject to the Mozilla Public License Version 1.1 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at <http://www.mozilla.org/MPL/>.

Software distributed under the License is distributed on an "AS IS" basis, WITHOUT WARRANTY OF ANY KIND, either express or implied. See the License for the specific language governing rights and limitations under the License.

The Original Code is the Network Security Services libraries.

The Initial Developer of the Original Code is Red Hat, Inc. Portions created by the Initial Developer are Copyright (C) 2009 the Initial Developer. All Rights Reserved.

Portions created by Netscape Communications Corporation are Copyright (C) 1994-2000 Netscape Communications Corporation. All Rights Reserved.

Contributor(s):

OpenSSL license

Copyright (c) 1998-2001 The OpenSSL Project. All rights reserved.

1. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:
2. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
3. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation or other materials provided with the distribution.
4. All advertising materials mentioning features or use of this software must display the following acknowledgment: "This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (<http://www.openssl.org/>)"
5. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact openssl-core@openssl.org.

E Cryptographic software

6. Products derived from this software may not be called “OpenSSL” nor may “OpenSSL” appear in their names without prior written permission of the OpenSSL Project.
7. Redistributions of any form whatsoever must retain the following acknowledgment: “This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>)”

THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT “AS IS” AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Cryptographic software

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com). This product includes software written by Tim Hudson (tjh@cryptsoft.com).

Original SSLeay License

Copyright (C) 1995-1998 Eric Young (eay@cryptsoft.com). All rights reserved.

This package is an SSL implementation written by Eric Young (eay@cryptsoft.com). The implementation was written so as to conform with Netscape’s SSL.

This library is free for commercial and non-commercial use as long as the following conditions are adhered to. The following conditions apply to all code found in this distribution, be it the RC4, RSA, lhash, DES, etc., code; not just the SSL code. The SSL documentation included with this distribution is covered by the same copyright terms except that the holder is Tim Hudson (tjh@cryptsoft.com).

Copyright remains Eric Young’s, and as such any Copyright notices in the code are not to be removed. If this package is used in a product, Eric Young should be given attribution as the author of the parts of the library used. This can be in the form of a textual message at program startup or in documentation (online or textual) provided with the package.

1. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:
2. Redistributions of source code must retain the copyright notice, this list of conditions and the following disclaimer.
3. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation or other materials provided with the distribution.
4. All advertising materials mentioning features or use of this software must display the following acknowledgement:

“This product includes cryptographic software written by Eric Young (eay@cryptsoft.com)” The word ‘cryptographic’ can be left out if the routines from the library being used are not cryptographic related:-).

5. If you include any Windows specific code (or a derivative thereof) from the apps directory (application code) you must include an acknowledgement:

“This product includes software written by Tim Hudson (tjh@cryptsoft.com)”

THIS SOFTWARE IS PROVIDED BY ERIC YOUNG “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The license and distribution terms for any publicly available version or derivative of this code cannot be changed. i.e. this code cannot simply be copied and put under another distribution license [including the GNU Public License.]

E Cryptographic software

NP Memory Errors

The Sysmon NP Memory Error Monitoring event monitors memory errors on interface modules. This appendix lists the interface cards that support NP memory error monitoring. It also details the different NP memory errors supported on each interface card.

The following interface cards support NP memory error monitoring:

- BR-MLX-40Gx4-X
- BR-MLX-10Gx24
- BR-MLX-100Gx2-X(100G)
- Gen-1
 - NI-XMR-10Gx4
 - NI-MLX-10Gx4)
- Gen-1.1
 - BR-MLX-1GCx24-X
 - BR-MLX-1GFx24-X
 - BR-MLX-10Gx4-X
- Gen-2
 - NI-MLX-10Gx8-M
 - NI-MLX-10Gx8-D
 - BR-MLX-10Gx8-X

TABLE 94 NP memory errors supported on BR-MLX-40Gx4-X interface cards

External Memory Errors	
1	PRAM Parity Errors
2	CAM2PRAM Parity Errors
3	LBLRAM Parity Errors
4	CAM wd10 Parity Error
5	CAM GIO Parity Error
6	CAM PEO Parity Error
7	CAM Operation Parity Error
8	CAM Result Bus Parity Error
Internal Memory Errors	
1	Tx Deframer MVLAN Flag FIFO Parity
2	Tx Deframer MVLAN control Packet FIFO Parity
3	Tx Deframer MVLAN replication table Parity
4	Tx Deframer MVLAN start offset FIFO Parity

