

Brocade FastIron Administration Guide, 08.0.30p

Supporting FastIron Software Release 08.0.30p

Copyright Notice and Proprietary Information

Copyright 2017 Brocade Communications Systems, Inc. All rights reserved.

No part of this documentation may be used, reproduced, transmitted, or translated, in any form or by any means, electronic, mechanical, manual, optical, or otherwise, without prior written permission of or as expressly provided by under license from Brocade.

Destination Control Statement

Technical data contained in this publication may be subject to the export control laws of the United States of America. Disclosure to nationals of other countries contrary to United States law is prohibited. It is the reader's responsibility to determine the applicable regulations and to comply with them.

Disclaimer

THIS DOCUMENTATION AND ALL INFORMATION CONTAINED HEREIN ("MATERIAL") IS PROVIDED FOR GENERAL INFORMATION PURPOSES ONLY. BROCADE and RUCKUS WIRELESS, INC. AND THEIR LICENSORS MAKE NO WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, WITH REGARD TO THE MATERIAL, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT AND FITNESS FOR A PARTICULAR PURPOSE, OR THAT THE MATERIAL IS ERROR-FREE, ACCURATE OR RELIABLE. BROCADE and RUCKUS RESERVE THE RIGHT TO MAKE CHANGES OR UPDATES TO THE MATERIAL AT ANY TIME.

Limitation of Liability

IN NO EVENT SHALL BROCADE or RUCKUS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL OR CONSEQUENTIAL DAMAGES, OR DAMAGES FOR LOSS OF PROFITS, REVENUE, DATA OR USE, INCURRED BY YOU OR ANY THIRD PARTY, WHETHER IN AN ACTION IN CONTRACT OR TORT, ARISING FROM YOUR ACCESS TO, OR USE OF, THE MATERIAL.

Trademarks

Ruckus Wireless, Ruckus, the bark logo, BeamFlex, ChannelFly, Dynamic PSK, FlexMaster, Simply Better Wireless, SmartCell, SmartMesh, SmartZone, Unleashed, ZoneDirector and ZoneFlex are trademarks of Ruckus Wireless, Inc. in the United States and in other countries. Brocade, the B-wing symbol, MyBrocade, and ICX are trademarks of Brocade Communications Systems, Inc. in the United States and in other countries. Other trademarks may belong to third parties.

Contents

Preface	11
Document conventions.....	11
Notes, cautions, and warnings.....	11
Text formatting conventions.....	11
Command syntax conventions.....	12
Brocade resources.....	12
Document feedback.....	12
Contacting Brocade Technical Support.....	13
Brocade customers.....	13
Brocade OEM customers.....	13
About This Document	15
Supported hardware and software.....	15
What's new in this document	15
How command information is presented in this guide.....	16
Management Applications	17
Management port overview.....	17
How the management port works.....	17
CLI Commands for use with the management port.....	18
Web Management Interface.....	19
Management VRFs.....	19
Source interface and management VRF compatibility.....	20
Supported management applications.....	20
Configuring a global management VRF.....	22
Displaying management VRF information.....	23
Basic Software Features	27
Basic system parameter configuration.....	27
Entering system administration information.....	27
SNMP parameter configuration.....	28
Displaying virtual routing interface statistics.....	30
User-login details in Syslog messages and traps.....	31
Cancelling an outbound Telnet session.....	32
Network Time Protocol Version 4 (NTPv4).....	32
Limitations.....	35
Network Time Protocol (NTP) leap second	35
How Brocade supports leap second handling for NTP.....	35
NTP and SNTP.....	36
NTP server.....	36
NTP Client.....	37
NTP peer.....	37
NTP broadcast server.....	38
NTP broadcast client.....	38
NTP associations.....	39
Synchronizing time.....	40
Authentication.....	40
VLAN and NTP.....	40

Configuring NTP.....	40
Basic port parameter configuration.....	49
Specifying a port address.....	49
Assigning port names.....	51
Displaying the port name for an interface.....	52
Enabling auto-negotiation maximum port speed advertisement and down-shift.....	55
Configuring port speed down-shift and auto-negotiation for a range of ports.....	57
Enabling port speed down-shift.....	58
Force mode configuration.....	58
MDI and MDIX configuration.....	59
Disabling or re-enabling a port.....	60
Disabling laser light emission on port.....	60
Flow control configuration.....	61
Symmetric flow control on FCX and ICX devices.....	64
PHY FIFO Rx and Tx depth configuration.....	67
Interpacket Gap (IPG) on a FastIron X Series switch.....	67
IPG on FastIron Stackable devices.....	69
Enabling and disabling support for 100BaseTX.....	70
Enabling and disabling support for 100BaseFX.....	70
Changing the Gbps fiber negotiation mode.....	71
Port priority (QoS) modification.....	72
Dynamic configuration of Voice over IP (VoIP) phones.....	72
Port flap dampening configuration.....	73
Port loop detection.....	76
CLI banner configuration.....	82
Setting a message of the day banner.....	82
Requiring users to press the Enter key after the message of the day banner.....	83
Setting a privileged EXEC CLI level banner.....	84
Displaying a console message when an incoming Telnet session is detected.....	84
Operations, Administration, and Maintenance.....	85
OAM Overview.....	85
Software versions installed and running on a device.....	86
Determining the flash image version running on the device.....	86
Displaying the boot image version running on the device.....	87
Displaying the image versions installed in flash memory.....	88
Flash image verification	88
Software Image file types.....	89
Flash timeout.....	90
Software upgrades.....	90
Boot code synchronization feature.....	90
Viewing the contents of flash files.....	90
Using SNMP to upgrade software.....	92
Software reboot.....	93
Software boot configuration notes.....	93
Displaying the boot preference.....	94
Loading and saving configuration files.....	94
Replacing the startup configuration with the running configuration.....	95
Replacing the running configuration with the startup configuration.....	95
Logging changes to the startup-config file.....	95
Copying a configuration file to or from a TFTP server.....	96

Dynamic configuration loading.....	96
Maximum file sizes for startup-config file and running-config.....	98
Loading and saving configuration files with IPv6.....	99
Using the IPv6 copy command.....	99
Copying a file from an IPv6 TFTP server.....	100
IPv6 copy command.....	101
IPv6 TFTP server file upload.....	101
Using SNMP to save and load configuration information.....	102
Erasing image and configuration files.....	103
System reload scheduling.....	103
Reloading at a specific time.....	104
Reloading after a specific amount of time.....	104
Displaying the amount of time remaining before a scheduled reload.....	104
Canceling a scheduled reload.....	104
Diagnostic error codes and remedies for TFTP transfers.....	104
Network connectivity testing.....	106
Pinging an IPv4 address.....	106
Tracing an IPv4 route.....	107
IEEE 802.3ah EFM-OAM.....	108
Network deployment use case.....	108
EFM-OAM protocol.....	108
Process overview.....	109
Remote failure indication.....	110
Remote loopback.....	110
EFM-OAM error disable recovery	111
Configuring EFM-OAM.....	111
Displaying OAM information.....	112
Displaying OAM statistics.....	114
EFM-OAM syslog messages.....	116
Hitless management on the FSX 800 and FSX 1600.....	116
Benefits of hitless management.....	117
Supported protocols and services for hitless management events.....	117
Hitless management configuration notes and feature limitations.....	119
Hitless reload or switchover requirements and limitations.....	119
What happens during a Hitless switchover or failover.....	120
Enabling hitless failover on the FSX 800 and FSX 1600.....	122
Executing a hitless switchover on the FSX 800 and FSX 1600.....	122
Hitless OS upgrade on the FSX 800 and FSX 1600.....	123
Syslog message for Hitless management events.....	125
Displaying diagnostic information.....	125
Displaying management redundancy information	126
Layer 3 hitless route purge	126
Setting the IPv4 hitless purge timer on the default VRF.....	127
Example for setting IPv4 hitless purge timer on the default VRF.....	127
Setting the IPv4 hitless purge timer on the non-default VRF.....	127
Example for setting the IPv4 hitless purge timer on the non-default VRF.....	127
Setting the IPv6 hitless purge timer on the default VRF.....	127
Example for setting the IPv6 hitless purge timer on the default VRF.....	127
Setting the IPv4 hitless purge timer on the non-default VRF.....	127
Example for setting the IPv6 hitless purge timer on the non-default VRF.....	128

DHCP Client-Based Auto-Configuration and Flash image update.....	128
Configuration notes and feature limitations for DHCP Client-Based Auto-Configuration.....	129
Energy Efficient Ethernet.....	137
Port support for Energy Efficient Ethernet.....	137
Enabling Energy Efficient Ethernet.....	137
Histogram information overview.....	137
Displaying CPU histogram information.....	138
External USB Hotplug.....	138
Using External USB Hotplug.....	139
IPv6.....	141
Static IPv6 route configuration.....	141
Configuring a static IPv6 route.....	141
Configuring a static route in a non-default VRF or User VRF.....	142
IPv6 over IPv4 tunnels.....	143
IPv6 over IPv4 tunnel configuration notes.....	143
Configuring a manual IPv6 tunnel.....	144
Clearing IPv6 tunnel statistics.....	145
Displaying IPv6 tunnel information.....	145
SNMP Access.....	149
SNMP overview.....	149
SNMP community strings.....	149
Encryption of SNMP community strings	150
Adding an SNMP community string.....	150
Displaying the SNMP community strings.....	151
User-based security model.....	152
Configuring your NMS.....	152
Configuring SNMP version 3 on Brocade devices.....	152
Defining the engine id.....	153
Defining an SNMP group.....	153
Defining an SNMP user account.....	154
Defining SNMP views.....	155
SNMP version 3 traps.....	156
Defining an SNMP group and specifying which view is notified of traps.....	157
Defining the UDP port for SNMP v3 traps.....	157
Trap MIB changes.....	158
Specifying an IPv6 host as an SNMP trap receiver.....	158
SNMP v3 over IPv6.....	158
Specifying an IPv6 host as an SNMP trap receiver	159
Viewing IPv6 SNMP server addresses.....	159
Displaying SNMP Information.....	159
Displaying the Engine ID.....	159
Displaying SNMP groups.....	160
Displaying user information.....	160
Interpreting varbinds in report packets.....	160
SNMP v3 configuration examples.....	161
Example 1.....	161
Example 2.....	161
Foundry Discovery Protocol (FDP) and Cisco Discovery Protocol (CDP) Packets	163
FDP overview.....	163

FDP configuration.....	163
Displaying FDP information.....	165
Clearing FDP and CDP information.....	167
CDP packets.....	167
Enabling interception of CDP packets globally.....	168
Enabling interception of CDP packets on an interface.....	168
Displaying CDP information.....	168
Clearing CDP information.....	170
LLDP and LLDP-MED.....	171
LLDP terms used in this chapter.....	171
LLDP overview.....	172
Benefits of LLDP.....	173
LLDP-MED overview.....	173
Benefits of LLDP-MED.....	174
LLDP-MED class.....	175
General LLDP operating principles.....	175
LLDP operating modes.....	175
LLDP packets.....	176
TLV support.....	176
MIB support.....	179
Syslog messages.....	180
LLDP configuration.....	180
LLDP configuration notes and considerations.....	180
Enabling and disabling LLDP.....	181
Enabling support for tagged LLDP packets.....	181
Changing a port LLDP operating mode.....	181
Configuring LLDP processing on 802.1x blocked port.....	183
Maximum number of LLDP neighbors	183
Enabling LLDP SNMP notifications and Syslog messages.....	184
Changing the minimum time between LLDP transmissions.....	185
Changing the interval between regular LLDP transmissions.....	185
Changing the holdtime multiplier for transmit TTL.....	185
Changing the minimum time between port reinitializations.....	186
LLDP TLVs advertised by the Brocade device.....	186
LLDP-MED configuration.....	192
Enabling LLDP-MED.....	193
Enabling SNMP notifications and Syslog messages for LLDP-MED topology changes.....	193
Changing the fast start repeat count.....	193
Defining a location id.....	194
Defining an LLDP-MED network policy.....	200
LLDP-MED attributes advertised by the Brocade device.....	201
LLDP-MED capabilities.....	202
Extended power-via-MDI information.....	202
Displaying LLDP statistics and configuration settings.....	204
LLDP configuration summary.....	204
Displaying LLDP statistics.....	205
Displaying LLDP neighbors.....	206
Displaying LLDP neighbors detail.....	207
Displaying LLDP configuration details.....	207
Resetting LLDP statistics.....	209

Clearing cached LLDP neighbor information.....	209
Hardware Component Monitoring.....	211
Traffic Limitations in Mixed Environments.....	211
Virtual cable testing.....	211
Virtual Cable Testing configuration notes.....	211
Virtual Cable Test command syntax.....	212
Viewing the results of the cable analysis.....	212
Digital optical monitoring.....	214
Digital optical monitoring configuration limitations.....	214
Enabling digital optical monitoring.....	214
Setting the alarm interval.....	214
Displaying information about installed media.....	215
Viewing optical monitoring information.....	216
Syslog messages for optical transceivers.....	218
FastIron Fiber-optic Transceivers.....	219
Network Monitoring.....	221
Basic system management.....	221
Viewing system information.....	221
Viewing configuration information.....	222
Enabling the display of the elapsed timestamp for port statistics reset.....	223
Viewing port statistics.....	223
Viewing STP statistics.....	226
Clearing statistics.....	226
Traffic counters for outbound traffic	227
Viewing egress queue counters on ICX 6610 and FCX devices.....	229
Viewing egress queue counters on ICX 7750 devices.....	230
Clearing the egress queue counters.....	231
RMON support.....	231
Maximum number of entries allowed in the RMON control table.....	231
Statistics (RMON group 1).....	232
History (RMON group 2).....	234
Alarm (RMON group 3).....	234
Event (RMON group 9).....	235
sFlow.....	235
sFlow version 5.....	235
sFlow support for IPv6 packets.....	236
sFlow configuration considerations.....	237
Configuring and enabling sFlow.....	239
Enabling sFlow forwarding.....	243
sFlow version 5 feature configuration.....	245
Configuring sFlow with Multi-VRFs.....	248
Displaying sFlow information.....	249
Utilization list for an uplink port.....	251
Utilization list for an uplink port command syntax.....	251
Displaying utilization percentages for an uplink.....	252
System Monitoring.....	253
Overview of system monitoring.....	253
Configuration notes and feature limitations.....	253
Configure system monitoring.....	254

disable system-monitoring all	254
enable system-monitoring all	254
sysmon timer	254
sysmon log-backoff	255
sysmon threshold	255
System monitoring on FCX and ICX devices.....	256
sysmon ecc-error	256
sysmon link-error	256
System monitoring for Fabric Adapters.....	257
sysmon fa error-count	257
sysmon fa link	258
System monitoring for Cross Bar.....	259
sysmon xbar error-count	259
sysmon xbar link	260
System monitoring for Packet Processors.....	260
sysmon pp error-count	261
clear sysmon counters	261
show sysmon logs	262
show sysmon counters	263
show sysmon config	266
show sysmon system sfm	267
Syslog.....	269
About Syslog messages.....	269
Displaying Syslog messages.....	269
Enabling real-time display of Syslog messages.....	270
Enabling real-time display for a Telnet or SSH session.....	270
Broadcast, unknown unicast, and multicast suppression Syslog and SNMP notification.....	271
Displaying real-time Syslog messages	272
Syslog service configuration.....	272
Displaying the Syslog configuration.....	273
Generating the Syslog specific to RFC 5424.....	276
Disabling or re-enabling Syslog.....	278
Specifying a Syslog server.....	278
Specifying an additional Syslog server.....	278
Disabling logging of a message level.....	278
Changing the number of entries the local buffer can hold.....	279
Changing the log facility.....	279
Displaying interface names in Syslog messages.....	280
Retaining Syslog messages after a soft reboot.....	280
Clearing the Syslog messages from the local buffer.....	281
Syslog messages.....	283
Brocade Syslog messages.....	283
Power over Ethernet	311
Power over Ethernet overview.....	311
Power over Ethernet terms used in this chapter.....	311
Methods for delivering Power over Ethernet.....	312
PoE autodiscovery.....	314
Power class.....	314
Dynamic upgrade of PoE power supplies.....	315

Power over Ethernet cabling requirements.....	317
Supported powered devices.....	317
Installing PoE firmware	318
PoE and CPU utilization.....	322
Enabling and disabling Power over Ethernet.....	322
Disabling support for PoE legacy power-consuming devices.....	323
Enabling the detection of PoE power requirements advertised through CDP.....	324
Command syntax for PoE power requirements.....	324
Setting the maximum power level for a PoE power-consuming device.....	325
Considerations for setting power levels.....	325
Configuring power levels command syntax.....	325
Setting the power class for a PoE power-consuming device.....	326
Setting the power class command syntax.....	326
Setting the power budget for a PoE interface module on an FSX device.....	327
Setting the inline power priority for a PoE port	327
Command syntax for setting the inline power priority for a PoE port.....	328
Resetting PoE parameters.....	329
Displaying Power over Ethernet information.....	329
Displaying PoE operational status	330
Displaying PoE data specific to PD ports	332
Displaying detailed information about PoE power supplies.....	334
Inline power on PoE LAG ports.....	340
Restriction.....	341
Configuring inline power on PoE ports in a LAG.....	341
Decouple PoE and datalink operations on PoE ports.....	342
Restriction.....	342
Decoupling of PoE and datalink operations on PoE LAG ports.....	343
Decoupling of PoE and datalink operations on regular PoE ports.....	344
40 Gbps Breakout Ports.....	347
Overview of 40 Gbps breakout ports.....	347
Configuring 40 Gbps breakout ports.....	347
Configuring sub-ports.....	348
Removing breakout configuration.....	350
Displaying information for breakout ports.....	352
OpenSSL License.....	353
OpenSSL license.....	353
Original SSLeay License.....	353
Joint Interoperability Test Command.....	355
JITC overview.....	355
AES-CTR encryption mode support for SSH.....	355
SHA1 authentication support for NTP.....	355
IPv6 ACL for SNMPv3 group.....	355

Preface

- Document conventions..... 11
- Brocade resources..... 12
- Document feedback..... 12
- Contacting Brocade Technical Support..... 13

Document conventions

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

Notes, cautions, and warnings

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

NOTE

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

ATTENTION

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



CAUTION

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



DANGER

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

Text formatting conventions

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

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

Format	Description
	Identifies command syntax examples.