TABLE 94 NP memory errors supported on BR-MLX-40Gx4-X interface cards

5	Tx Deframer MVLAN sop FIFO Parity
6	Tx Deframer MVLAN payload Data FIFO Parity
7	Tx Packet Edit Data FIFO Parity
8	Tx Packet Edit Next Hop Table Parity
9	ACL PRAM Results FIFO Parity
10	ACL Data FIFO Parity
11	ACL Control FIFO Parity
12	ACL QoS Done FIFO Parity
13	ACL Port Number FIFO Parity
14	ACL Priet Table Parity
15	ACL Tx VLAN Table Parity
16	Tx Priet Lookup Result Parity
17	MAC0 Frame LSTD Parity
18	MAC0 Frame Data Parity
19	MAC0 Frame Control Parity
20	MAC1 Frame LSTD Parity
21	MAC1 Frame Data Parity
22	MAC1 Frame Control Parity
23	Tx Packet Edit Data FIFO Parity
24	Tx Packet Edit Control FIFO Parity
25	Tx Packet Edit nhlk FIFO Parity
26	Tx Packet Edit pipe LBLc FIFO Parity
27	Start Offset Table CPU Read Parity
28	Replacement Table CPU Read Parity
29	Next Hop Table CPU Read Parity
30	Tx VLAN Table CPU Read Parity
31	Priet Table CPU Read Parity
32	Rx MAC Data FIFO Read Parity
33	Rx MAC Flag FIFO Read Parity
34	Rx MAC Data FIFO Read Parity
35	Rx MAC Flag FIFO Read Parity
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

TABLE 94 NP memory errors supported on BR-MLX-40Gx4-X interface cards

42	CAM2 Lookup FIFO1 Overflow
43	Rx Port Pipeline HQoS Data Parity
44	Rx Port Pipeline Rx Data-in Parity
45	Rx Port Pipeline Rxctrl FIFO Read Data Parity
46	Rx Port Pipeline Read Rx QoS Id FIFO Parity
47	Rx Port Pipeline Rx portnum FIFO Parity
48	Rx Port Pipeline Rx QoS Done FIFO Parity
49	Rx Port Pipeline Rx Flag FIFO Parity
50	Rx Port Pipeline Rx Header FIFO Parity
51	Rx Port Pipeline PRAM Packet Id Status
52	Rx Port Pipeline Data Path Packet Id Status
53	EXM IP Address FIFO Overflow
54	Rx Packet Header FIFO Parity
55	LBL Lookup FIFO Overflow
56	Indicates Parity in the Packet Decode FIFO
57	PRAM Result Error Flag evalclk
58	ECMP FIFO Overflow
59	LBL2 ECMP FIFO Overflow
60	Rx port Pipeline Rx QoS Id FIFO Overflow Status
61	Rx port Pipeline Rx Flag FIFO Overflow Status
62	Rx port Pipeline Rx Header FIFO Overflow Status
63	Rx CAM Result FIFO Parity
64	Rx CAM Result FIFO Underflow Status
65	Rx CAM Result FIFO Overflow Status
66	Rx Packet Decode FIFO Overflow Status
67	Rx Packet Decode FIFO underflow Status
68	Rx Packet Decode FIFO Parity
69	Rx topotos FIFO Parity
70	Rx topotos FIFO Underflow Status
71	Rx topotos FIFO Overflow Status
72	CAM1 Lookup FIFO 1
73	CAM1 Lookup FIFO 2
74	CAM1 Lookup FIFO 3
75	CAM1 Asc FIFO
76	CAM2 Lookup FIFO 1
77	CAM2 Lookup FIFO 2
78	CAM2 Lookup FIFO 3

TABLE 94 NP memory errors supported on BR-MLX-40Gx4-X interface cards

79	CAM2 Asc FIFO
80	LBLram srvt Lookup FIFO Underflow
81	LBLram extd Service rd Data Parity
82	LBLram LBL Lookup FIFO Underflow
83	LBLram srvt Lookup FIFO Overflow
84	LBLram srvt Lookup FIFO Underflow
85	LBLram rdrequest FIFO Overflow
86	LBLram rdrequest FIFO Underflow
87	LBLram extd Service Read FIFO Overflow
88	LBLram extd Service Read FIFO Underflow
89	srvt Lookup FIFO rd Data Parity
90	cam Result Scheduler FIFO Underflow
91	CAM1 Result FIFO Overflow
92	CAM1 Result FIFO Underflow
93	CAM2 Result FIFO Overflow
94	CAM2 Result FIFO Underflow
95	Label Result FIFO Overflow
96	Label Result FIFO Underflow
97	LBL hold FIFO Underflow
98	LBL hold FIFO Overflow
99	ECMP FIFO Underflow
100	LBL Result sync FIFO Underflow
101	LBL Result sync FIFO Overflow
102	LBL2 ECMP FIFO Underflow
103	EXM cpu2hashbkt rData Parity
104	EXM hwsrch hashbkt rData Parity
105	EXM hash idx table Parity
106	EXM IP Address FIFO Underflow
107	EXM IP Address FIFO Parity
108	EXM VPN Id FIFO Underflow
109	EXM VPN Id FIFO Parity
110	Stats Block0 Rx FIFO Underflow
111	Stats Block0 tx FIFO Overflow
112	Stats Block0 tx FIFO Underflow
113	Stats Block0 Rx FIFO Overflow
114	Stats Block1 Rx FIFO Underflow
115	Stats Block1 tx FIFO Overflow

TABLE 94 NP memory errors supported on BR-MLX-40Gx4-X interface cards

116	Stats Block1 tx FIFO Underflow
117	Stats Block1 Rx FIFO Overflow
118	Stats Block2 Rx FIFO Underflow
119	Stats Block2 tx FIFO Overflow
120	Stats Block2 tx FIFO Underflow
121	Stats Block2 Rx FIFO Overflow
122	Stats Block3 Rx FIFO Underflow
123	Stats Block3 tx FIFO Overflow
124	Stats Block3 tx FIFO Underflow
125	Stats Block3 Rx FIFO Overflow
126	Stats Block4 Rx FIFO Underflow
127	Stats Block4 tx FIFO Overflow
128	Stats Block4 tx FIFO Underflow
129	Stats Block4 Rx FIFO Overflow
130	Stats Block5 Rx FIFO Underflow
131	Stats Block5 tx FIFO Overflow
132	Stats Block5 tx FIFO Underflow
133	Stats Block5 Rx FIFO Overflow
134	CAM2PRAM QDR Interface mw FIFO Parity
135	CAM2PRAM QDR Interface mw FIFO Underflow
136	CAM2PRAM QDR Interface mw FIFO Overflow
137	CAM2PRAM QDR Interface rdrequest FIFO Parity
138	CAM2PRAM QDR Interface rdrequest FIFO Underflow
139	CAM2PRAM QDR Interface rdrequest FIFO Overflow
140	CAM2PRAM cpu FIFO Parity
141	CAM2PRAM CAM Interface Data FIFO Underflow
142	CAM2PRAM CAM Interface Count FIFO Underflow
143	CAM2PRAM Result FIFOs Parity
144	CAM2PRAM Result FIFOs Underflow
145	CAM2PRAM Result FIFOs Overflow
146	CAM2PRAM Result Scheduler Underflow
147	CAM2PRAM Result Scheduler Overflow
148	PRAM CAM Interface Data FIFO Underflow 0
149	PRAM CAM Interface Data FIFO Overflow 0
150	PRAM CAM Interface Data FIFO Underflow 1
151	PRAM CAM Interface Data FIFO Overflow 1
152	PRAM CAM Interface Data FIFO Underflow 2