Command syntax conventions

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

Convention	Description
bold text	Identifies command names, keywords, and command options.
<i>italic text</i>	Identifies a variable.
value	In Fibre Channel products, a fixed value provided as input to a command option is printed in plain text, for example, <code>--show WWN</code> .
[]	Syntax components displayed within square brackets are optional. Default responses to system prompts are enclosed in square brackets.
{ x y z }	A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options. In Fibre Channel products, square brackets may be used instead for this purpose.
x y	A vertical bar separates mutually exclusive elements.
< >	Nonprinting characters, for example, passwords, are enclosed in angle brackets.
...	Repeat the previous element, for example, <code>member[member...]</code> .
\	Indicates a "soft" line break in command examples. If a backslash separates two lines of a command input, enter the entire command at the prompt without the backslash.

Brocade resources

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

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

Click the **Support** tab and select **Document Library** to access product documentation on MyBrocade or www.brocade.com. You can locate documentation by product or by operating system.

Release notes are bundled with software downloads on MyBrocade. Links to software downloads are available on the MyBrocade landing page and in the Document Library.

Document feedback

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

- Through the online feedback form in the HTML documents posted on www.brocade.com
- By sending your feedback to documentation@brocade.com

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

Contacting Brocade Technical Support

As a Brocade customer, you can contact Brocade Technical Support 24x7 online or by telephone. Brocade OEM customers should contact their OEM/solution provider.

Brocade customers

For product support information and the latest information on contacting the Technical Assistance Center, go to www.brocade.com and select **Support**.

If you have purchased Brocade product support directly from Brocade, use one of the following methods to contact the Brocade Technical Assistance Center 24x7.

Online	Telephone
<p>Preferred method of contact for non-urgent issues:</p> <ul style="list-style-type: none"> • Case management through the MyBrocade portal. • Quick Access links to Knowledge Base, Community, Document Library, Software Downloads and Licensing tools 	<p>Required for Sev 1-Critical and Sev 2-High issues:</p> <ul style="list-style-type: none"> • Continental US: 1-800-752-8061 • Europe, Middle East, Africa, and Asia Pacific: +800-AT FIBREE (+800 28 34 27 33) • Toll-free numbers are available in many countries. • For areas unable to access a toll-free number: +1-408-333-6061

Brocade OEM customers

If you have purchased Brocade product support from a Brocade OEM/solution provider, contact your OEM/solution provider for all of your product support needs.

- OEM/solution providers are trained and certified by Brocade to support Brocade® products.
- Brocade provides backline support for issues that cannot be resolved by the OEM/solution provider.
- Brocade Supplemental Support augments your existing OEM support contract, providing direct access to Brocade expertise. For more information, contact Brocade or your OEM.
- For questions regarding service levels and response times, contact your OEM/solution provider.

About This Document

- Supported hardware and software.....15
- What's new in this document15
- How command information is presented in this guide..... 16

Supported hardware and software

This guide supports the following product families for the FastIron 08.0.30 release:

- FCX Series
- FastIron X Series (FastIron SX 800, FastIron SX 1600)
- ICX 6610 Series
- ICX 6430 Series (ICX 6430, ICX 6430-C12)
- ICX 6450 Series (ICX 6450, ICX 6450-C12-PD)
- ICX 6650 Series
- ICX 7250 Series
- ICX 7450 Series
- ICX 7750 Series

NOTE

The Brocade ICX 6430-C switch supports the same feature set as the Brocade ICX 6430 switch unless otherwise noted.

NOTE

The Brocade ICX 6450-C12-PD switch supports the same feature set as the Brocade ICX 6450 switch unless otherwise noted.

For information about the specific models and modules supported in a product family, refer to the hardware installation guide for that product family.

What's new in this document

The following tables include descriptions of new information added to this guide for the FastIron 08.0.30p release.

TABLE 1 Summary of enhancements in FastIron release 08.0.30p

Feature	Description	Described in
Simultaneous Firmware download	On ICX 7250 and ICX 7450 devices, PoE Firmware download can be initiated on all PoE units or multiple stacks simultaneously.	Installing PoE firmware on page 318

How command information is presented in this guide

For all new content supported in FastIron release 08.0.20 and later, command information is documented in a standalone command reference guide.

In the *Brocade FastIron Command Reference*, the command pages are in alphabetical order and follow a standard format to present syntax, parameters, mode, usage guidelines, examples, and command history.

NOTE

Many commands introduced before FastIron release 08.0.20 are also included in the guide.

Management Applications

• Management port overview.....	17
• Web Management Interface.....	19
• Management VRFs.....	19

Management port overview

NOTE

The management port applies to FCX, SX 800, SX 1600, ICX 6430, and ICX 6450 devices.

The management port is an out-of-band port that customers can use to manage their devices without interfering with the in-band ports. The management port is widely used to download images and configurations, for Telnet sessions and for Web management.

For FCX devices, the MAC address for the management port is derived from the base MAC address of the unit, plus the number of ports in the base module. For example, on a 48-port FCX standalone device, the base MAC address is 0000.0034.2200. The management port MAC address for this device would be 0000.0034.2200 plus 0x30, or 0000.0034.2230. The 0x30 in this case equals the 48 ports on the base module.

For SX 800 and SX 1600 devices, the MAC address for the management port is derived as if the management port is the last port on the management module where it is located. For example, on a 2 X 10G management module, the MAC address of the management port is that of the third port on that module.

How the management port works

The following rules apply to management ports:

- Only packets that are specifically addressed to the management port MAC address or the broadcast MAC address are processed by the Layer 2 switch or Layer 3 switch. All other packets are filtered out.
- No packet received on a management port is sent to any in-band ports, and no packets received on in-band ports are sent to a management port.
- A management port is not part of any VLAN
- Configuring a strict management VRF disables certain features on the management port.
- Protocols are not supported on the management port.
- Creating a management VLAN disables the management port on the device.
- For FCX and ICX devices, all features that can be configured from the global configuration mode can also be configured from the interface level of the management port. Features that are configured through the management port take effect globally, not on the management port itself.

For switches, any in-band port may be used for management purposes. A router sends Layer 3 packets using the MAC address of the port as the source MAC address.

For stacking devices, (for example, an FCX stack) each stack unit has one out-of band management port. Only the management port on the Active Controller will actively send and receive packets. If a new Active Controller is elected, the new Active Controller management port will become the active management port. In this situation, the MAC address of the old Active Controller and the MAC address of the new controller will be different.

CLI Commands for use with the management port

The following CLI commands can be used with a management port.

To display the current configuration, use the **show running-config interface management** command.

Syntax: show running-config interface management *num*

```
device(config-if-mgmt)# ip addr 10.44.9.64/24
device(config)# show running-config interface management 1
interface management 1
ip address 10.44.9.64 255.255.255.0
```

To display the current configuration, use the **show interfaces management** command.

Syntax: show interfaces management *num*

```
device(config)# show interfaces management 1
GigEthernetmgmt1 is up, line protocol is up
Hardware is GigEthernet, address is 0000.0076.544a (bia 0000.0076.544a)
Configured speed auto, actual 1Gbit, configured duplex fdx, actual fdx
Configured mdi mode AUTO, actual none
BPRU guard is disabled, ROOT protect is disabled
Link Error Dampening is Disabled
STP configured to OFF, priority is level0, MAC-learning is enabled
Flow Control is config disabled, oper enabled
Mirror disabled, Monitor disabled
Not member of any active trunks
Not member of any configured trunks
No port name
IPG MII 0 bits-time, IPG GMII 0 bits-time
IP MTU 1500 bytes
300 second input rate: 83728 bits/sec, 130 packets/sec, 0.01% utilization
300 second output rate: 24 bits/sec, 0 packets/sec, 0.00% utilization
39926 packets input, 3210077 bytes, 0 no buffer
Received 4353 broadcasts, 32503 multicasts, 370 unicasts
0 input errors, 0 CRC, 0 frame, 0 ignored
0 runts, 0 giants
22 packets output, 1540 bytes, 0 underruns
Transmitted 0 broadcasts, 6 multicasts, 16 unicasts
0 output errors, 0 collisions
```

To display the management interface information in brief form, enter the **show interfaces brief management** command.

Syntax: show interfaces brief management *num*

```
device# show interfaces brief management 1
Port  Link  State Dupl  Speed Trunk  Tag  Pri  MAC  Name
mgmt1 Up      None  Full  1G      None  No   0    0000.0076.544a
```

To display management port statistics, enter the **show statistics management** command.

Syntax: show statistics management *num*

```
device# show statistics management 1
Port  Link  State Dupl  Speed Trunk  Tag  Pri  MAC  Name
mgmt1 Up      None  Full  1G      None  No   0    0000.0076.544a
Port mgmt1 Counters:
  InOctets  3210941  OutOctets  1540
  InPkts   39939   OutPackets  22
InBroadcastPkts  4355  OutbroadcastPkts  0
InMultiastPkts  35214  OutMulticastPkts  6
InUnicastPkts  370  OutUnicastPkts  16
InBadPkts  0
InFragments  0
InDiscards  0  OutErrors  0
CRC  0  Collisions  0
InErrors  0  LateCollisions  0
InGiantPkts  0
```

```

InShortPkts      0
InJabber         0
InFlowCtrlPkts  0   OutFlowCtrlPkts  0
InBitsPerSec    83728   OutBitsPerSec    24
InPktsPerSec    130   OutPktsPerSec    0
InUtilization   0.01%   OutUtilization   0.00%

```

To display the management interface statistics in brief form, enter the **show statistics brief management** command.

Syntax: **show statistics brief management num**

```

device(config)# show statistics brief management 1
Port      In Packets      Out PacketsTrunk      In Errors      Out Errors
mgmt1     39946          22      0      0
Total     39945          22      0      0

```

Web Management Interface

The Web Management Interface is a browser-based interface that allows administrators to manage and monitor a single Brocade device or a group of Brocade devices connected together.

For many of the features on a Brocade device, the Web Management Interface can be used as an alternate to the CLI for creating new configurations, modifying existing ones, and monitoring the traffic on a device.

For more information on how to log in and use the Web Management Interface, refer to the *FastIron SX, FCX, and ICX Web Management Interface User Guide*.

Management VRFs

Virtual routing and forwarding (VRF) allows routers to maintain multiple routing tables and forwarding tables on the same router. A management VRF can be configured to control the flow of management traffic as described in this section.

NOTE

For information on configuring Multi-VRF, sometimes called VRF-Lite or Multi-VRF CE, refer to the *FastIron Ethernet Switch Layer 3 Routing Configuration Guide*.

A management VRF is used to provide secure management access to the device by sending inbound and outbound management traffic through the VRF specified as a global management VRF and through the out-of-band management port, thereby isolating management traffic from the network data traffic.

By default, the inbound traffic is unaware of VRF and allows incoming packets from any VRF, including the default VRF. Outbound traffic is sent only through the default VRF. The default VRF consists of an out-of-band management port and all the LP ports that do not belong to any other VRFs.

Any VRF, except the default VRF, can be configured as a management VRF. When a management VRF is configured, the management traffic is allowed through the ports belonging to the specified VRF and the out-of-band management port. The management traffic through the ports belonging to the other VRFs and the default VRF are dropped, and the rejection statistics are incremented.

If the management VRF is not configured, the management applications follows default behavior. The management VRF is configured the same way for IPv4 and IPv6 management traffic.

A management VRF instance cannot be configured on any of the following:

- Physical interface
- Management interface

The management VRF is supported by the following management applications:

- SNMP server
- SNMP trap generator
- Telnet server
- SSH server
- Telnet client
- RADIUS client
- TACACS+ client
- TFTP
- SCP
- Syslog

NOTE

Any **ping** or **traceroute** commands use the VRF specified in the command or the default VRF if no VRF is specified.

Source interface and management VRF compatibility

A source interface must be configured for management applications. When a source interface is configured, management applications use the lowest configured IP address of the specified interface as the source IP address in all the outgoing packets. If the configured interface is not part of the management VRF, the response packet does not reach the destination. If the compatibility check fails while either the management VRF or the source interface is being configured, the following warning message is displayed. However, the configuration command is accepted.

```
The source-interface for Telnet, TFTP is not part of the management-vrf
```

Supported management applications

This section explains the management VRF support provided by the management applications.

SNMP server

When the management VRF is configured, the SNMP server receives SNMP requests and sends SNMP responses only through the ports belonging to the management VRF and through the out-of-band management port.

Any change in the management VRF configuration becomes immediately effective for the SNMP server.

SNMP trap generator

When the management VRF is configured, the SNMP trap generator sends traps to trap hosts through the ports belonging to the management VRF and through the out-of-band management port.

Any change in the management VRF configuration takes effect immediately for the SNMP trap generator.

NOTE

The SNMP source interface configuration command **snmp-server trap-source** must be compatible with the management VRF configuration.

SSH server

When the management VRF is configured, the incoming SSH connection requests are allowed only from the ports belonging to the management VRF and from the out-of-band management port. Management VRF enforcement occurs only while a connection is established.

To allow the incoming SSH connection requests only from the management VRF and not from the out-of-band management port, enter the following command.

```
device(config)# ip ssh strict-management-vrf
```

The **ip ssh strict-management-vrf** command is applicable only when the management VRF is configured. If not, the command issues the following warning message.

```
Warning - Management-vrf is not configured.
```

For the SSH server, changing the management VRF configuration or configuring the **ip ssh strict-management-vrf** command does not affect the existing SSH connections. The changes are applied only to new incoming connection requests.

Telnet client

To allow the incoming Telnet connection requests only from the management VRF and not from the out-of-band management port, enter the following command.

```
device(config)# telnet strict-management-vrf
```

Syntax: telnet strict-management-vrf

RADIUS client

When the management VRF is configured, the RADIUS client sends RADIUS requests or receives responses only through the ports belonging to the management VRF and through the out-of-band management port.

Any change in the management VRF configuration takes effect immediately for the RADIUS client.

NOTE

The RADIUS source interface configuration command **ip radius source-interface** must be compatible with the management VRF configuration.

TACACS+ client

When the management VRF is configured, the TACACS+ client establishes connections with TACACS+ servers only through the ports belonging to the management VRF and the out-of-band management port.

For the TACACS+ client, a change in the management VRF configuration does not affect the existing TACACS+ connections. The changes are applied only to new TACACS+ connections.

NOTE

The TACACS+ source interface configuration command **ip tacacs source-interface** must be compatible with the management VRF configuration.

TFTP

When the management VRF is configured, TFTP sends or receives data and acknowledgments only through ports belonging to the management VRF and through the out-of-band management port.

Any change in the management VRF configuration takes effect immediately for TFTP. You cannot change in the management VRF configuration while TFTP is in progress.

NOTE

The TFTP source interface configuration command **ip tftp source-interface** must be compatible with the management VRF configuration.

SCP

SCP uses SSH as the underlying transport. The behavior of SCP is similar to the SSH server.

Syslog

When the management VRF is configured, the Syslog module sends log messages only through the ports belonging to the management VRF and the out-of-band management port.

Any change in the management VRF configuration takes effect immediately for Syslog.

NOTE

The Syslog source interface configuration command **ip syslog source-interface** must be compatible with the management VRF configuration.

Configuring a global management VRF

To configure a VRF as a global management VRF, enter the following command.

```
device(config)# management-vrf mvrf
```

Syntax: **[no] management-vrf** *vrf-name*

The *vrf-name* parameter must specify the name of a pre-configured VRF. If the VRF is not pre-configured, command execution fails, and the following error message is displayed.

```
Error - VRF <vrf-name>
doesn't exist
```

When the management VRF is configured, the following Syslog message is displayed.

```
SYSLOG: VRF <vrf-name>
has been configured as management-vrf
```

Enter the **no** form of the command to remove the management VRF. When the management VRF is deleted, the following Syslog message is displayed.

```
SYSLOG: VRF <vrf-name>
has been un-configured as management-vrf
```

Configuration notes

Consider the following configuration notes:

- If a management VRF is already configured, you must remove the existing management VRF configuration before configuring a new one. If not, the system displays the following error message.

```
device(config)# management-vrf red
Error - VRF mvrf already configured as management-vrf
```

- If you try to delete a management VRF that was not configured, the system displays the following error message.

```
device(config)# no management-vrf red
Error - VRF red is not the current management-vrf
```

- If a VRF is currently configured as the management VRF, it cannot be deleted or modified. Attempting to do so causes the system to return the following error message.

```
device(config)# no vrf mvrf
Error - Cannot modify/delete a VRF which is configured as management-vrf
```

Displaying management VRF information

To display IP Information for a specified VRF, enter the following command at any level of the CLI.

```
device(config)# show vrf mvrf
VRF mvrf, default RD 1100:1100, Table ID 11
Configured as management-vrf
IP Router-Id: 1.0.0.1
  Interfaces:
    ve3300 ve3400
  Address Family IPv4
    Max Routes: 641
    Number of Unicast Routes: 2
  Address Family IPv6
    Max Routes: 64
    Number of Unicast Routes: 2
```

Syntax: `show vrf vrf-name`

The *vrf-name* parameter specifies the VRF for which you want to display IP information.

TABLE 2 show vrf output descriptions

This field	Displays
VRF <i>vrf-name</i>	The name of the VRF.
default RD	The default route distinguisher for the VRF.
Table ID	The table ID for the VRF.
Routes	The total number of IPv4 and IPv6 Unicast routes configured on this VRF.
Configured as management-vrf	Indicates that the specified VRF is configured as a management VRF.
IP Router-Id	The 32-bit number that uniquely identifies the router.
Number of Unicast Routes	The number of Unicast routes configured on this VRF.

The **show who** command displays information about the management VRF from which the Telnet or SSH connection has been established.

```
device(config)# show who
Console connections:
  established, monitor enabled, privilege super-user, in config mode
  1 minutes 47 seconds in idle
```

```

Telnet server status: Enabled
Telnet connections (inbound):
 1      established, client ip address 10.53.1.181, user is lab, privilege super-user
        using vrf default-vrf.
        2 minutes 46 seconds in idle
 2      established, client ip address 10.20.20.2, user is lab, privilege super-user
        using vrf mvrf.
        16 seconds in idle
 3      closed
 4      closed
 5      closed
Telnet connections (outbound):
 6      established, server ip address 10.20.20.2, from Telnet session 2, , privilege super-user
        using vrf mvrf.
        12 seconds in idle
 7      closed
 8      closed
 9      closed
10      closed
SSH server status: Enabled
SSH connections:
 1      established, client ip address 10.53.1.181, privilege super-user
        using vrf default-vrf.
        you are connecting to this session
        3 seconds in idle
 2      established, client ip address 10.20.20.2, privilege super-user
        using vrf mvrf.
        48 seconds in idle
 3      closed
 4      closed
 5      closed
 6      closed
 7      closed
 8      closed
 9      closed
10      closed
11      closed
12      closed
13      closed
14      closed
15      closed
16      closed
    
```

Syntax: show who

To display packet and session rejection statistics due to failure in management VRF validation, enter the following command.

```

device(config)# show management-vrf

Management VRF name : sflow
Management Application      Rx Drop Pkts      Tx Drop Pkts
SNMP Engine                 0                  11
RADIUS Client                0                  0
TFTP Client                  0                  0
Traps                        -                  0
SysLogs                      -                  0

TCP Connection rejects:
Telnet      :          0
SSH        (Strict):    685
TACACS+ Client :          0
    
```

Syntax: show management-vrf

TABLE 3 show management-vrf output descriptions

This field	Displays
Management VRF name	Displays the configured management VRF name.
Management Application	Displays the management application names.

TABLE 3 show management-vrf output descriptions (continued)

This field	Displays
Rx Drop Pkts	Displays the number of packets dropped in the inbound traffic.
Tx Drop Pkts	Displays the number of packets dropped in the outbound traffic.
TCP Connection rejects	Displays the number of TCP connections per application rejected due to management VRF validation.

Make sure that the management VRF is configured before executing the **show management-vrf** command. If not, the system displays the following error message.

```
Error - Management VRF is not configured.
```

To clear the management VRF rejection statistics, enter the following command.

```
device(config)# clear management-vrf-stats
```

Syntax: clear management-vrf-stats

Basic Software Features

- Basic system parameter configuration..... 27
- Network Time Protocol Version 4 (NTPv4)..... 32
- Basic port parameter configuration..... 49
- CLI banner configuration..... 82
- Requiring users to press the Enter key after the message of the day banner..... 83
- Setting a privileged EXEC CLI level banner..... 84
- Displaying a console message when an incoming Telnet session is detected..... 84

Basic system parameter configuration

Brocade devices are configured at the factory with default parameters that allow you to begin using the basic features of the system immediately. However, many of the 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. If you use the Command Line Interface (CLI) to configure system parameters, you can find these system level parameters at the Global CONFIG level of the CLI.

NOTE

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

NOTE

For information about configuring IP addresses, DNS resolver, DHCP assist, and other IP-related parameters, refer to "IP Configuration" chapter in the *FastIron Ethernet Switch Layer 3 Routing Configuration Guide*

NOTE

For information about the Syslog buffer and messages, refer to [Basic system parameter configuration](#).

The procedures in this section describe how to configure the basic system parameters listed in [Basic Software Features](#) on page 27.

Entering system administration information

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

The name, contact, and location each can be up to 255 alphanumeric characters.

Here is an example of how to configure a system name, system contact, and location.

```
device(config)# hostname zappa
zappa(config)# snmp-server contact Support Services
zappa(config)# snmp-server location Centerville
zappa(config)# end
zappa# write memory
```

Syntax:hostname *string*

Syntax: snmp-server contact *string*

Syntax: snmp-server location *string*

The text strings can contain blanks. The SNMP text strings do not require quotation marks when they contain blanks but the host name does.

NOTE

The **chassis name** command does not change the CLI prompt. Instead, the command assigns an administrative ID to the device.

SNMP parameter configuration

Use the procedures in this section to perform the following configuration tasks:

- Specify a Simple Network Management Protocol (SNMP) trap receiver.
- Specify a source address and community string for all traps sent by the device.
- 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/TACACS+ server.

NOTE

To add and modify "get" (read-only) and "set" (read-write) community strings, refer to "Security Access" chapter in the *FastIron Ethernet Switch Security Configuration Guide*.

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.

When you add a trap receiver, the software automatically encrypts the community string you associate with the receiver when the string is displayed by the CLI or Web Management Interface. If you want the software to show the community string in the clear, you must explicitly specify this when you add a trap receiver. In either case, the software does not encrypt the string in the SNMP traps sent to the receiver.

To specify the host to which the device sends all SNMP traps, use one of the following methods.

To add a trap receiver and encrypt the display of the community string, enter commands such as the following.

To specify an SNMP trap receiver and change the UDP port that will be used to receive traps, enter a command such as the following.

```
device(config)# snmp-server host 10.2.2.2 0 mypublic port 200
device(config)# write memory
```

Syntax: `snmp-server host ip-addr { 0 | 1 } string [port value]`

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

The *0/1* parameter specifies whether you want the software to encrypt the string (**1**) or show the string in the clear (**0**). The default is **0**.

The *string* parameter specifies an SNMP community string configured on the Brocade device. The string can be a read-only string or a read-write string. The string is not used to authenticate access to the trap host but is instead 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 different Brocade devices based on the community strings.

The command in the example above adds trap receiver 10.2.2.2 and configures the software to encrypt display of the community string. When you save the new community string to the startup-config file (using the **write memory** command), the software adds the following command to the file.

```
snmp-server host 10.2.2.2 1
encrypted-string
```

To add a trap receiver and configure the software to encrypt display of the community string in the CLI and Web Management Interface, enter commands such as the following.

```
device(config)# snmp-server host 10.2.2.2 0 FastIron-12
device(config)# write memory
```

The *port value* parameter allows you to specify which UDP port will be used by the trap receiver. This parameter allows you to configure several trap receivers in a system. With this parameter, a network management application can coexist in the same system. 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 Layer 3 switch use the same source IP address. For configuration details, refer to "Specifying a single source interface for specified packet types" section in the *FastIron Ethernet Switch Layer 3 Routing Configuration Guide*.

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 device might not be able to reach the servers, in which case the messages are lost.

By default, a 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 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.

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

The command in this example changes the holddown time for SNMP traps to 30 seconds. The 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 seconds`

The `secs` parameter specifies the number of seconds and can be from 1 - 600 (ten minutes). The default is 60 seconds.

Disabling SNMP traps

Brocade devices come with SNMP trap generation enabled by default for all traps. You can selectively disable one or more of the following traps.

NOTE

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

SNMP Layer 2 traps

The following traps are generated on devices running Layer 2 software:

- SNMP authentication keys
- Power supply failure
- Fan failure
- Cold start

- Link up
- Link down
- Bridge new root
- Bridge topology change
- Locked address violation

SNMP Layer 3 traps

The following traps are generated on devices running Layer 3 software:

- SNMP authentication key
- Power supply failure
- Fan failure
- Cold start
- Link up
- Link down
- Bridge new root
- Bridge topology change
- Locked address violation
- BGP4
- OSPF
- VRRP
- VRRP-E

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

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

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

SNMP ifIndex

On Brocade IronWare 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 64. On all IronWare devices, the system automatically assign 64 indexes to each module on the device. This value is not configurable.

Displaying virtual routing interface statistics

NOTE

This feature is supported on FastIron X Series and ICX 6650 devices only.

You can enable SNMP to extract and display virtual routing interface statistics from the ifXTable (64-bit counters).

The following describes the limitations of this feature:

- The Brocade device counts traffic from all virtual interfaces (VEs). For example, in a configuration with two VLANs (VLAN 1 and VLAN 20) on port 1, when traffic is sent on VLAN 1, the counters (VE statistics) increase for both VE 1 and VE 20.
- The counters include all traffic on each virtual interface, even if the virtual interface is disabled.

- The counters include traffic that is denied by ACLs or MAC address filters.

To enable SNMP to display VE statistics, enter the **enable snmp ve-statistics** command.

```
device(config)# enable snmp ve-statistics
```

Syntax: [no] enable snmp ve-statistics

Use the **no** form of the command to disable this feature once it is enabled.

Note that the above CLI command enables SNMP to display virtual interface statistics. It does not enable the CLI or Web Management Interface to display the statistics .

User-login details in Syslog messages and traps

Brocade devices send Syslog messages and SNMP traps when a user logs into or out of the User EXEC or Privileged EXEC level of the CLI. The feature applies to users whose access is authenticated by an authentication-method list based on a local user account, RADIUS server, or TACACS/TACACS+ server.

To view the user-login details in the Syslog messages and traps, you must enable the **logging enable user-login** command.

```
device(config)# logging enable user-login
```

Syntax: [no] logging enable user-login

NOTE

The Privileged EXEC level is sometimes called the "Enable" level, because the command for accessing this level is **enable**.

Examples of Syslog messages for CLI access

When a user whose access is authenticated by a local user account, a RADIUS server, or a TACACS or TACACS+ server logs into or out of the CLI User EXEC or Privileged EXEC mode, the software generates a Syslog message and trap containing the following information:

- The time stamp
- The user name
- Whether the user logged in or out
- The CLI level the user logged into or out of (User EXEC or Privileged EXEC level)

NOTE

Messages for accessing the User EXEC level apply only to access through Telnet. The device does not authenticate initial access through serial connections but does authenticate serial access to the Privileged EXEC level. Messages for accessing the Privileged EXEC level apply to access through the serial connection or Telnet.

The following examples show login and logout messages for the User EXEC and Privileged EXEC levels of the CLI.

```
device# show logging
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
Buffer logging: level ACDMEINW, 12 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):
Oct 15 18:01:11:info:dg logout from USER EXEC mode
Oct 15 17:59:22:info:dg logout from PRIVILEGE EXEC mode
Oct 15 17:38:07:info:dg login to PRIVILEGE EXEC mode
Oct 15 17:38:03:info:dg login to USER EXEC mode
```

Syntax: show logging

The first message (the one on the bottom) indicates that user "dg" logged in to the CLI User EXEC level on October 15 at 5:38 PM and 3 seconds (Oct 15 17:38:03). The same user logged into the Privileged EXEC level four seconds later.

The user remained in the Privileged EXEC mode until 5:59 PM and 22 seconds. (The user could have used the CONFIG modes as well. Once you access the Privileged EXEC level, no further authentication is required to access the CONFIG levels.) At 6:01 PM and 11 seconds, the user ended the CLI session.

Removing user-login details from the Syslog messages and traps

If you want to disable the logging of user-login details from the system log, enter the following commands.

```
device(config)# no logging enable user-login
device(config)# write memory
device(config)# end
device# reload
```

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 doing the following.

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.

Network Time Protocol Version 4 (NTPv4)

NTPv4 feature synchronizes the local system clock in the device with the UTC. The synchronization is achieved by maintaining a loop-free timing topology computed as a shortest-path spanning tree rooted on the primary server. NTP does not know about local time zones or daylight-saving time. A time server located anywhere in the world can provide synchronization to a client located anywhere else in the world. It allows clients to use different time zone and daylight-saving properties. Primary servers are synchronized by wire or radio to national standards such as GPS. Timing information is conveyed from primary servers to secondary servers and clients in the network. NTP runs on UDP, which in turn runs on IP.

NTP has a hierarchical structure. NTP uses the concept of a stratum to describe how many NTP hops away a machine is from an authoritative time source. A stratum 1 time server typically has an authoritative time source such as a radio or atomic clock, or a Global Positioning System [GPS] time source directly attached. A stratum 2 time server receives its time through NTP from a stratum 1 time server and so on. 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 use stratum 2 or 3 references.

NTP operates on a client-server basis. The current implementation runs NTP as a secondary server and/or a NTP Client. As a secondary server, the device operates with one or more upstream servers and one or more downstream servers or clients. A client device synchronizes to one or more upstream servers, but does not provide synchronization to dependant clients. Secondary servers at each lower level are assigned stratum numbers one greater than the preceding level. As stratum number increases, the accuracy decreases. Stratum one is assigned to Primary servers.

NTP uses the concept of associations to describe communication between two machines running NTP. NTP associations are statistically configured. On startup or on the arrival of NTP packets, associations are created. Multiple associations are created by the protocol to communicate with multiple servers. NTP maintains a set of statistics for each of the server or the client it is associated with. The statistics

represent measurements of the system clock relative to each server clock separately. NTP then determines the most accurate and reliable candidates to synchronize the system clock. The final clock offset applied for clock adjustment is a statistical average derived from the set of accurate sources.

When multiple sources of time (hardware clock, manual configuration) are available, NTP is always considered to be more authoritative. NTP time overrides the time that is set by any other method.


NTPv4 obsoletes NTPv3 (RFC1305) and SNTP (RFC4330). SNTP is a subset of NTPv4. RFC 5905 describes NTPv4.

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 to maintain time synchronization.

NOTE

Network Time Protocol (NTP) commands must be configured on each individual device.

FIGURE 1 NTP Hierarchy



- NTP implementation conforms to RFC 5905.
- NTP can be enabled in server and client mode simultaneously.
- The NTP uses UDP port 123 for communicating with NTP servers/peers.
- NTP server and client can communicate using IPv4 or IPv6 address
- NTP implementation supports below association modes.
 - Client
 - Server
 - Symmetric active/passive
 - Broadcast server
 - Broadcast client
- NTP supports maximum of 8 servers and 8 peers. The 8 peers includes statically configured and dynamically learned.

- NTP can operate in authenticate or non-authenticate mode. Only symmetric key authentication is supported.
- By default, NTP operates in default VLAN and it can be changed.

Limitations

- FastIron devices cannot operate as primary time server (or stratum 1). It only serves as secondary time server (stratum 2 to 15).
- NTP server and client cannot communicate using hostnames.
- NTP is not supported on VRF enabled interface.
- Autokey public key authentication is not supported.
- The NTP version 4 Extension fields are not supported. The packets containing the extension fields are discarded.
- The NTP packets having control (6) or private (7) packet mode is not supported. NTP packets with control and private modes will be discarded.
- On reboot or switchover, all the NTP state information will be lost and time synchronization will start fresh.
- NTP multicast server/client and manycast functionalities are not supported.
- NTP versions 1 and 2 are not supported.
- NTP MIB is not supported.

Network Time Protocol (NTP) leap second

A leap second is a second added to Coordinated Universal Time (UTC) in order to keep it synchronized with astronomical time (UT1).

There are two main reasons that cause leap seconds to occur. The first is that the atomic second defined by comparing cesium clocks to the Ephemeris Time (ET) scale was incorrect, as the duration of the ephemeris second was slightly shorter than the mean solar second and this characteristic was passed along to the atomic second. The second reason for leap seconds is that the speed of the Earth's rotation is not constant. It sometimes speeds up, and sometimes slows down, but when averaged over long intervals the trend indicates that it is gradually slowing. This gradual decrease in the rotational rate is causing the duration of the mean solar second to gradually increase with respect to the atomic second.

Leap seconds are added in order to keep the difference between UTC and astronomical time (UT1) to less than 0.9 seconds. The International Earth Rotation and Reference Systems Service (IERS), measures Earth's rotation and publishes the difference between UT1 and UTC. Usually leap seconds are added when UTC is ahead of UT1 by 0.4 seconds or more.

How Brocade supports leap second handling for NTP

The obvious question raised is what happens during the NTP leap second itself.

Specifically, a positive leap second is inserted between second 23:59:59 of a chosen UTC calendar date (the last day of a month, usually June 30 or December 31) and second 00:00:00 of the following date. This extra second is displayed on UTC clocks as 23:59:60. On clocks that display local time tied to UTC, the leap second may be inserted at the end of some other hour (or half-hour or quarter-hour), depending on the local time zone. When ever there is a leap second the NTP server notifies by setting the NTP leap second bits.

On Brocade devices when ever there is a negative leap second, the clock is set once second backward of the following date as described here. On positive leap second the clock suppress second 23:59:59 of the last day of a chosen month, so that second 23:59:58 of that date would be followed immediately by second 00:00:00 of the following date.

Because the Earth's rotation speed varies in response to climatic and geological events, UTC leap seconds are irregularly spaced and unpredictable. Insertion of each UTC leap second is usually decided about six months in advance by the International Earth Rotation and

Reference Systems Service (IERS), when needed to ensure that the difference between the UTC and UT1 readings will never exceed 0.9 second.

NTP and SNTP

FastIron 07.3.00c and earlier releases implements SNTP for time synchronization. In FastIron 07.3.00d, NTP can be used for time synchronization in FCX devices with router images. From FastIron 8.0 release onwards, NTP can be used for time synchronization in all FastIron devices with both router and switch images.

NTP and SNTP implementations cannot operate at the same time and one of them has to be disabled.

On downgrading from FastIron 07.3.00d to FastIron 07.3.00c or lower version, the entire NTP configuration is lost.

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 functionality is enabled when you use the `ntp` command, provided SNTP configuration is already removed.

When the NTP server is enabled, it will start listening on the NTP port for client requests and responds with the reference time. Its stratum number will be the upstream time server's stratum + 1. The stratum 1 NTP server is the time server which is directly attached to the authoritative time source.

The device cannot be configured as primary time server with stratum 1. It can be configured as secondary time server with stratum 2 to 15 to serve the time using the local clock.

The NTP server is stateless and will not maintain any NTP client information.

System as an Authoritative NTP Server

The NTP server can operate in master mode to serve time using the local clock, when it has lost synchronization. Serving local clock can be enabled using the `master` command. In this mode, the NTP server stratum number is set to the configured stratum number. When the `master` command is configured and the device was never synchronized with an upstream time server and the clock setting is invalid, the server will respond to client's request with the stratum number set to 16. While the device is operating in the master mode and serving the local clock as the reference time, if synchronization with the upstream server takes place it will calibrate the local clock using the NTP time. The stratum number will switch to that of the synchronized source +1. And when synchronization is lost, the device switches back to local clock time with stratum number as specified manually (or the default).

NOTE

Local time and time zone has to be configured before configuring the `master` command.

- The following scenarios are observed when the `master` command is not configured and the NTP upstream servers are configured:
- If the synchronization with the NTP server/peer is active, the system clock is synchronized and the reference time is the NTP time.
- If the NTP server/peer is configured but not reachable and if the local clock is valid, the server will respond to client's request with the stratum number set to 16.
- If there is no NTP server/peer configured and if the local clock is valid, the server will respond to client's request with the stratum number set to 16.

- If there is no NTP server/peer configured and if the local clock is invalid, the system clock is not synchronized.

The following scenarios are observed when the master command is configured and the NTP upstream servers are also configured:

- If the synchronization with the time server/peer is active, system clock is synchronized and the reference time is the NTP time. If the NTP server/peer is configured but not reachable, the system clock is synchronized. If the local time is valid then the reference time is the local clock time.
- If the NTP server/peer is not configured, the system clock is synchronized. If the local clock is valid, then the reference time is the local clock time.
- If the NTP server/peer is not configured and the local clock is invalid, system clock is not synchronized.