TABLE 94 NP memory errors supported on BR-MLX-40Gx4-X interface cards

153	PRAM CAM Interface Data FIFO Overflow 2
154	PRAM CAM Interface Data FIFO Underflow 3
155	PRAM CAM Interface Data FIFO Overflow 3
156	PRAM Channel0 rdParity Flag
157	PRAM Channel1 rdParity Flag
158	PRAM Channel2 rdParity Flag
159	PRAM Channel3 rdParity Flag
160	CAM2Age L2 Underflow 0
161	CAM2Age L2 Underflow 1
162	CAM2Age ACL Underflow 0
163	CAM2Age ACL Underflow 1
164	CAM2Age L3 Underflow
165	L2 Aged FIFO Underflow
166	L2 Aged FIFO Overflow
167	L2 Aged FIFO Parity
168	L2 Aged mem Parity
169	L3 Aged FIFO Underflow
170	L3 Aged FIFO Overflow
171	L3 Aged FIFO Parity
172	L3 Aged mem Parity
173	ACL Aged FIFO Underflow
174	ACL Aged FIFO Overflow
175	ACL Aged FIFO Parity
176	ACL Aged mem Parity
177	Rx QoS Id ff Underflow
178	Rx QoS Id ff Overflow
179	Rx Flag ff Underflow
180	Rx QoS Done ff Underflow
181	Rx QoS Done ff Overflow
182	PRAM QDR Interface rdrequest FIFO Parity
183	PRAM QDR Interface rdrequest FIFO Underflow
184	PRAM QDR Interface rdrequest FIFO Overflow
185	PRAM QDR Interface cpu rd FIFO Parity
186	PRAM Result FIFO Parity
187	PRAM Result FIFO Overflow
188	CAM2Age L3 Overflow
189	CAM2Age L2 Overflow 0

TABLE 94 NP memory errors supported on BR-MLX-40Gx4-X interface cards

190	CAM2Age L2 Overflow 1
191	CAM2Age ACL Overflow 0
192	CAM2Age ACL Overflow 1
193	CAM2PRAM Data FIFO Overflow
194	CAM2PRAM Count FIFO Overflow
195	CAM1 Lookup FIFO1 Underflow
196	CAM1 Lookup FIFO2 Underflow
197	CAM1 Lookup FIFO3 Underflow
198	CAM2 Lookup FIFO1 Underflow
199	CAM2 Lookup FIFO2 Underflow
200	CAM2 Lookup FIFO3 Underflow
201	CAM1 Asc FIFO Underflow
202	CAM1 Asc FIFO Overflow
203	CAM2 Asc FIFO Underflow
204	CAM2 Asc FIFO Overflow
205	EXM VPN Id FIFO Overflow
206	CAM3 Lookup FIFO1 Underflow
207	CAM3 Lookup FIFO2 Underflow
208	CAM3 Lookup FIFO3 Underflow
209	CAM3 Asc FIFO Underflow
210	CAM3 Asc FIFO Overflow
211	CAM3 Lookup FIFO1 Parity
212	CAM3 Lookup FIFO2 Parity
213	CAM3 Lookup FIFO3 Parity
214	CAM3 Asc FIFO Parity
215	Cmd rdData Parity
216	Cmpl Data Parity
217	CAM3 Lookup FIFO1 Overflow
218	CAM3 Lookup FIFO2 Overflow
219	CAM3 Lookup FIFO3 Overflow
220	Service CAM Block Mux FIFO Overflow
221	Service CAM Block Mux FIFO Underflow
222	eACL CAM Block Mux FIFO1 Overflow
223	eACL CAM Block Mux FIFO1 Underflow
224	eACL CAM Block Mux FIFO2 Overflow
225	eACL CAM Block Mux FIFO2 Underflow
226	Service CAM Result FIFO Overflow

TABLE 94 NP memory errors supported on BR-MLX-40Gx4-X interface cards

227	Service CAM Result FIFO Underflow
228	eACL CAM Result FIFO Overflow
229	eACL CAM Result FIFO Underflow
230	eACL CAM Block Mux FIFO2 Parity
231	eACL CAM Block Mux FIFO1 Parity
232	Service CAM Block Mux FIFO Parity
233	CAM3 Service Result FIFO Parity
234	CAM3 eACL Result FIFO Parity

TABLE 95 NP memory errors supported on BR-MLX-10Gx24 interface cards.

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
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 Frame 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

TABLE 95 NP memory errors supported on BR-MLX-10Gx24 interface cards.