NOTE

Use the master command with caution. It is very easy to override valid time sources using this command, especially if a low stratum number is configured. Configuring multiple machines in the same network with the master command can cause instability in timekeeping if the machines do not agree on the time.

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.

The NTP client can be enabled when we enter the **ntp** command and configure one or more NTP servers/peers.

The NTP client maintains the server and peer state information as association. The server and peer association is mobilized at the startup or whenever user configures. The statically configured server/peer associations are not demobilized unless user removes the configuration. The symmetric passive association is mobilized upon arrival of NTP packet from peer which is not statically configured. The associations will be demobilized on error or time-out.

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 loses a reference source, the time values can flow from the surviving peers to all the others. Each device operates with one or more primary reference sources, such as a radio clock, or a subset of reliable NTP secondary servers. When one of the devices lose all reference sources or simply cease operation, the other peers automatically reconfigures so that time values can flow from the surviving peers to others.

When the NTP server or peer is configured with burst mode, client will send burst of up to 8 NTP packets in each polling interval. The burst number of packets in each interval increases as the polling interval increases from minimum polling interval towards maximum interval.

The NTP peer can operate in:

- Symmetric Active—When the peer is configured using the peer command.
- Symmetric Passive—Dynamically learned upon arrival of a NTP packet from the peer which is not configured. The symmetric passive association is removed on timeout or error.

The following scenarios are observed when the upstream server is not reachable after retries:

- If the NTP server/peer is configured and the master command is not configured, then the system clock is synchronized. When the system clock is synchronized, the server will respond to client's request with the stratum number set to +1. And when the system clock is unsynchronized, the server will respond to client's request with the stratum number set to 16.

- If the NTP server/peer is configured and the master command is configured, then the system clock is synchronized. When the system clock is synchronized, the reference time is the local clock time. If the local clock is valid then the server will respond to client's request with the specified stratum number if it is configured otherwise with the default stratum number.

The following scenarios are observed when you remove the last NTP server/peer under the conditions - the NTP server/peer is configured, master command is not configured, system clock is synchronized and the reference time is the NTP time:

- If the local clock is not valid, the system clock is not synchronized.
- If the local clock is valid, the system clock is synchronized and the reference time is the local clock. The server will respond to the client's request with the specified stratum number if it is configured otherwise with the default stratum number.

NOTE

To create a symmetric active association when a passive association is already formed, disable NTP, configure peer association and then enable NTP again.

NTP broadcast server

An NTP server can also operate in a broadcast mode. 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 packets are sent to the configured subnet when the NTP broadcast server is configured on the interface which is up and the IP address is configured for the broadcast subnet under the following conditions:
 - The local clock is valid and the system clock is synchronized
 - The local clock is valid and the system clock is not synchronized
 - Authentication key is configured, the system clock is synchronized and the local clock is valid
- NTP broadcast packets are not sent in the following cases:
 - NTP broadcast server is configured on the interface which is down even if the system clock is synchronized and the local clock is valid.
 - NTP broadcast server is configured on the interface which is up and no IP address is configured for the broadcast subnet even if the system clock is synchronized and the local clock is valid.
 - NTP broadcast server is configured on the interface which is not present and no IP address is configured for the broadcast subnet even if the system clock is synchronized and the local clock is valid.
 - NTP broadcast server without authentication key is configured on the interface which is up and the IP address is configured for the broadcast subnet even when NTP authentication is enforced and the system clock is synchronized and the local clock is valid.

NTP broadcast client

An NTP broadcast client listens for NTP packets on a broadcast address. When the first packet is received, the client attempts to quantify the delay to the server, 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 interchanges occur, the client has an idea of the network delay and thereafter can estimate the time based only on broadcast packets.

NTP associations

Networking devices running NTP can be configured to operate in variety of association modes when synchronizing time with reference time sources. A networking device can obtain time information on a network in two ways—by polling host servers and by listening to NTP broadcasts. That is, there are two types of associations—poll-based and broadcast-based.

NTP poll-based associations

The following modes are the NTP polling based associations:

1. Server mode
2. Client mode
3. Symmetric Active/Passive

The server mode requires no prior client configuration. The server responds to client mode NTP packets. Use the master command to set the device to operate in server mode when it has lost the synchronization.

When the system is operating in the client mode, it polls all configured NTP servers and peers. The device selects a host from all the polled NTP servers to synchronize with. Because the relationship that is established in this case is a client-host relationship, the host will not capture or use any time information sent by the local client device. This mode is most suited for file-server and workstation clients that are not required to provide any form of time synchronization to other local clients. Use the server and peer to individually specify the time server that you want the networking device to consider synchronizing with and to set your networking device to operate in the client mode.

Symmetric active/passive mode is intended for configurations where group devices operate as mutual backups for each other. Each device operates with one or more primary reference sources, such as a radio clock, or a subset of reliable NTP secondary servers. If one of the devices lose all reference sources or simply cease operation, the other peers automatically reconfigures. This helps the flow of time value from the surviving peers to all the others.

When a networking device is operating in the symmetric active mode, it polls its assigned time-serving hosts for the current time and it responds to polls by its hosts. Because symmetric active mode is a peer-to-peer relationship, the host will also retain time-related information of the local networking device that it is communicating with. When many mutually redundant servers are interconnected via diverse network paths, the symmetric active mode should be used. Most stratum 1 and stratum 2 servers on the Internet adopt the symmetric active form of network setup. The FastIron device operates in symmetric active mode, when the peer information is configured using the peer command and specifying the address of the peer. The peer is also configured in symmetric active mode in this way by specifying the FastIron device information. If the peer is not specifically configured, a symmetric passive association is activated upon arrival of a symmetric active message.

The specific mode that you should set for each of your networking devices depends primarily on the role that you want them to assume as a timekeeping device (server or client) and the device's proximity to a stratum 1 timekeeping server. A networking device engages in polling when it is operating as a client or a host in the client mode or when it is acting as a peer in the symmetric active mode. An exceedingly large number of ongoing and simultaneous polls on a system can seriously impact the performance of a system or slow the performance of a given network. To avoid having an excessive number of ongoing polls on a network, you should limit the number of direct, peer-to-peer or client-to-server associations. Instead, you should consider using NTP broadcasts to propagate time information within a localized network.

NTP broadcast-based associations

The broadcast-based NTP associations should be used in configurations involving potentially large client population. Broadcast-based NTP associations are also recommended for use on networks that have limited bandwidth, system memory, or CPU resources.

The devices operating in the broadcast server mode broadcasts the NTP packets periodically which can be picked up by the devices operating in broadcast client mode. The broadcast server is configured using the **broadcast** command.

A networking device operating in the broadcast client mode does not engage in any polling. Instead, the device receives the NTP broadcast server packets from the NTP broadcast servers in the same subnet. The NTP broadcast client forms a temporary client association with the NTP broadcast server. A broadcast client is configured using the **broadcast client** command. For broadcast client mode to work, the broadcast server and the clients must be located on the same subnet.

Synchronizing time

After the system peer is chosen, the system time is synchronized based on the time difference with system peer:

- If the time difference with the system peer is 128 msec and < 1000 sec, the system clock is stepped to the system peer reference time and the NTP state information is cleared.

Authentication

The time kept on a machine is a critical resource, so it is highly recommended to use the encrypted authentication mechanism.

The NTP can be configured to provide cryptographic authentication of messages with the clients/peers, and with its upstream time server. Symmetric key scheme is supported for authentication. The scheme uses MD5 keyed hash algorithm.

The authentication can be enabled using the **authenticate** command. The set of symmetric key and key string is specified using the **authentication-key** command.

If authentication is enabled, NTP packets not having a valid MAC address are dropped.

If the NTP server/peer is configured without authentication keys, the NTP request is not sent to the configured server/peer.

NOTE

The same set or subset of key id and key string should be installed on all NTP devices.

VLAN and NTP

When VLAN is configured,

- NTP time servers should be reachable through the interfaces which belong to the configured VLAN. Otherwise, NTP packets are not transmitted. This is applicable to both the unicast and the broadcast server/client.
- NTP broadcast packets are sent only on the interface which belongs to the configured VLAN.
- The received unicast or broadcast NTP packet are dropped if the interface on which packet has been received does not belong to the configured VLAN

Configuring NTP

NTP services are disabled on all interfaces by default.

Prerequisites:

- 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).
- Disable SNTP by removing all the SNTP configurations.

Enabling NTP

NTP and SNTP implementations cannot operate simultaneously. By default, SNTP is enabled. To disable SNTP and enable NTP, use the **ntp** command in configuration mode. This command enables the NTP client and server mode if SNTP is disabled.

```
Brocade(config)# ntp
Brocade(config-ntp)#
```

Syntax: [no] ntp

Use the **no** form of the command to disable NTP and remove the NTP configuration.

NOTE

The **no ntp** command removes all the configuration which are configured statistically and learned associations from NTP neighbors.

NOTE

You cannot configure the **ntp** command if SNTP is enabled. If SNTP is enabled, configuring the **ntp** command will display the following message: "SNTP is enabled. Disable SNTP before using NTP for time synchronization"

Disabling NTP

To disable the NTP server and client mode, use the **disable** command in NTP configuration mode. Disabling the NTP server or client mode will not remove the configurations.

```
Brocade(config-ntp)# disable
```

Syntax: [no] disable [serve]

If the **serve** keyword is specified, then NTP will not serve the time to downstream devices. The **serve** keyword disables the NTP server mode functionalities. If the **serve** keyword is not specified, then both NTP client mode and NTP server mode functionalities are disabled.

Use the **no** form of the command to enable NTP client and server mode. To enable the client mode, use the **no disable** command. To enable the client and server mode, use the **no disable serve** command. The **no disable** command enables both client and server, if the client is already enabled and server is disabled at that time "no disable server " enables the server.

NOTE

The **disable** command disables the NTP server and client mode; it does not remove the NTP configuration.

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.

By default, authentication is disabled.

```
Brocade(config-ntp)# [no] 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.

By default, authentication keys are not configured.

```
Brocade(config-ntp)# authentication-key key-id 1 md5 moof
```

Syntax: `[no] authentication-key key-id [md5 | sha1] 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 **sha1** keyword specifies that the SHA1 keyed hash algorithm is used for NTP authentication.

NOTE

If JITC is enabled, only the **sha1** option is available.

The key-string option is the value of the MD5 or SHA1 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

When the system sends an NTP packet, the source IP address is normally set to the address of the interface through which the NTP packet is sent. Use the **source-interface** command to configure a specific interface from which the IP source address will be taken. To remove the specified source address, use the no form of this command.

This interface will be used for the source address for all packets sent to all destinations. If a source address is to be used for a specific association, use the source keyword in the peer or server 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. Source IP address of a tunnel interface is not supported.

```
Brocade(config-ntp)# source-interface ethernet 1/3/1
```

Syntax: `[no] source-interface ethernet { port | loopback num | ve num }`

Specify the *port* parameter in the format stack-unit/slotnum/portnum.

The *loopback num* parameter specifies the loopback interface number.

The *ve num* 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 remove the NTP server configuration, use the **no** form of this command.

By default, no servers are configured.

```
Brocade(config-ntp)#server 1.2.3.4 key 1234
```

Syntax: **[no] server** { *ipv4-address* | *ipv6-address* } [**version num**] [**key key-id**] [**minpoll interval**] [**maxpoll interval**] [**burst**]

The *ipv4-address* or *ipv6-address* parameter is the IP address of the server providing the clock synchronization.

The *version num* option defines the Network Time Protocol (NTP) version number. Valid values are 3 or 4. If the **num** option is not specified, the default is 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, and so on).

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, and so on).

The *burst* option sends a burst of packets to the server at each polling interval.

Configuring the master

To configure the FastIron device 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**]

The number variable is a number from 2 to 15. It indicates the NTP stratum number that the system will claim.

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.

This **peer** command is not effective if the NTP is enabled in client-only mode.

NOTE

If the peer is a member of symmetric passive association, then configuring the **peer** command will fail.

```
Brocade(config-ntp)# peer 1.2.3.4 key 1234
```

Syntax: **[no] peer** { *ipv4-address* | *ipv6-address* } [**version num**] [**key key-id**] [**minpoll interval**] [**maxpoll interval**] [**burst**]

The *ipv4-address* or *ipv6-address* parameter is the IP address of the peer providing the clock synchronization.

The *version num* option defines the Network Time Protocol (NTP) version number. Valid values are 3 and 4. If this option is not specified, then the default is 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, and so on).

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, and so on).

The *burst* option sends a burst of packets to the peer at each polling interval.

NOTE

When the NTP server/peer is configured, the **master** command is not configured; on configuring the **clock set** command the system clock is not synchronized. When the **master** command is configured, on configuring the **clock set** command the system clock is synchronized and the reference time will be the local clock.

To have active peers at both the ends, you need to disable NTP, configure the peers and enable the NTP using the **no disable** command.

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 1/2/13
Brocade(config-ntp-if-e1000-1/2/13)# exit
Brocade(config-ntp)# ntp-interface management 1
Brocade(config-ntp-mgmt-1)# exit
Brocade(config-ntp)# ntp-interface ve 100
Brocade(config-ntp-ve-100)#
```

Syntax: [no] ntp-interface { management 1 | ethernet port | ve id }

The *management 1* parameter is the management port 1.

The *ethernet port* parameter specifies the ethernet port number. Specify the port parameter in the format stack-unit/slotnum/portnum.

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 the 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. The 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 the NTP broadcast server packets are not sent. To disable this capability, use the no form of this command.

By default, the broadcast mode is not enabled.

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 *num*]

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 num* option defines the Network Time Protocol (NTP) version number. If this option is not specified, then the default value is 4.

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
```

The following table provides descriptions of the **show ntp status** command output.

TABLE 4 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.

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
~2001:235::234
INIT 16 - 64 0 0.00 0.000 15937
* synced, # selected, + candidate, - outlier, x falseticker, ~ configured
```

The following table provides descriptions of the **show ntp associations** command output.

TABLE 5 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.
address	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
2001:1:99:30::1 configured server, sys peer, stratum 3
ref ID 204.235.61.9, 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** command with the appropriate parameters to display the NTP servers and peers association information for a specific IP address.

```
Brocade# show ntp association detail 1.99.40.1
1.99.40.1 configured server, candidate, stratum 3
ref ID 216.45.57.38, 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 following table provides descriptions of the **show ntp associations detail** command output.

TABLE 6 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.
reach peer	The peer reachability (bit string in octal).
Delay	Round-trip delay to peer.
offset	Offset of a peer clock relative to this clock.
Dispersion	Dispersion of a peer clock.
precision	Precision of a peer clock.
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.

TABLE 6 NTP associations detail command output descriptions (continued)

Field	Description
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 list configuration examples to configure the Brocade device.

NTP server and client mode configuration

Sample CLI commands to configure the Brocade device in NTP server and client modes.

```
Brocade(config-ntp)# server 10.1.2.3 minpoll 5 maxpoll 10
Brocade(config-ntp)# server 11::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 11::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-id 1 md5 key123
Brocade(config-ntp)# server 10.1.2.4 key 1
```

NTP loose authentication configuration

Sample CLI commands to configure the Brocade device in loose authentication mode. This allows some of the servers or clients to use the authentication keys.

```
Brocade(config-ntp)# authentication-key key-id 1 md5 key123
Brocade(config-ntp)# server 10.1.2.4 key 1
Brocade(config-ntp)# server 10.1.2.7
```

NTP interface context for the broadcast server or client mode

Sample CLI commands to 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 1/1/3
```



```
Brocade(config-ntp-if-e1000-1/1/3)# 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 1/1/5
Brocade(config-ntp-if-e1000-1/1/5)# broadcast client
Brocade(config-ntp)# ntp-interface ve 100
Brocade(config-ntp-ve-100)# broadcast client
```

Basic port parameter configuration

The procedures in this section describe how to configure the port parameters shown in [Basic Software Features](#) on page 27.

All Brocade ports are pre-configured with default values that allow the device to be fully operational at initial startup without any additional configuration. However, in some cases, changes to the port parameters may be necessary to adjust to attached devices or other network requirements.

Specifying a port address

You can specify a port address for an uplink (data) port, stacking port, or a management port.

ICX 6430 and ICX 6450

Specifying a data port

The port address format is stack unit/slot/port, where:

- *stack unit* --Specifies the stack unit ID. For the ICX 6430, range is from 1 to 4. For the ICX 6450, range is from 1 to 8. If the device is not part of a stack, the stack unit ID is 1.
- *slot* --Specifies the slot number. Can be 1 or 2.
- *port* --Specifies the port number in the slot. Range is from 1 to 24 (24-port models) or 1 to 48 (48-port models).

This example shows how to specify port 2 in slot 1 of a device that is not part of a stack:

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

Specifying a stacking port

The port address format is stack unit/slot/port, where:

- *stack unit* --Specifies the stack unit ID. For the ICX 6430, range is from 1 to 4. For the ICX 6450, range is from 1 to 8.
- *slot* --Specifies the slot number. Stacking ports are in slot 2.
- *port* --Specifies the port number in the slot. Stacking ports are 1, 2, 3, and 4.

This example shows how to specify stacking port 3 in slot 2 of unit 3 in a stack:

```
Brocade (config) # interface ethernet 3/2/3
```

Specifying a management port

The management port number is always 1. This example shows how to specify the management port:

```
Brocade (config) # interface management 1
```

ICX 6610

Specifying a data port

The port address format is stack unit/slot/port, where:

- *stack unit* --Specifies the stack unit ID. Range is from 1 to 8. If the device is not part of a stack, the stack unit ID is 1.
- *slot* --Specifies the slot number. Can be 1 or 3.
- *port* --Specifies the port number in the slot. Range is from 1 to 24 (24-port models) or 1 to 48 (48-port models).

This example shows how to specify port 2 in slot 1 of a device that is not part of a stack:

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

Specifying a stacking port

The port address format is stack unit/slot/port, where:

- *stack unit* --Specifies the stack unit ID. Range is from 1 to 8.
- *slot* --Specifies the slot number. Stacking ports are in slot 2.
- *port* --Specifies the port number in the slot. Dedicated stacking ports are 1, 2, 6, and 7.

This example shows how to specify stacking port 2 in slot 2 of unit 3 in a stack:

```
Brocade (config) # interface ethernet 3/2/2
```

Specifying a management port

The management port number is always 1. This example shows how to specify the management port:

```
Brocade (config) # interface management 1
```

FCX

Specifying a data port

The port address format is stack unit/slot/port, where:

- *stack unit* --Specifies the stack unit ID. Range is from 1 to 8. If the device is not part of a stack, the stack unit ID is 1.
- *slot* --Specifies the slot number. Can be 1 or 3.
- *port* --Specifies the port number in the slot. Range is from 1 to 24 (24-port models) or 1 to 48 (48-port models).

This example shows how to specify port 2 in slot 1 of a device that is not part of a stack:

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

Specifying a stacking port

The port address format is stack unit/slot/port, where:

- `stack unit` --Specifies the stack unit ID. Range is from 1 to 8.
- `slot` --Specifies the slot number. Default stacking ports are in slot 2 (FCX S/S-F) and slot3 (FCX E/I).
- `port` --Specifies the port number in the slot. Default stacking ports in slot 2 and slot 3 are ports 1 and 2.

This example shows how to specify port 2 in slot 2 of unit 3 in a stack:

```
Brocade (config) # interface ethernet 3/2/2
```

Specifying a management port

The management port number is always 1. This example shows how to specify the management port:

```
Brocade (config) # interface management 1
```

FSX

Specifying a data port

The port address format is slot/port, where:

- `slot` --Specifies the interface slot number. Range is from 1 to 8 (FSX 800) or 1 to 16 (FSX 1600).
- `port` --Specifies the port number in the slot. Range is from 1 to 48 depending on the interface module.

This example shows how to specify port 2 in slot 1:

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

Specifying a management port

The management port number is always 1. This example shows how to specify the management port:

```
Brocade (config) # interface management 1
```

NOTE

Stacking is not supported on FSX devices.

Assigning port names

You can assign text strings as port names, which help you identify ports with meaningful names. You can assign port names to individual ports or to a group of ports. You can assign a port name to physical ports, virtual interfaces, and loopback interfaces.

Assigning a port name

To assign a name to a port, enter commands such as the following:

```
device(config)# interface ethernet 2
device(config-if-e1000-2) # port-name Marsha
```

Syntax: `port-name text`

The *text* parameter is an alphanumeric string. The name can be up to 255 characters long. The name can contain blanks. You do not need to use quotation marks around the string, even when it contains blanks. The port name can contain special characters as well, but the percentage character (%), if it appears at the end of the port name, is dropped.

Assigning the same name to multiple ports

To assign a name to a range of ports, enter commands such as the following:

```
Brocade (config)# interface ethernet 1/1/1 to 1/1/10
Brocade (config-mif-1/1/1-1/1/10)# port-name connected-to-the nearest device
```

Syntax: [no] port-name *text*

To remove the assigned port name, use **no** form of the command.

The *text* parameter is an alphanumeric string, up to 255 characters long. The name can contain blanks. You do not need to use quotation marks around the string, even when it contains blanks.

You can also specify the individual ports, separated by space.

To assign a name to multiple specific ports, enter commands such as the following:

```
Brocade (config)# interface ethernet 1/1/1 ethernet 1/1/5 ethernet 1/1/7
Brocade (config-mif-1/1/1, 1/1/5, 1/1/7)# port-name connected-to-the nearest device
```

Displaying the port name for an interface

You can use the **show interface brief** command to display the name assigned to the port. If any of the ports have long port names, they are truncated. To show full port names, use the **show interfaces brief wide** command.

```
Brocade# show interfaces brief
Port   Link   State   Dupl Speed Trunk Tag Pvid Pri   Name
MAC
1/1/23 Up     Forward Full 1G   None No  1   0   748e.f82d.7a16 connected-
1/1/47 Up     Forward Full 1G   None No  1   0   748e.f82d.7a2e
mgmt1  Up     None    Full 1G   None No  None 0   748e.f82d.7a00
```

In this output, the port name for interface 1/1/23 is truncated.

Use the **show interface brief wide** command to avoid truncating long port names.

To display the complete port name for an interface, enter the following command.

```
Brocade# show interface brief wide
Port   Link   State   Dupl Speed Trunk Tag Pvid Pri   Name
MAC
1/1/23 Up     Forward Full 1G   None No  1   0   748e.f82d.7a16 connected-
to-the nearest device
1/1/47 Up     Forward Full 1G   None No  1   0   748e.f82d.7a2e
mgmt1  Up     None    Full 1G   None No  None 0   748e.f82d.7a00
```

Syntax: show interface brief [wide] [ethernet *stack-unit/slot/port* | loopback *port* | management *port* | slot *port* | tunnel *port* | ve *port*]

The *ethernet stack-unit/slot/port* parameter specifies the Ethernet port for which you want to display the interface information.

The *loopback* option specifies the loopback port for which you want to display the interface information.

The *management* option specifies the management port for which you want to display the interface information.

The *slot* option specifies all the ports in a slot for which you want to display the interface information.

The *tunnel* option specifies the tunnel port for which you want to display the interface information.

The *ve* option specifies the virtual routing (VE) port for which you want to display the interface information.

The following table describes the output parameters of the **show interface brief wide** command.

TABLE 7 Output parameters of the show interface brief wide command

Field	Description
Port	Specifies the port number.
Link	Specifies the link state.
Port-State	Specifies the current port state.
Speed	Specifies the link speed.
Tag	Specifies if the port is tagged or not.
Pvid	Specifies the port VLAN ID.
Pri	Specifies the priority.
MAC	Specifies the MAC address.
Name	Specifies the port name.

To display the complete port name for an Ethernet interface, enter a command such as the following.

```
Brocade# show interface brief wide ethernet 1/1/23
PPort   Link   State Dupl Speed Trunk Tag Pvid Pri MAC           Name
1/1/23          Up    Forward Full 1G   None No  1   0   748e.f82d.7a16 connected- to-FCX
```

Syntax: `show interface brief wide ethernet stack-unit/slot/port`

For more information about field descriptions of the command output, refer [Displaying the port name for an interface](#).

Port speed and duplex mode modification

The Gigabit Ethernet copper ports are designed to auto-sense and auto-negotiate the speed and duplex mode of the connected device. If the attached device does not support this operation, you can manually enter the port speed to operate at either 10, 100, or 1000 Mbps. This configuration is referred to as force mode. The default and recommended setting is 10/100/1000 auto-sense. Port duplex mode and port speed are modified by the same command

NOTE

You can modify the port speed of copper ports only; this feature does not apply to fiber ports.

NOTE

For optimal link operation, copper ports on devices that do not support 803.3u must be configured with like parameters, such as speed (10,100,1000), duplex (half, full), MDI/MDIX, and Flow Control.

Port speed and duplex mode configuration syntax

The following commands change the port speed of copper interface 8 on a FastIron device from the default of 10/100/1000 auto-sense, to 100 Mbps operating in full-duplex mode.

```
device(config)# interface ethernet 1/1/8
device(config-if-e1000-1/1/8)# speed-duplex 100-full
```

Syntax: `speed-duplex value`

The value variable can be one of the following values:

- **10-full** - 10 Mbps, full duplex
- **10-half** - 10 Mbps, half duplex
- **100-full** - 100 Mbps, full duplex

- **100-half** - 100 Mbps, half duplex
- **1000-full-master** - 1 Gbps, full duplex, master
- **1000-full-slave** - 1 Gbps, full duplex, slave
- **10g-full** - 10 Gbps, full duplex
- **auto** - auto-negotiation

The default is **auto** (auto-negotiation).

Use the **no** form of the command to restore the default.

NOTE

1G speed is not supported on 10G ports with 10G SFP+ optics.

NOTE

On Brocade ICX 7450-32ZP, the command options **10-half**, **10-full**, and **100-half** are not supported for 2.5G ports.

NOTE

On Brocade ICX 7450 and Brocade ICX 7250-24G, the command options **10-half** and **100-half** are not supported on 1 Gbps fiber ports with mini-GBIC (SFPs) for copper.

NOTE

On FastIron devices, when setting the speed and duplex-mode of an interface to **1000-full**, configure one side of the link as master (**1000-full-master**) and the other side as slave (**1000-full-slave**).

NOTE

On Brocade ICX 6610 and ICX 6650 devices, after you remove the 10 Gbps speed from the running configuration, plugging in a 1Gbps optic SFP transceiver into a 10 Gbps port causes the software to fail to revert the ports back from the default 10Gbps mode to the 1 Gbps speed. Remove the 1Gbps SFP transceiver and plug in the 10Gbps optic SFP+transceiver so that the devices go into the default 10 Gbps mode.

NOTE

When you use fixed speed and duplex configuration, you should use the non-auto MDI-MDIX configuration.

NOTE

On ICX 7450, **speed-duplex 1000-full** must be configured on both of the SFP sides for the front 4x10G module to link-up the port as 1G speed.

Configuration considerations for port speed and duplex mode

The following considerations apply to the port speed and duplex mode configuration:

- On Brocade ICX 7250, uplink ports using E1MG-TX (1G GBIC Copper SFP) transceivers, do not support 10 Mbps speed.
- When a local partner issues a **speed-duplex 100-full** or **speed-duplex 10-full** command, if the remote partner does not issue the same commands, it becomes **100-half** or **10-half**, and may receive collision errors. The local partner may receive In Errors such as CRC, fragments, or bad packets.
- When a local partner issues a **speed-duplex 100-full** or **speed-duplex 10-full** command, if the remote partner issues the same command, the port may or may not come up because both sides enter the force mode and want to force the partner to accept these conditions. If both sides come up, they may not receive any In or Out Errors.
- When both local and remote partners have a force mode configuration such as **100-full/half** or **10-full/half**, for example, ICX6610-24F 1/1/1 (local link **100-full**)<->(100-full remote link) FCX 1/1/1, if another force mode such as **10-full** is entered in a local or remote partner, the remote or local partner link may or may not come up. This is an IEEE force mode

standard. To resolve the force mode changing, it is recommended that you first change to auto mode on one side, before switching to another force mode configuration.

- Brocade ICX 7750, Brocade ICX 7450 and Brocade ICX 7250 devices do not support half duplex configuration as the MAC addresses of these platforms do not support half duplex. For example, 10-half or 100-half configuration on fixed copper port is an invalid configuration.
- Same speed clock setting on both sides of the link on fixed copper port is an invalid configuration and is not supported. For example, speed set to 1000-full-master on both sides of the link is an invalid configuration. Instead, configure 1000-full-master at one end and 1000-full-slave at the other.

The following tables provide a list of invalid configurations on copper ports.

TABLE 8 List of invalid configurations

ICX 7450 / ICX 7250 configuration	Link Partner - ICX 7450 / ICX 7250 / ICX 7750 configuration
Auto-negotiation (10-full)	10-half
Auto-negotiation (100-full)	100-half
10-half	10-half
100-half	100-half
1000-full-slave	1000-full-slave
1000-full-master	1000-full-master
10g-full-master	10g-full-master
10g-full-slave	10g-full-slave

TABLE 9 List of invalid configurations

ICX 7450 / ICX 7250 configuration	Link Partner - ICX 6430 / ICX 6450 / ICX 6610 / ICX 6650 / FCX / FSX 800 / FSX 1600 configuration
Auto-negotiation (10-full)	10-half
Auto-negotiation (100-full)	100-half
10-half	AN (10-half)
10-full	AN (10-half)
100-half	AN (100-half)
10-half	10-half
100-half	100-half
1000-full-slave	1000-full-slave
1000-full-master	1000-full-master
1000-full	1000-full
10g-full-master	10g-full-master
10g-full-slave	10g-full-slave
10g-full	10g-full

Enabling auto-negotiation maximum port speed advertisement and down-shift

NOTE

For optimal link operation, link ports on devices that do not support 802.3u must be configured with like parameters, such as speed (10,100,1000), duplex (half, full), MDI/MDIX, and Flow Control.

Maximum Port speed advertisement is an enhancement to the auto-negotiation feature, a mechanism for accommodating multi-speed network devices by automatically configuring the highest performance mode of inter-operation between two connected devices.

Port speed down-shift enables Gbps copper ports on the Brocade device to establish a link at 1000. Mbps over a 4-pair wire when possible, or to down-shift to 100 Mbps if the medium is a 2-pair wire.

Maximum port speed advertisement enables you to configure an auto-negotiation maximum speed that Gbps copper ports on the Brocade device will advertise to the connected device. You can configure a port to advertise a maximum speed of either 100 Mbps or 10 Mbps. When the maximum port speed advertisement feature is configured on a port that is operating at 100 Mbps maximum speed, the port will advertise 10/100 Mbps capability to the connected device. Similarly, if a port is configured at 10 Mbps maximum speed, the port will advertise 10 Mbps capability to the connected device.

The maximum port speed and down-shift advertisement features operate dynamically at the physical link layer between two connected network devices. They examine the cabling conditions and the physical capabilities of the remote link, then configure the speed of the link segment according to the highest physical-layer technology that both devices can accommodate.

The maximum port speed and down-shift advertisement features operate independently of logical trunk group configurations. Although Brocade recommends that you use the same cable types and auto-negotiation configuration on all members of a trunk group, you could utilize the auto-negotiation features conducive to your cabling environment. For example, in certain circumstances, you could configure each port in a trunk group to have its own auto-negotiation maximum port speed advertisement or port speed down-shift configuration.

Maximum port speed advertisement and down-shift application notes

- The maximum port speed advertisement works only when auto-negotiation is enabled (CLI command **speed-duplex auto**). If auto-negotiation is OFF, the device will reject the maximum port speed advertisement configuration.
- When the maximum port speed advertisement is enabled on a port, the device will reject any configuration attempts to set the port to a forced speed mode (100 Mbps or 1000 Mbps).
- When port speed down-shift or maximum port speed advertisement is enabled on a port, the device will reject any configuration attempts to set the port to a forced speed mode (100 Mbps or 1000 Mbps).

Configuring maximum port speed advertisement

NOTE

This feature is not supported on Brocade ICX 7750.

To configure a maximum port speed advertisement of 10 Mbps on a port that has auto-negotiation enabled, enter a command such as the following at the Global CONFIG level of the CLI.

```
device(config)
# link-config gig copper autoneg-control 10m ethernet 1
```

To configure a maximum port speed advertisement of 100 Mbps on a port that has auto-negotiation enabled, enter the following command at the Global CONFIG level of the CLI.

```
device(config)
# link-config gig copper autoneg-control 100m ethernet 2
```

Syntax: [no] link-config gig copper autoneg-control { 100m-auto | 10m-auto } ethernet stack-id/slot/port [to stack-id/slot/port | [ethernet stack-id/slot/port to stack-id/slot/port | ethernet stack-id/slot/port] ...]

You can enable maximum port speed advertisement on one or two ports at a time.

To disable maximum port speed advertisement after it has been enabled, enter the **no** form of the command.

Configuring port speed down-shift and auto-negotiation for a range of ports

Port speed down-shift and auto-negotiation can be configured for an entire range of ports with a single command.

For example, to configure down-shift on ports 1/1/1 to 1/1/10 and 1/1/15 to 1/1/20 on the device, enter the following.

```
Brocade(config)# link-config gig copper autoneg-control down-shift ethernet 1/1/1
to 1/1/10 ethernet 1/1/15 to 1/1/20
```

To configure down-shift on ports 5 to 13 and 17 to 19 on a compact switch, enter the following.

```
Brocade(config)# link-config gig copper autoneg-control down-shift ethernet 5 to 13 ethernet 17 to 19
```

Syntax: [no] link-config gig copper autoneg-control [down-shift | 100m-auto | 10m-auto] ethernet *port-list*

NOTE

The <port-list> variable represents the list of ports to which the command will be applied.

For <port-list>, specify the ports in the following format:

- <unit-id/slotnum/portnum>

You can list all of the ports individually, use the keyword **to** to specify ranges of ports, or a combination of both. To apply the configuration to all ports on the device, use the keyword **all** instead of listing the ports individually.

The output from the **show run** command for this configuration will resemble the following.

```
Brocade# show run

Current configuration:
!
ver 08.0.40q017T213
!
module 1 icx7450-24-port-management-module
module 2 icx7400-xgf-4port-40g-module
!
link-config gig copper autoneg-control down-shift ethernet 1/1/1 to 1/1/10
ethernet 1/1/15 to 1/1/20
!
!
ip address 10.44.9.11 255.255.255.0
ip default-gateway 10.53.5.1
!
end
```

To disable selective auto-negotiation of 100m-auto on ports 1/1/21 to 1/1/25 and 1/1/30, enter the following.

```
Brocade(config)# no link-config gig copper autoneg-control 100m-auto ethernet
1/1/21 to 1/1/25 ethernet 1/1/30
```

Enabling port speed down-shift

Enable port speed down-shift on a port that has auto-negotiation enabled.

Enter the context of your task here (optional).

1. At the Global CONFIG level of the CLI, enter the following:

```
Brocade(config)# link-config gig copper autoneg-control down-shift ethernet 1 ethernet 2
```

The above command configures Gbps copper ports 1 and 2 to establish a link at 1000 Mbps over a 4-pair wire when possible, or to down-shift (reduce the speed) to 100 Mbps when the medium is a 2-pair wire.

Syntax: `[no] link-config gig copper autoneg-control down-shift ethernet port [ethernet port] to port`

2. Specify the port variable in one of the following formats:
 - FWS and FCX stackable switches – `<stack-unit/slotnum/portnum>`
 - FSX 800 and FSX 1600 chassis devices – `<slotnum/portnum>`
 - FESX compact switches – `<portnum>`

NOTE

To list all of the ports individually, use the keyword in order to specify ranges of ports, or a combination of both. You can enable port speed down-shift on one or two ports at a time.

3. To disable port speed down-shift, enter the no form of the command.

Force mode configuration

You can manually configure a 10/100 Mbps port to accept either full-duplex (bi-directional) or half-duplex (uni-directional) traffic.

NOTE

You can modify the port duplex mode of copper ports only. This feature does not apply to fiber ports.

Port duplex mode and port speed are modified by the same command.

Force mode configuration syntax

To change the port speed of interface 8 from the default of 10/100/1000 auto-sense to 10 Mbps operating at full-duplex, enter the following.

```
device(config)
# interface ethernet 8
device(config-if-e1000-8)# speed-duplex 10-full
```

Syntax: `speed-duplex value`

The value can be one of the following:

- 10-full
- 10-half
- 100-full
- 100-half
- auto (default)

NOTE

On Brocade ICX 7450 and Brocade ICX 7250-24G, the command options **10-half** and **100-half** are not supported on 1G fiber ports with mini-GBIC (SFPs) for copper.