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 IOR PCIe Read Parity Error
31	PIB IOW PCIe Read Parity Error
32	PIB IO 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
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

TABLE 95 NP memory errors supported on BR-MLX-10Gx24 interface cards.

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
83	CCM Session Table Read Uncorrectible ECC Error
84	CCM Session Table Read Correctible ECC Error

TABLE 95 NP memory errors supported on BR-MLX-10Gx24 interface cards.

85	CCM Hash Table Read Uncorrectable ECC Error
86	CCM Hash Table Read Correctable ECC Error
87	CCM Hash Bucket Read Uncorrectable ECC Error
88	CCM Hash Bucket Read Correctable ECC Error
89	BFD Session Table Read Uncorrectable ECC Error
90	BFD Session Table Read Correctable 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 96 NP memory errors supported on BR-MLX-100Gx2-X(100G) interface cards.

External Memory Errors	
1	LBLRAM Parity Errors

TABLE 96 NP memory errors supported on BR-MLX-100Gx2-X(100G) interface cards.

2	Age RAM 1 Parity Errors
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	Nexthtable 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

TABLE 96 NP memory errors supported on BR-MLX-100Gx2-X(100G) interface cards.

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
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

TABLE 96 NP memory errors supported on BR-MLX-100Gx2-X(100G) interface cards.

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
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 rdrequest FIFO Parity
99	PRAM1 CPU rd FIFO Parity
100	PRAM2 rdrequest FIFO Parity
101	PRAM2 CPU rd FIFO Parity
102	PRAM3 rdrequest FIFO Parity
103	PRAM3 CPU rd FIFO Parity
104	CAM2PRAM1 rdrequest FIFO Parity

TABLE 96 NP memory errors supported on BR-MLX-100Gx2-X(100G) interface cards.

105	CAM2PRAM1 ECMP FIFO rdData Parity
106	CAM2PRAM1 CPU FIFO rdData Parity
107	CAM2PRAM1 mw FIFO rd Parity
108	CAM2PRAM2 rdrequest FIFO Parity
109	CAM2PRAM2 ECMP FIFO rdData Parity
110	CAM2PRAM2 CPU FIFO rdData Parity
111	CAM2PRAM2 mw FIFO rd Parity
112	CAM2PRAM3 rdrequest 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
118	Rx MAC Flag FIFO Parity Status

TABLE 97 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 98 NP memory errors supported on Gen-2 interface cards.

External Memory Errors	
1	PRAM Parity Errors
2	CAM2PRAM Parity Errors
3	LBLRAM /TXPRAM Parity Errors
4	AGERAM Parity Errors
5	CAM1 Interface Parity Errors
6	CAM2 Interface Parity Errors
7	CAM3 Interface Parity Errors
Internal Memory Errors	
1	Tx ACL PRAM Results FIFO Parity
2	Tx VLAN Result Parity
3	Tx Frame ctrl Parity
4	Tx Frame Data Parity
5	Tx nexthoptable lkup Data Parity

TABLE 98 NP memory errors supported on Gen-2 interface cards.

6	Stats Data Parity
7	Spix Multicast VLAN replace rData Parity
8	Rx Dispatch Parity
9	Rx MAC0 Data FIFO Parity
10	Rx MAC1 Data FIFO Parity
11	Rx MAC2 Data FIFO Parity
12	Rx MAC3 Data FIFO Parity
13	Rx MAC0 ctrl FIFO Parity
14	Rx MAC1 ctrl FIFO Parity
15	Rx MAC2 ctrl FIFO Parity
16	Rx MAC3 ctrl FIFO Parity
17	SPIO Multicast VLAN sopFIFO Parity
18	SPIO Multicast VLAN sfsFIFO Parity
19	SPIO Multicast VLAN repTable Parity
20	SPIO Multicast VLAN cpktFIFO Parity
21	SPIO Multicast VLAN Flag FIFO Parity
22	SPI1 Multicast VLAN sopFIFO Parity
23	SPI1 Multicast VLAN sfsFIFO Parity
24	SPI1 Multicast VLAN repTable Parity
25	SPI1 Multicast VLAN cpktFIFO Parity
26	SPI1 Multicast VLAN Flag FIFO Parity
27	SPI2 Multicast VLAN sopFIFO Parity
28	SPI2 Multicast VLAN sfsFIFO Parity
29	SPI2 Multicast VLAN repTable Parity
30	SPI2 Multicast VLAN cpktFIFO Parity
31	SPI2 Multicast VLAN Flag FIFO Parity
32	SPI3 Multicast VLAN sopFIFO Parity
33	SPI3 Multicast VLAN sfsFIFO Parity
34	SPI3 Multicast VLAN repTable Parity
35	SPI3 Multicast VLAN cpktFIFO Parity
36	SPI3 Multicast VLAN Flag FIFO Parity
37	Agezero Read Data Parity
38	CAM3 Async FIFO rbus Parity
39	CAM3 SyncFIFO rdData lo Parity
40	CAM3 SyncFIFO rdData hi Parity
41	CAM3 Lookup FIFO Parity
42	CAM2 Async FIFO rbus Parity