Force Mode Configuration Considerations

The following considerations apply to the force mode configuration.

- When a local partner issues a **speed-dup 100-full** or **speed-dup 10-full** command, if the remote partner does not issue the same commands it becomes 100-half or 10-half, and may receive collision errors. The local partner may receive InErrors such as CRC, Fragment or Bad packets.
- When a local partner issues a **speed-dup 100-full** or **speed-dup 10-full** command, if the remote partner issues the same command, the port may or may not come up, since both sides enter the force mode and want to force the partner to accept these conditions. If both sides come up, they may not receive any In or Out Errors.
- When a local partner is a force mode configuration such as 100-full/half or 10-full-half and the remote partner is also a force mode configuration, for example, ICX6610-24F 1/1/1 (local link 100-full)<->(100-full remote link) FCX 1/1/1, if another force mode in a local or remote partner such as 10-full is entered, the remote or local partner link may or may not come up. This is an IEEE force mode standard. To resolve force mode changing, it is recommended that you change to auto mode first on one side before switching to another force mode configuration.

MDI and MDIX configuration

Brocade devices support automatic Media Dependent Interface (MDI) and Media Dependent Interface Crossover (MDIX) detection on all Gbps Ethernet Copper ports.

MDI/MDIX is a type of Ethernet port connection using twisted pair cabling. The standard wiring for end stations is MDI, whereas the standard wiring for hubs and switches is MDIX. MDI ports connect to MDIX ports using straight-through twisted pair cabling. For example, an end station connected to a hub or a switch uses a straight-through cable. MDI-to-MDI and MDIX-to-MDIX connections use crossover twisted pair cabling. So, two end stations connected to each other, or two hubs or switches connected to each other, use crossover cable.

The auto MDI/MDIX detection feature can automatically correct errors in cable selection, making the distinction between a straight-through cable and a crossover cable insignificant.

MDI and MDIX configuration notes

- This feature applies to copper ports only.
- The **mdi-mdix mdi** and **mdi-mdix mdix** commands work independently of auto-negotiation. Thus, these commands work whether auto-negotiation is turned ON or OFF.

MDI and MDIX configuration syntax

The auto MDI/MDIX detection feature is enabled on all Gbps copper ports by default. For each port, you can disable auto MDI/MDIX, designate the port as an MDI port, or designate the port as an MDIX port.

To turn off automatic MDI/MDIX detection and define a port as an MDI only port.

```
device(config-if-e1000-2)# mdi-mdix mdi
```

To turn off automatic MDI/MDIX detection and define a port as an MDIX only port.

```
device(config-if-e1000-2)# mdi-mdix mdix
```

To turn on automatic MDI/MDIX detection on a port that was previously set as an MDI or MDIX port.

```
device(config-if-e1000-2)# mdi-mdix auto
```

Syntax: `mdi-mdix[mdi | mdix | auto]`

After you enter the **mdi-mdix** command, the Brocade device resets the port and applies the change.

To display the MDI/MDIX settings, including the configured value and the actual resolved setting (for `mdi-mdix auto`), enter the command **show interface** at any level of the CLI.

Disabling or re-enabling a port

A port can be made inactive (disable) or active (enable) by selecting the appropriate status option. The default value for a port is enabled.

To disable port 8 of a Brocade device, enter the following.

```
device(config)
# interface ethernet 8
device(config-if-e1000-8)# disable
```

You also can disable or re-enable a virtual interface. To do so, enter commands such as the following.

```
device(config)
# interface ve v1
device(config-vif-1)# disable
```

Syntax: `disable`

To re-enable a virtual interface, enter the **enable** command at the Interface configuration level. For example, to re-enable virtual interface `v1`, enter the **enable** command.

```
device(config-vif-1)# enable
```

Syntax: `enable`

Disabling laser light emission on port

This feature enables you to switch off the laser light emission, when a port is disabled.

This overcomes the situation of laser light continuing to emit even when the port is disabled. You can disable the laser light emission using the **port-down-disable-laser** command. The **no** form of the command re-enables the default behavior.

```
device(config-if-e1000-1/1/1)# port-down-disable-laser
device(config-if-e1000-1/1/1)# no port-down-disable-laser
```

The command is present in the running configuration and is applicable per port at the interface level. The command persist across reloads. Laser is switched on or off based on the CLI configuration and the port status (enable or disable). If the CLI configuration is present, the laser is switched off automatically when the port is disabled. You can apply the CLI configuration irrespective of the port state (enable or disable).

NOTE

Before applying the CLI command, make sure that the optics are present on the interface.

The command supports stacking environments and range configuration.

Feature support

This feature is supported on ports of ICX 7450 and ICX 7250 only. The CLI is supported only on 1G SFP and 10G SFP+ fiber ports. It is not supported on copper ports and CGBIC. Also, this is not supported on ICX6610, ICX6450, ICX6430, ICX7750, SXL, ICX 6650 and on 40G ports.

Turning off the port laser light

Execute the following steps to turn off the port laser light.

1. Issue the **port-down-disable-laser** command on the interface.

```
device(config-if-e1000-1/1/1)# port-down-disable-laser
```

2. Issue the **disable** CLI command on that interface to disable the port.

```
device(config-if-e1000-1/1/1)# disable
```

Limitations

If the **port-down-disable-laser** command is configured on the port, the laser emission is switched off when the port is disabled. However, if the command is removed from the disabled port, turning on the laser light must be taken care manually when the port is enabled. The command does not support global configuration.

Flow control configuration

Flow control (802.3x) is a QoS mechanism created to manage the flow of data between two full-duplex Ethernet devices. Specifically, a device that is oversubscribed (is receiving more traffic than it can handle) sends an 802.3x PAUSE frame to its link partner to temporarily reduce the amount of data the link partner is transmitting. Without flow control, buffers would overflow, packets would be dropped, and data retransmission would be required.

All FastIron devices support *asymmetric* flow control, meaning they can receive PAUSE frames but cannot transmit them. In addition, FCX and ICX devices also support *symmetric* flow control, meaning they can both receive and transmit 802.3x PAUSE frames. For details about symmetric flow control, refer to [Symmetric flow control on FCX and ICX devices](#) on page 64.

Flow control configuration notes

- Auto-negotiation of flow control is not supported on 10 Gbps and 40 Gbps ports, fiber ports, and copper or fiber combination ports.
- When any of the flow control commands are applied to a port that is up, the port will be disabled and re-enabled.
- For 10 Gbps and 40 Gbps ports, the **show interface** command with the appropriate parameters shows whether Flow Control is enabled or disabled, depending on the configuration.
- When flow-control is enabled, the hardware can only advertise PAUSE frames. It does not advertise Asym.
- On ICX 7750 devices the default packet-forwarding method is cut-through, in which port flow control (IEEE 802.3x) is not supported but priority-based flow control (PFC) is supported. You can configure the **store-and-forward** command in global configuration mode to enable the store-and-forward method for packet-forwarding.

NOTE

You must save the configuration and reload for the change to take effect. See the description of the **store-and-forward** command in the *FastIron Command Reference* for more information.

Disabling or re-enabling flow control

You can configure the Brocade device to operate with or without flow control. Flow control is enabled by default globally and on all full-duplex ports. You can disable and re-enable flow control at the Global CONFIG level for all ports. When flow control is enabled globally, you can disable and re-enable it on individual ports.

To disable flow control, enter the **no flow-control** command.

```
device(config)# no flow-control
```

To turn the feature back on, enter the **flow-control** command.

```
device(config)# flow-control
```

Syntax: [no] flow-control

NOTE

For optimal link operation, link ports on devices that do not support 803.3u must be configured with like parameters, such as speed (10,100,1000), duplex (half, full), MDI/MDIX, and Flow Control.

Negotiation and advertisement of flow control

By default, when flow control is enabled globally and auto-negotiation is ON, flow control is enabled and advertised on 10/100/1000M ports. If auto-negotiation is OFF or if the port speed was configured manually, then flow control is not negotiated with or advertised to the peer. For details about auto-negotiation, refer to [Port speed and duplex mode modification](#) on page 53.

To disable flow control capability on a port, enter the following commands.

```
device(config)
# interface ethernet 0/1/21
device(config-if-e1000-0/1/21)# no flow-control
```

To enable flow control negotiation, enter the following commands.

```
device(config)# interface ethernet 0/1/21
device(config-if-e1000-0/1/21)# flow-control neg-on
```

Syntax: [no] flow-control [neg-on]

- *flow-control* [default] - Enable flow control, flow control negotiation, and advertise flow control
- *no flow-control neg-on* - Disable flow control negotiation
- *no flow-control* - Disable flow control, flow control negotiation, and advertising of flow control

After flow control negotiation is enabled using the **flow-control neg-on** command option, flow control is enabled or disabled depending on the peer advertisement.

Commands may be entered in IF (single port) or MIF (multiple ports at once) mode.

```
device(config)# interface ethernet 0/1/21
device(config-if-e1000-0/1/21)# no flow-control
```

This command disables flow control on port 0/1/21.

```
device(config)# interface ethernet 0/1/11 to 0/1/15
device(config-mif-0/1/11-0/1/15)# no flow-control
```

This command disables flow control on ports 0/1/11 to 0/1/15.

Displaying flow-control status

The **show interface** command with the appropriate parameters displays configuration, operation, and negotiation status where applicable.

For example, on a FastIron Stackable device, issuing the command for 10/100/1000M port 0/1/21 displays the following output.

```
device# show interfaces ethernet 0/1/21
GigabitEthernet0/1/21 is up, line protocol is up
Port up for 30 minutes 20 seconds
Hardware is GigabitEthernet, address is 0000.0004.4014 (bia 0000.0004.4014)
  Configured speed auto, actual 100Mbit, configured duplex fdx, actual fdx
  Configured mdi mode AUTO, actual MDIX
  Member of L2 VLAN ID 1, port is untagged, port state is LISTENING
  BPDU Guard is disabled, Root Protect is disabled
  STP configured to ON, priority is level0
  Flow Control is config enabled, oper enabled, negotiation disabled

Mirror disabled, Monitor disabled
Not member of any active trunks
Not member of any configured trunks
No port name
Inter-Packet Gap (IPG) is 96 bit times
300 second input rate: 0 bits/sec, 0 packets/sec, 0.00% utilization
300 second output rate: 0 bits/sec, 0 packets/sec, 0.00% utilization
0 packets input, 0 bytes, 0 no buffer
Received 0 broadcasts, 0 multicasts, 0 unicasts
0 input errors, 0 CRC, 0 frame, 0 ignored
0 runts, 0 giants
5 packets output, 320 bytes, 0 underruns
Transmitted 0 broadcasts, 5 multicasts, 0 unicasts
0 output errors, 0 collisions
```

NOTE

The port up/down time is required only for physical ports and not for loopback/ ve/ tunnel ports.

Issuing the **show interface** command with the appropriate parameters on a FSX device displays the following output:

```
device# show interface ethernet 18/1
GigabitEthernet18/1 is up, line protocol is up
Port up for 50 seconds
  Hardware is GigabitEthernet, address is 0000.0028.0600 (bia 0000.0028.0798)
  Configured speed auto, actual 1Gbit, configured duplex fdx, actual fdx
  Configured mdi mode AUTO, actual MDIX
  Member of 4 L2 VLANs, port is tagged, port state is FORWARDING
  BPDU guard is Disabled, ROOT protect is Disabled
  Link Error Dampening is Disabled
  STP configured to ON, priority is level0, flow control enabled
  Flow Control is config enabled, oper enabled, negotiation disabled
  mirror disabled, monitor disabled
  Not member of any active trunks
  Not member of any configured trunks
  No port name
  IPG MII 96 bits-time, IPG GMII 96 bits-time
  IP MTU 1500 bytes, encapsulation ethernet
  300 second input rate: 0 bits/sec, 0 packets/sec, 0.00% utilization
  300 second output rate: 848 bits/sec, 0 packets/sec, 0.00% utilization
  0 packets input, 0 bytes, 0 no buffer
  Received 0 broadcasts, 0 multicasts, 0 unicasts
  0 input errors, 0 CRC, 0 frame, 0 ignored
  0 runts, 0 giants
  10251 packets output, 1526444 bytes, 0 underruns
  Transmitted 1929 broadcasts, 8293 multicasts, 29 unicasts
  0 output errors, 0 collisions
```

The line highlighted in bold will resemble one of the following, depending on the configuration:

- If flow control negotiation is enabled (and a neighbor advertises "Pause-Not Capable"), the display shows:

```
Flow Control is config enabled, oper disabled, negotiation enabled
```

- If flow control negotiation is enabled (and a neighbor advertises "Pause-Capable"), the display shows:

```
Flow Control is config enabled, oper enabled, negotiation enabled
```

- If flow control is enabled, and flow control negotiation is disabled, the display shows:

```
Flow Control is config enabled, oper enabled, negotiation disabled
```

- If flow control is disabled, the display shows:

```
Flow control is config disabled, oper disabled
```

Symmetric flow control on FCX and ICX devices

In addition to *asymmetric* flow control, FCX and ICX devices support *symmetric* flow control, meaning they can both receive and transmit 802.3x PAUSE frames.

By default on FCX devices, packets are dropped from the end of the queue at the egress port (tail drop mode), when the maximum queue limit is reached. Conversely, when symmetric flow control is enabled, packets are guaranteed delivery since they are managed at the ingress port and no packets are dropped.

Symmetric flow control addresses the requirements of a lossless service class in an Internet Small Computer System Interface (iSCSI) environment. It is supported on FCX and ICX standalone units as well as on all FCX and ICX units in a traditional stack.

About XON and XOFF thresholds

An 802.3x PAUSE frame is generated when the buffer limit at the ingress port reaches or exceeds the port's upper watermark threshold (XOFF limit). The PAUSE frame requests that the sender stop transmitting traffic for a period of time. The time allotted enables the egress and ingress queues to be cleared. When the ingress queue falls below the port's lower watermark threshold (XON limit), an 802.3x PAUSE frame with a quanta of 0 (zero) is generated. The PAUSE frame requests that the sender resume sending traffic normally.

Each 1G, 10G, and 40G port is configured with a default total number of buffers as well as a default XOFF and XON threshold. The defaults are different for 1G ports versus 10G or 40G ports. Also, the default XOFF and XON thresholds are different for jumbo mode versus non-jumbo mode. The defaults are shown in [About XON and XOFF thresholds](#).

TABLE 10 XON and XOFF default thresholds

	Limit when Jumbo disabled / % of buffer limit	Limit when Jumbo enabled / % of buffer limit
1G ports		
Total buffers	272	272
XOFF	240 / 91%	216 / 82%
XON	200 / 75%	184 / 70%
10G ports		
Total buffers	416	416
XOFF	376 / 91%	336 / 82%
XON	312 / 75%	288 / 70%
40G ports		
Total buffers	960	960

TABLE 10 XON and XOFF default thresholds (continued)

	Limit when Jumbo disabled / % of buffer limit	Limit when Jumbo enabled / % of buffer limit
XOFF	832 (87%)	832 (87%)
XON	720 (75%)	720 (75%)

If necessary, you can change the total buffer limits and the XON and XOFF default thresholds. Refer to [Changing the total buffer limits](#) on page 66 and [Changing the XON and XOFF thresholds](#) on page 66, respectively.

Configuration notes and feature limitations for symmetric flow control

Note the following configuration notes and feature limitations before enabling symmetric flow control.

- Symmetric flow control is supported on FCX and ICX devices only. It is not supported on other FastIron models.
- Symmetric flow control is supported on all 1G,10G, and 40G data ports on FCX and ICX devices.
- Symmetric flow control is not supported on stacked ports or across units in a stack. If you are using symmetric flow control on stacked ports or across units in a stack be aware that:
 - It is unrealistic to infer that lossless service exists across stacked units.
 - Symmetric flow control is not priority aware; oversubscription of one priority may cause the dropping of higher priority controls in stacked links. The loss of these priority controls results in a broken stack.
 - The system depends on buffer resources to ensure quality of service. Under symmetric flow control, persistent congestions may leave a buffer resource vulnerable to exhaustion. An example is where bandwidth of ingress ports is greater than egress ports — a packet receives on a 10G port, but then forwards the packet to a 1G port. If the buffers are exhausted, there is no guarantee of quality of service. The end result is an unstable system with flapping protocols.
 - In a stacked environment, pause frames are not propagated from one stack unit to another, as a result they may hold buffers up to a core limit due to multiple port congestions. Under this condition, the stack may break.
 - Not propagating pause frames also prevents head-of-line (HOL) blocking conditions for stacked ports, which are normally used as aggregation links. Stacked ports or trunks are flow control disabled for both transmit and receive, HOL blocking may occur when symmetric flow control is enabled. This means that a peer can stop transmitting traffic streams unrelated to the congestion stream.
- To use this feature, 802.3x flow control must be enabled globally and per interface on FCX and ICX devices. By default, 802.3x flow control is enabled, but can be disabled with the **no flow-control** command.
- The following QoS features are not supported together with symmetric flow control:
 - Dynamic buffer allocation — CLI commands **qd-descriptor** and **qd-buffer**)
 - Buffer profiles — CLI command **buffer-profile port-region**)
 - DSCP-based QoS — CLI command **trust dscp**)

NOTE

Although the above QoS features are not supported with symmetric flow control, the CLI will still accept these commands. The last command issued will be the one placed into effect on the device. For example, if **trust dscp** is enabled after **symmetric-flow-control** is enabled, symmetric flow control will be disabled and trust dscp will be placed into effect. Make sure you do not enable incompatible QoS features when symmetric flow control is enabled on the device.

Enabling and disabling symmetric flow control

By default, symmetric flow control is disabled and tail drop mode is enabled. However, because flow control is enabled by default on all full-duplex ports, these ports will always honor received 802.3x Pause frames, whether or not symmetric flow control is enabled.

To enable symmetric flow control globally on all full-duplex data ports of a standalone unit, enter the **symmetric-flow-control enable** command.

```
device(config)# symmetric-flow-control enable
```

To enable symmetric flow control globally on all full-duplex data ports of a particular unit in a traditional stack, enter the **symmetric-flow-control enable** command with the appropriate parameters.

```
device(config)# symmetric-flow-control enable unit 4
```

Syntax: **[no] symmetric-flow-control enable** [unit *stack-unit*]

The *stack-unit* parameter specifies one of the units in a stacking system. Master/Standby/Members are examples of a stack-unit.

To disable symmetric flow control once it has been enabled, use the **no** form of the command.

Changing the XON and XOFF thresholds

This section describes how to change the XON and XOFF thresholds described in [About XON and XOFF thresholds](#) on page 64.

To change the thresholds for all 1G ports, enter a command such as the following.

```
device(config)# symmetric-flow-control set 1 xoff 91 xon 75
```

To change the thresholds for all 10G ports, enter a command such as the following.

```
device(config)# symmetric-flow-control set 2 xoff 91 xon 75
```

In the above configuration examples, when the XOFF limit of 91% is reached or exceeded, the Brocade device will send PAUSE frames to the sender telling it to stop transmitting data temporarily. When the XON limit of 75% is reached, the Brocade device will send PAUSE frames to the sender telling it to resume sending data.

Syntax: **symmetric-flow-control set { 1 | 2 } xoff % xon %**

symmetric-flow-control set 1 sets the XOFF and XON limits for 1G ports.

symmetric-flow-control set 2 sets the XOFF and XON limits for 10G ports.

For *xoff %*, the % minimum value is 60% and the maximum value is 95%.

For *xon %*, the % minimum value is 50% and the maximum value is 90%.

Use the **show symmetric** command to view the default or configured XON and XOFF thresholds. Refer to [Displaying symmetric flow control status](#) on page 67.

Changing the total buffer limits

This section describes how to change the total buffer limits described in [About XON and XOFF thresholds](#) on page 64. You can change the limits for all 1G ports and for all 10G ports.

To change the total buffer limit for all 1G ports, enter a command such as the following.

```
device(config)# symmetric-flow-control set 1 buffers 320
Total buffers modified, 1G: 320, 10G: 128
```

To change the total buffer limit for all 10G ports, enter a command such as the following.

```
device(config)# symmetric-flow-control set 2 buffers 128
Total buffers modified, 1G: 320, 10G: 128
```

Syntax: **symmetric-flow-control set { 1 | 2 } buffers value**

symmetric-flow-control set 1 buffers value sets the total buffer limits for 1G ports. The default value is 272. You can specify a number from 64 - 320.

symmetric-flow-control set 2 buffers value sets the total buffer limits for 10G ports. The default value is 416. You can specify a number from 64 - 1632.

Use the **show symmetric** command to view the default or configured total buffer limits. Refer to [Displaying symmetric flow control status](#) on page 67.

Displaying symmetric flow control status

The **show symmetric-flow-control** command displays the status of symmetric flow control as well as the default or configured total buffer limits and XON and XOFF thresholds.

```
device(config)# show symmetric
Symmetric Flow Control Information:
-----
Symmetric Flow Control is enabled on units: 2 3
Buffer parameters:
1G Ports:
    Total Buffers : 272
    XOFF Limit    : 240 (91%)
    XON Limit     : 200 (75%)
10G Ports:
    Total Buffers : 416
    XOFF Limit    : 376 (91%)
    XON Limit     : 312 (75%)
```

Syntax: **show symmetric-flow-control**

PHY FIFO Rx and Tx depth configuration

PHY devices on Brocade devices contain transmit and receive synchronizing FIFOs to adjust for frequency differences between clocks. The **phy-fifo-depth** command allows you to configure the depth of the transmit and receive FIFOs. There are 4 settings (0-3) with 0 as the default. A higher setting indicates a deeper FIFO.

The default setting works for most connections. However, if the clock differences are greater than the default will handle, CRCs and errors will begin to appear on the ports. Raising the FIFO depth setting will adjust for clock differences.

Brocade recommends that you disable the port before applying this command, and re-enable the port. Applying the command while traffic is flowing through the port can cause CRC and other errors for any packets that are actually passing through the PHY while the command is being applied.

Syntax: **[no] phy-fifo-depth setting**

- *setting* is a value between 0 and 3. (0 is the default.)

This command can be issued for a single port from the IF config mode or for multiple ports from the MIF config mode.

NOTE

Higher settings give better tolerance for clock differences with the partner phy, but may marginally increase latency as well.

Interpacket Gap (IPG) on a FastIron X Series switch

IPG is the time delay, in bit time, between frames transmitted by the device. You configure IPG at the interface level. The command you use depends on the interface type on which IPG is being configured.

The default interpacket gap is 96 bits-time, which is 9.6 microseconds for 10 Mbps Ethernet, 960 nanoseconds for 100 Mbps Ethernet, 96 nanoseconds for 1 Gbps Ethernet, and 9.6 nanoseconds for 10 Gbps Ethernet.

IPG on a FastIron X series switch configuration notes

- The CLI syntax for IPG differs on FastIron X Series devices compared to FastIron Stackable devices. This section describes the configuration procedures for FastIron X Series devices. For FastIron Stackable devices, refer to [IPG on FastIron Stackable devices](#) on page 69.
- IPG configuration commands are based on "port regions". All ports within the same port region should have the same IPG configuration. If a port region contains two or more ports, changes to the IPG configuration for one port are applied to all ports in the same port region. When you enter a value for IPG, the CLI displays the ports to which the IPG configuration is applied.

```
device(config-if-e1000-7/1)# ipg-gmii 120
IPG 120(112) has been successfully configured for ports 7/1 to 7/12
```

- When you enter a value for IPG, the device applies the closest valid IPG value for the port mode to the interface. For example, if you specify 120 for a 1 Gbps Ethernet port in 1 Gbps mode, the device assigns 112 as the closest valid IPG value to program into hardware.

Configuring IPG on a Gbps Ethernet port

On a Gbps Ethernet port, you can configure IPG for 10/100 mode and for Gbps Ethernet mode.

10/100M mode

To configure IPG on a Gbps Ethernet port for 10/100M mode, enter the following command.

```
device(config)# interface ethernet 7/1
device(config-if-e1000-7/1)# ipg-mii 120
IPG 120(120) has been successfully configured for ports 7/1 to 7/12
```

Syntax: [no] ipg-mii *bit-time*

Enter 12-124 for *bit time* . The default is 96 bit time.

1G mode

To configure IPG on a Gbps Ethernet port for 1-Gbps Ethernet mode, enter commands such as the following.

```
device(config)# interface ethernet 7/1
device(config-if-e1000-7/1)# ipg-gmii 120
IPG 120(112) has been successfully configured for ports 0/7/1 to 7/12
```

Syntax: [no] ipg-gmii *bit-time*

Enter 48 - 112 for *bit time* . The default is 96 bit time.

Configuring IPG on a 10 Gbps Ethernet interface

To configure IPG on a 10 Gbps Ethernet interface, enter commands such as the following.

```
device(config)# interface ethernet 9/1
device(config-if-e10000-9/1)# ipg-xgmii 120
IPG 120(128) has been successfully configured for port 9/1
```

Syntax: [no] ipg-xgmii *bit-time*

Enter 96-192 for *bit time* . The default is 96 bit time.

IPG on FastIron Stackable devices

On FCX and ICX devices, you can configure an IPG for each port. An IPG is a configurable time delay between successive data packets. You can configure an IPG with a range from 48-120 bit times in multiples of 8, with a default of 96. The IPG may be set from either the interface configuration level or the multiple interface level.

IPG configuration notes

- The CLI syntax for IPG differs on FastIron Stackable devices compared to FastIron X Series devices. This section describes the configuration procedures for FastIron Stackable devices. For FastIron X Series devices, refer to [Interpacket Gap \(IPG\) on a FastIron X Series switch](#) on page 67.
- When an IPG is applied to a trunk group, it applies to all ports in the trunk group. When you are creating a new trunk group, the IPG setting on the primary port is automatically applied to the secondary ports.
- This feature is supported on 10/100/1000M ports.

Configuring IPG on a 10/100/1000M port

To configure an IPG of 112 on Ethernet interface 0/1/21, for example, enter the following command.

```
device(config)# interface ethernet 0/1/21
device(config-if-e1000-0/1/21)# ipg 112
```

For multiple interface levels, to configure IPG for ports 0/1/11 and 0/1/14 through 0/1/17, enter the following commands.

```
device(config)# interface ethernet 0/1/11 ethernet 0/1/14 to 0/1/17
device(config-mif-0/1/11,0/1/14-0/1/17)# ipg 104
```

Syntax: [no] ipg *value*

For *value* , enter a number in the range from 48-120 bit times in multiples of 8. The default is 96.

As a result of the above configuration, the output from the show interface Ethernet 0/1/21 command is as follows.

```
device# show interfaces ethernet 0/1/21
GigabitEthernet 0/1/21 is up, line protocol is up
Port up for 40 seconds
Hardware is GigabitEthernet, address is 0000.0004.4014 (bia 0000.0004.4014)
Configured speed auto, actual 100Mbit, configured duplex fdx, actual fdx
Configured mdi mode AUTO, actual MDIX
Member of L2 VLAN ID 1, port is untagged, port state is FORWARDING
BPDU Guard is disabled, Root Protect is disabled
STP configured to ON, priority is level0
Flow Control is config enabled, oper enabled, negotiation disabled
Mirror disabled, Monitor disabled
Not member of any active trunks
Not member of any configured trunks
No port name
Inter-Packet Gap (IPG) is 112 bit times
IP MTU 10222 bytes
300 second input rate: 0 bits/sec, 0 packets/sec, 0.00% utilization
300 second output rate: 248 bits/sec, 0 packets/sec, 0.00% utilization
0 packets input, 0 bytes, 0 no buffer
Received 0 broadcasts, 0 multicasts, 0 unicasts
0 input errors, 0 CRC, 0 frame, 0 ignored
0 runts, 0 giants
80 packets output, 5120 bytes, 0 underruns
Transmitted 0 broadcasts, 80 multicasts, 0 unicasts
0 output errors, 0 collisions
```

Enabling and disabling support for 100BaseTX

For FastIron X Series devices, you can configure a 1000Base-TX SFP (part number E1MG-TX) to operate at a speed of 100 Mbps. To do so, enter the **100-tx** command at the Interface level of the CLI.

```
device(config-if-e1000-11)# 100-tx
```

After the link is up, it will be in 100M/full-duplex mode, as shown in the following example.

```
device# show interface brief ethernet 11
Port Link State Dupl Speed
Trunk Tag Priori MAC Name
11 Up Forward Full 100M
None No level10 0000.0013.c74b
```

The **show media** command will display the SFP transceiver as *1G M-TX*.

Syntax: [no] 100-tx

To disable support, enter the **no** form of the command.

100BaseTX configuration notes

- This feature requires that autonegotiation be enabled on the other end of the link.
- Although combo ports (ports 1 - 4) on Hybrid Fiber (HF) models support the 1000Base-TX SFP, they cannot be configured to operate at 100 Mbps. The 100 Mbps operating speed is supported only with non-combo ports (ports 5-24).
- The FCX624S-F is the only FCX model that supports the 1000Base-TX SFP module, and only on the non-combo ports (ports 5-24). The FCX624S-F does not have a specific command to enable the 1000Base-TX SFP optic at 100 Mbps. You must manually configure it with the **speed-duplex 100-full** command. Refer to [Port speed and duplex mode configuration syntax](#) on page 53.
- 1000Base-TX modules must be configured individually, one interface at a time.
- 1000Base-TX modules do not support Digital Optical Monitoring.
- This module requires a Cat5 cable and uses an RJ45 connector.
- Hotswap is supported for this module when it is configured in 100M mode.

Enabling and disabling support for 100BaseFX

Some Brocade devices support 100BaseFX fiber transceivers. After you physically install a 100BaseFX transceiver, you must enter a CLI command to enable it. For information about supported SFP and SFP+ transceivers on ICX devices, refer to the *Brocade Optics Family Datasheet* on the Brocade website.

Enabling and disabling 100BaseFX on Chassis-based and stackable devices

NOTE

The following procedure applies to Stackable devices and to Chassis-based 100/1000 Fiber interface modules only. The CLI syntax for enabling and disabling 100BaseFX support on these devices differs than on a Compact device. Make sure you refer to the appropriate procedures. These are not supported on ICX 6430 and ICX 6450 devices.

FastIron devices support the following types of SFPs for 100BaseFX:

- *Multimode SFP* - maximum distance is 2 kilometers
- *Long Reach (LR)* - maximum distance is 40 kilometers

- *Intermediate Reach (IR)* - maximum distance is 15 kilometers

For information about supported SFP and SFP+ transceivers on FastIron devices, refer to the *Brocade Optics Family Datasheet* on the Brocade website.

NOTE

Connect the 100BaseFX fiber transceiver *after* configuring both sides of the link. Otherwise, the link could become unstable, fluctuating between up and down states.

To enable support for 100BaseFX on an FSX fiber port or on a Stackable switch, enter commands such as the following.

```
device(config)# interface ethernet 1/1/6
device(config-if-1/1/6)# 100-fx
```

The above commands enable 100BaseFX on port 6 in slot 1.

Syntax: [no] 100-fx

To disable 100BaseFX support on a fiber port, enter the **no** form of the command. Note that you must disable 100BaseFX support before inserting a different type of module in the same port. Otherwise, the device will not recognize traffic traversing the port.

Changing the Gbps fiber negotiation mode

The globally configured Gbps negotiation mode is the default mode for all Gbps fiber ports. You can override the globally configured default and set individual ports to the following:

NOTE

Gbps negotiation is not supported on ICX 6430, ICX 6450, and ICX 6650 devices.

- **Negotiate-full-auto** - The port first tries to perform a handshake with the other port to exchange capability information. If the other port does not respond to the handshake attempt, the port uses the manually configured configuration information (or the defaults if an administrator has not set the information). This is the default.
- **Auto-Gbps** - The port tries to perform a handshake with the other port to exchange capability information.
- **Negotiation-off** - The port does not try to perform a handshake. Instead, the port uses configuration information manually configured by an administrator.

To change the mode for individual ports, enter commands such as the following.

```
device(config)
# interface ethernet 1 to 4
device(config-mif-1-4)# gig-default auto-gig
```

This command overrides the global setting and sets the negotiation mode to auto-Gbps for ports 1 - 4.

Syntax: gig-default{ neg-full-auto | auto-gig | neg-off }

NOTE

When Gbps negotiation mode is turned off (CLI command **gig-default neg-off**), the Brocade device may inadvertently take down both ends of a link. This is a hardware limitation for which there is currently no workaround.

Configuration considerations for Gbps fiber negotiation mode

For Fiber ports, the configuration is considered invalid if the Gbps negotiation mode is enabled on one end of the link and Gbps negotiation mode is turned off at the other end.

The following tables provide a list of invalid configurations on fiber ports.

TABLE 11 List of invalid configurations

ICX 7450 / ICX 7250 (1G fiber port) configuration	Link Partner - ICX 7450 / ICX 7250 / ICX 6430 / ICX 6450 / ICX 6610 / ICX 6650 / FCX / FSX 800 / FSX 1600 (1G fiber port) configuration
100-fx	1000-full
100-fx	neg-off

TABLE 12 List of invalid configurations

ICX 7450 / ICX 7750 (10G fiber port) configuration	Link Partner - ICX 6430 / ICX 6450 / ICX 6610 / ICX 6650 / FCX / FSX 800 / FSX 1600 / ICX 7450 / ICX 7250 (1G fiber port) configuration
1000-full + neg-off	1000-full
1000-full (with default auto-gig)	neg-off

TABLE 13 List of invalid configurations

ICX 7450 / ICX 7750 (10G fiber port) configuration	Link Partner - ICX 6430 / ICX 6450 / ICX 6610 / ICX 6650 / FCX / FSX 800 / FSX 1600 / ICX 7450 / ICX 7750 / ICX 7250 (10G fiber port) configuration
1000-full (with default auto-gig)	1000-full and neg-off

Port priority (QoS) modification

You can give preference to the inbound traffic on specific ports by changing the Quality of Service (QoS) level on those ports. For information and procedures, refer to "Quality of Service" chapter in the *FastIron Ethernet Switch Traffic Management Guide*.

Dynamic configuration of Voice over IP (VoIP) phones

You can configure a FastIron device to automatically detect and re-configure a VoIP phone when it is physically moved from one port to another within the same device. To do so, you must configure a *voice VLAN ID* on the port to which the VoIP phone is connected. The software stores the voice VLAN ID in the port database for retrieval by the VoIP phone.

The dynamic configuration of a VoIP phone works in conjunction with the VoIP phone discovery process. Upon installation, and sometimes periodically, a VoIP phone will query the Brocade device for VoIP information and will advertise information about itself, such as, device ID, port ID, and platform. When the Brocade device receives the VoIP phone query, it sends the voice VLAN ID in a reply packet back to the VoIP phone. The VoIP phone then configures itself within the voice VLAN.

As long as the port to which the VoIP phone is connected has a voice VLAN ID, the phone will configure itself into that voice VLAN. If you change the voice VLAN ID, the software will immediately send the new ID to the VoIP phone, and the VoIP phone will re-configure itself with the new voice VLAN.

VoIP configuration notes

- This feature works with any VoIP phone that:
 - Runs CDP
 - Sends a VoIP VLAN query message
 - Can configure its voice VLAN after receiving the VoIP VLAN reply
- Automatic configuration of a VoIP phone will not work if one of the following applies:
 - You do not configure a voice VLAN ID for a port with a VoIP phone

- You remove the configured voice VLAN ID from a port without configuring a new one
- You remove the port from the voice VLAN
- Make sure the port is able to intercept CDP packets (**cdp run** command).
- Some VoIP phones may require a reboot after configuring or re-configuring a voice VLAN ID. For example, if your VoIP phone queries for VLAN information only once upon boot up, you must reboot the VoIP phone before it can accept the VLAN configuration. If your phone is powered by a PoE device, you can reboot the phone by disabling then re-enabling the port.

Enabling dynamic configuration of a Voice over IP (VoIP) phone

You can create a voice VLAN ID for a port, or for a group of ports.

To create a voice VLAN ID for a port, enter commands such as the following.

```
device(config)
# interface ethernet 2
device(config-if-e1000-2)# voice-vlan 1001
```

To create a voice VLAN ID for a group of ports, enter commands such as the following.

```
device(config)
# interface ethernet 1-8
device(config-mif-1-8)# voice-vlan 1001
```

Syntax: [**no**] **voice-vlan** *voice-vlan-num*

where *voice-vlan-num* is a valid VLAN ID between 1 - 4095.

To remove a voice VLAN ID, use the **no** form of the command.