TABLE 98 NP memory errors supported on Gen-2 interface cards.

43	CAM2 SyncFIFO rdData lo Parity
44	CAM2 SyncFIFO rdData hi Parity
45	CAM2 Lookup FIFO Parity
46	CAM1 Async FIFO rbus Parity
47	CAM1 SyncFIFO rdData lo Parity
48	CAM1 SyncFIFO rdData hi Parity
49	CAM1 Lookup FIFO Parity
49	Eval0 Trunk group Table Parity
50	Eval1 Trunk group Table Parity
51	LBLPRAM Result Scheduler FIFO Parity
52	PRAM Result Scheduler FIFO Parity
53	Rx topotos FIFO Parity
54	Aged FIFO Read Data Parity
55	cpu2replace rdData Parity
56	Rxctrl FIFO Read Data Parity
57	config Read Data Parity
58	cmd rdData Parity
59	cmpl Data Parity
60	Rx pktdecode FIFO Parity
61	pkt Tablerd Parity[3:0]
62	PRAM Result FIFO0 Parity
63	PRAM Result FIFO1 Parity
64	PRAM Result FIFO2 Parity
65	PRAM Result FIFO3 Parity
66	Rx pkthdr FIFO Parity
67	Packet Table Read Parity
68	NextHop Table Read Data Parity
69	TxVLAN Read Data Parity
70	LBLikup Lookup FIFO Overflow
71	LBLRAM rdrequest FIFO Parity
72	LBLRAM Txp lkupFIFO Parity
73	LBLRAM LBL lkupFIFO Parity
74	LBLRAM cpu FIFO rdData Parity
75	PRAM ecmp FIFO rd Parity
76	PRAM rdrequest FIFO Parity
77	PRAM cpu rdFIFO Parity
78	Rx CAMResult FIFO Parity

TABLE 98 NP memory errors supported on Gen-2 interface cards.

79	CAM2PRAM mwFIFO Parity
80	CAM2PRAM rdrequest FIFO Parity
81	CAM2PRAM ecmp FIFO rdData Parity
82	CAM2PRAM cpu FIFO rdData Parity
83	sCAM Result ReadData Parity
84	mCAM Result ReadData Parity
85	LBLkup Lookup FIFO Underflow
86	Txplkup Lookup FIFO Underflow
87	Txplkup Lookup FIFO Overflow
88	LBLRAM rdrequest FIFO Underflow
89	LBLRAM rdrequest FIFO Overflow
90	MAC0 Frame ctrl Parity
91	MAC0 Frame Data Parity
92	MAC1 Frame ctrl Parity
93	MAC1 Frame Data Parity
94	MAC2 Frame ctrl Parity
95	MAC2 Frame Data Parity
96	MAC3 Frame ctrl Parity
97	MAC3 Frame Data Parity
98	Sp0 Tx Frame ctrl Parity
99	Sp0 Tx Frame Data Parity
100	Sp1 Tx Frame ctrl Parity
101	Sp1 Tx Frame Data Parity
102	Sp2 Tx Frame ctrl Parity
103	Sp2 Tx Frame Data Parity
104	Sp3 Tx Frame ctrl Parity
105	Sp3 Tx Frame Data Parity