Viewing voice VLAN configurations

You can view the configuration of a voice VLAN for a particular port or for all ports.

To view the voice VLAN configuration for a port, specify the port number with the **show voice-vlan** command. The following example shows the command output results.

```
device# show voice-vlan ethernet 2
Voice vlan ID for port 2: 1001
```

The following example shows the message that appears when the port does not have a configured voice VLAN.

```
device# show voice-vlan ethernet 2
Voice vlan is not configured for port 2.
```

To view the voice VLAN for all ports, use the **show voice-vlan** command. The following example shows the command output results.

```
device# show voice-vlan
Port ID      Voice-vlan
2            1001
8            150
15           200
```

Syntax: **show voice-vlan** [**ethernet** *port*]

Port flap dampening configuration

Port Flap Dampening increases the resilience and availability of the network by limiting the number of port state transitions on an interface.

If the port link state toggles from up to down for a specified number of times within a specified period, the interface is physically disabled for the specified wait period. Once the wait period expires, the port link state is re-enabled. However, if the wait period is set to zero (0) seconds, the port link state will remain disabled until it is manually re-enabled.

Port flap dampening configuration notes

- When a flap dampening port becomes a member of a trunk group, that port, as well as all other member ports of that trunk group, will inherit the primary port configuration. This means that the member ports will inherit the primary port flap dampening configuration, regardless of any previous configuration.
- The Brocade device counts the number of times a port link state toggles from "up to down", and not from "down to up".
- The sampling time or window (the time during which the specified toggle threshold can occur before the wait period is activated) is triggered when the first "up to down" transition occurs.
- "Up to down" transitions include UDLD-based toggles, as well as the physical link state.

Configuring port flap dampening on an interface

This feature is configured at the interface level.

```
device(config)# interface ethernet 1/2/1
device(config-if-e10000-1/2/1)# link-error-disable 10 3 10
```

Syntax: `[no] link-error-disable toggle-threshold sampling-time-in-sec wait-time-in-sec`

The *toggle-threshold* is the number of times a port link state goes from up to down and down to up before the wait period is activated. Enter a value from 1 - 50.

The *sampling-time-in-sec* is the amount of time during which the specified toggle threshold can occur before the wait period is activated. The default is 0 seconds. Enter 1 - 65535 seconds.

The *wait-time-in-sec* is the amount of time the port remains disabled (down) before it becomes enabled. Enter a value from 0 - 65535 seconds; 0 indicates that the port will stay down until an administrative override occurs.

Configuring port flap dampening on a trunk

You can configure the port flap dampening feature on the primary port of a trunk using the **link-error-disable** command. Once configured on the primary port, the feature is enabled on all ports that are members of the trunk. You cannot configure port flap dampening on port members of the trunk.

Enter commands such as the following on the primary port of a trunk.

```
device(config)# interface ethernet 1/2/1
device(config-if-e10000-1/2/1)# link-error-disable 10 3 10
```

Re-enabling a port disabled by port flap dampening

A port disabled by port flap dampening is automatically re-enabled once the wait period expires; however, if the wait period is set to zero (0) seconds, you must re-enable the port by entering the following command on the disabled port.

```
device(config)# interface ethernet 1/2/1
device(config-if-e10000-1/2/1)# no link-error-disable 10 3 10
```

Displaying ports configured with port flap dampening

Ports that have been disabled due to the port flap dampening feature are identified in the output of the **show link-error-disable** command. The following shows an example output.

```
device# show link-error-disable
Port 1/2/1 is forced down by link-error-disable.
```

Use the **show link-error-disable all** command to display the ports with the port flap dampening feature enabled.

For FastIron Stackable devices, the output of the command shows the following.

```
device# show link-error-disable all
Port1/8/1 is configured for link-error-disable
      threshold:1, sampling_period:10, waiting_period:0
Port1/8/2 is configured for link-error-disable
      threshold:1, sampling_period:10, waiting_period:0
Port1/8/3 is configured for link-error-disable
      threshold:1, sampling_period:10, waiting_period:0
Port1/8/4 is configured for link-error-disable
      threshold:1, sampling_period:10, waiting_period:0
Port1/8/5 is configured for link-error-disable
      threshold:4, sampling_period:10, waiting_period:2
Port1/8/9 is configured for link-error-disable
      threshold:2, sampling_period:20, waiting_period:0
```

For FastIron X Series devices, the output of the command shows the following.

```
device# show link-error-disable all
Port #      Threshold  Sampling-Time  Shutoff-Time  State  Counter
-----
11         3           120           600          Idle   N/A
12         3           120           500          Down   424
```

The following table defines the port flap dampening statistics displayed by the **show link-error-disable all** command.

TABLE 14 Output of show link-error-disable

Column	Description
Port #	The port number.
Threshold	The number of times the port link state will go from up to down and down to up before the wait period is activated.
Sampling-Time	The number of seconds during which the specified toggle threshold can occur before the wait period is activated.
Shutoff-Time	The number of seconds the port will remain disabled (down) before it becomes enabled. A zero (0) indicates that the port will stay down until an administrative override occurs.
State	The port state can be one of the following: <ul style="list-style-type: none"> Idle - The link is normal and no link state toggles have been detected or sampled. Down - The port is disabled because the number of sampled errors exceeded the configured threshold. Err - The port sampled one or more errors.
Counter	<ul style="list-style-type: none"> If the port state is Idle, this field displays N/A. If the port state is Down, this field shows the remaining value of the shutoff timer. If the port state is Err, this field shows the number of errors sampled.

Syntax: show link-error-disable [all]

Also, in FastIron X Series devices, the **show interface** command indicates if the port flap dampening feature is enabled on the port.

```
device# show interface ethernet 15
GigabitEthernet15 is up, line protocol is up
Link Error Dampening is Enabled
Port up for 6 seconds
  Hardware is GigabitEthernet, address is 0000.0000.010e (bia 0000.0000.010e)
  Configured speed auto, actual 1Gbit, configured duplex fdx, actual fdx
  Configured mdi mode AUTO, actual MDIX
device# show interface ethernet 17
GigabitEthernet17 is ERR-DISABLED, line protocol is down
Link Error Dampening is Enabled
Port down for 40 seconds
  Hardware is GigabitEthernet, address is 0000.0000.010e (bia 0000.0000.010e)
  Configured speed auto, actual unknown, configured duplex fdx, actual unknown
```

The line "Link Error Dampening" displays "Enabled" if port flap dampening is enabled on the port or "Disabled" if the feature is disabled on the port. The feature is enabled on the ports in the two examples above. Also, the characters "ERR-DISABLED" is displayed for the "GbpsEthernet" line if the port is disabled because of link errors.

Syntax: show interface ethernet port-number

In addition to the show commands above, the output of the **show interface brief** command for FastIron X Series indicates if a port is down due to link errors.

```
device# show interface brief e17
Port  Link   State   Dupl Speed Trunk Tag Priori MAC           Name
17    ERR-DIS
None      None None   15    Yes level0 0000.0000.010e
```

The ERR-DIS entry under the "Link" column indicates the port is down due to link errors.

NOTE

If a port name is longer than five characters, the port name is truncated in the output of the **show interface brief** command.

Syslog messages for port flap dampening

The following Syslog messages are generated for port flap dampening.

- If the threshold for the number of times that a port link toggles from "up" to "down" then "down" to "up" has been exceeded, the following Syslog message is displayed.

```
0d00h02m10s:I:ERR_DISABLE: Link flaps on port ethernet 16 exceeded threshold; port in err-disable state
```

- If the wait time (port is down) expires and the port is brought up the following Syslog message is displayed.

```
0d00h02m41s:I:ERR_DISABLE: Interface ethernet 16, err-disable recovery timeout
```

Port loop detection

This feature allows the Brocade device to disable a port that is on the receiving end of a loop by sending test packets. You can configure the time period during which test packets are sent.

Types of loop detection

There are two types of loop detection; Strict Mode and Loose Mode. In Strict Mode, a port is disabled only if a packet is looped back to that same port. Strict Mode overcomes specific hardware issues where packets are echoed back to the input port. In Strict Mode, loop detection must be configured on the physical port.

In Loose Mode, loop detection is configured on the VLAN of the receiving port. Loose Mode disables the receiving port if packets originate from any port or VLAN on the same device. The VLAN of the receiving port must be configured for loop detection in order to disable the port.

Recovering disabled ports

Once a loop is detected on a port, it is placed in Err-Disable state. The port will remain disabled until one of the following occurs:

- You manually disable and enable the port at the Interface Level of the CLI.
- You enter the command **clear loop-detection**. This command clears loop detection statistics and enables all Err-Disabled ports.
- The device automatically re-enables the port. To set your device to automatically re-enable Err-Disabled ports, refer to [Configuring the device to automatically re-enable ports](#) on page 78.

Port loopback detection configuration notes

- Loopback detection packets are sent and received on both tagged and untagged ports. Therefore, this feature cannot be used to detect a loop across separate devices.

The following information applies to Loose Mode loop detection:

- With Loose Mode, two ports of a loop are disabled.
- Different VLANs may disable different ports. A disabled port affects every VLAN using it.
- Loose Mode floods test packets to the entire VLAN. This can impact system performance if too many VLANs are configured for Loose Mode loop detection.

NOTE

Brocade recommends that you limit the use of Loose Mode. If you have a large number of VLANs, configuring loop detection on all of them can significantly affect system performance because of the flooding of test packets to all configured VLANs. An alternative to configuring loop detection in a VLAN-group of many VLANs is to configure a separate VLAN with the same tagged port and configuration, and enable loop detection on this VLAN only.

NOTE

When loop detection is used with Layer 2 loop prevention protocols, such as spanning tree (STP), the Layer 2 protocol takes higher priority. Loop detection cannot send or receive probe packets if ports are blocked by Layer 2 protocols, so it does not detect Layer 2 loops when STP is running because loops within a VLAN have been prevented by STP. Loop detection running in Loose Mode can detect and break Layer 3 loops because STP cannot prevent loops across different VLANs. In these instances, the ports are not blocked and loop detection is able to send out probe packets in one VLAN and receive packets in another VLAN. In this way, loop detection running in Loose Mode disables both ingress and egress ports.

Enabling loop detection

Use the **loop-detection** command to enable loop detection on a physical port (Strict Mode) or a VLAN (Loose Mode). Loop detection is disabled by default. The following example shows a Strict Mode configuration.

```
device(config)# interface ethernet 1/1/1
device(config-if-e1000-1/1/1)# loop-detection
```

The following example shows a Loose Mode configuration.

```
device(config)# vlan20
device(config-vlan-20)# loop-detection
```

By default, the port will send test packets every one second, or the number of seconds specified by the **loop-detection-interval** command. Refer to [Configuring a global loop detection interval](#) on page 78.

Syntax: [no] loop-detection

Use the [no] form of the command to disable loop detection.

Configuring a global loop detection interval

The loop detection interval specifies how often a test packet is sent on a port. When loop detection is enabled, the loop detection time unit is 0.1 second, with a default of 10 (one second). The range is from 1 (one tenth of a second) to 100 (10 seconds). You can use the **show loop-detection status** command to view the loop detection interval.

To configure the global loop detection interval, enter a command similar to the following.

```
device(config)# loop-detection-interval 50
```

This command sets the loop-detection interval to 5 seconds (50 x 0.1).

To revert to the default global loop detection interval of 10, enter one of the following.

```
device(config)# loop-detection-interval 10
```

OR

```
device(config)# no loop-detection-interval 50
```

Syntax: [no] loop-detection-interval *number*

where *number* is a value from 1 to 100. The system multiplies your entry by 0.1 to calculate the interval at which test packets will be sent.

Configuring the device to automatically re-enable ports

To configure the Brocade device to automatically re-enable ports that were disabled because of a loop detection, enter the **errdisable recovery cause loop-detection** command.

```
device(config)# errdisable recovery cause loop-detection
```

The above command will cause the Brocade device to automatically re-enable ports that were disabled because of a loop detection. By default, the device will wait 300 seconds before re-enabling the ports. You can optionally change this interval to a value from 10 to 65535 seconds. Refer to [Specifying the recovery time interval](#) on page 78.

Syntax: [no] errdisable recovery cause loop-detection

Use the [no] form of the command to disable this feature.

Specifying the recovery time interval

The recovery time interval specifies the number of seconds the Brocade device will wait before automatically re-enabling ports that were disabled because of a loop detection. (Refer to [Configuring the device to automatically re-enable ports](#) on page 78.) By default, the device will wait 300 seconds. To change the recovery time interval, enter a command such as the following.

```
device(config)# errdisable recovery interval 120
```

The above command configures the device to wait 120 seconds (2 minutes) before re-enabling the ports.

To revert back to the default recovery time interval of 300 seconds (5 minutes), enter one of the following commands.

```
device(config)# errdisable recovery interval 300
```

OR

```
device(config)# no errdisable recovery interval 120
```

Syntax: [no] errdisable recovery interval *seconds*

where *seconds* is a number from 10 to 65535.

Clearing loop-detection

To clear loop detection statistics and re-enable all ports that are in Err-Disable state because of a loop detection, enter the **clear loop-detection** command.

```
device# clear loop-detection
```

Displaying loop-detection information

Use the **show loop-detection status** command to display loop detection status, as shown.

```
device# show loop-detection status
loop detection packets interval: 10 (unit 0.1 sec)
Number of err-disabled ports: 3
You can re-enable err-disable ports one by one by "disable" then "enable"
under interface config, re-enable all by "clear loop-detect", or
configure "errdisable recovery cause loop-detection" for automatic recovery
index port/vlan  status          #errdis  sent-pkts  recv-pkts
1      1/13      untag, LEARNING          0          0          0
2      1/15      untag, BLOCKING          0          0          0
3      1/17      untag, DISABLED          0          0          0
4      1/18      ERR-DISABLE by itself    1          6          1
5      1/19      ERR-DISABLE by vlan 12   0          0          0
6      vlan12    2 ERR-DISABLE ports      2         24          2
```

If a port is errdisabled in Strict mode, it shows "ERR-DISABLE by itself". If it is errdisabled due to its associated vlan, it shows "ERR-DISABLE by vlan ?"

The following command displays the current disabled ports, including the cause and the time.

```
device# show loop-detection disable
Number of err-disabled ports: 3
You can re-enable err-disable ports one by one by "disable" then "enable"
under interface config, re-enable all by "clear loop-detect", or
configure "errdisable recovery cause loop-detection" for automatic recovery
index  port      caused-by  disabled-time
1      1/18      itself     00:13:30
2      1/19      vlan 12    00:13:30
3      1/20      vlan 12    00:13:30
```

This example shows the disabled ports, the cause, and the time the port was disabled. If loop-detection is configured on a physical port, the disable cause will show "itself". For VLANs configured for loop-detection, the cause will be a VLAN.

The following command shows the hardware and software resources being used by the loop-detection feature.

```
Vlans configured loop-detection use 1 HW MAC
Vlans not configured but use HW MAC: 1 10
configuration pool      alloc in-use  avail get-fail  limit  get-mem  size  init
linklist pool           16      6      10      0      3712     6     15   16
```

Displaying loop detection resource information

Use the **show loop-detection resource** command to display the hardware and software resource information on loop detection.

```
device# show loop-detection resource
Vlans configured loop-detection use 1 HW MAC
Vlans not configured but use HW MAC: 1 10
configuration pool      alloc in-use  avail get-fail  limit  get-mem  size  init
linklist pool          16      6    10      0    3712    6    15   16
linklist pool          16     10     6      0    3712   10    16   16
```

Syntax: show loop-detection resource

The following table describes the output fields for this command.

TABLE 15 Field definitions for the **show loop-detection resource** command

Field	Description
alloc	Memory allocated
in-use	Memory in use
avail	Available memory
get-fail	The number of get requests that have failed
limit	The maximum memory allocation
get-mem	The number of get-memory requests
size	The size
init	The number of requests initiated

Displaying loop detection configuration status on an interface

Use the **show interface** command to display the status of loop detection configuration on a particular interface.

```
Brocade# show interface ethernet 1/2/1
10GigabitEthernet1/2/1 is up, line protocol is up
Port up for 1 day 22 hours 43 minutes 5 seconds
Hardware is 10GigabitEthernet, address is 0000.0089.1100 (bia 0000.0089.1118)
Configured speed 10Gbit, actual 10Gbit, configured duplex fdx, actual fdx
Member of 9 L2 VLANs, port is tagged, port state is FORWARDING
BPDU guard is Disabled, ROOT protect is Disabled
Link Error Dampening is Disabled
STP configured to ON, priority is level0
Loop Detection is ENABLED
Flow Control is enabled
Mirror disabled, Monitor disabled
Member of active trunk ports 1/2/1,1/2/2, primary port
Member of configured trunk ports 1/2/1,1/2/2, primary port
No port name
IPG XGMII 96 bits-time
MTU 1500 bytes, encapsulation ethernet
ICL port for BH1 in cluster id 1
300 second input rate: 2064 bits/sec, 3 packets/sec, 0.00% utilization
300 second output rate: 768 bits/sec, 1 packets/sec, 0.00% utilization
171319 packets input, 12272674 bytes, 0 no buffer
Received 0 broadcasts, 63650 multicasts, 107669 unicasts
0 input errors, 0 CRC, 0 frame, 0 ignored
0 runts, 0 giants
51094 packets output, 3925313 bytes, 0 underruns
Transmitted 2 broadcasts, 42830 multicasts, 8262 unicasts
0 output errors, 0 collisions
Relay Agent Information option: Disabled
```


Syslog message due to disabled port in loop detection

The following message is logged when a port is disabled due to loop detection. This message also appears on the console.

```
loop-detection: port ?/?/? vlan ?, detect, putting into err-disable state
```

Shutdown prevention for loop-detection on an interface

Prevents shut down for loop-detect on an interface.

In prior FastIron releases, when a loop detection probe packet was received back on an interface, the corresponding interface would be shut down either permanently or for a specific duration configured by the user. The new shut down prevention for loop-detection functionality allows users to disable the shutdown of a port when the loop detection probe packet is received on an interface. This provides control over deciding which port is allowed to enter into an error-disabled state and go into a shutdown state when a loop is detected. This function can also be used as a test tool to detect Layer 2 and Layer 3 loops in network current data packet flow.

Limitations of shutdown prevention for loop-detection

Shutdown prevention for loop-detection does not allow any corrective action to be taken on the loop. There could be network instability due to the presence of network loops, if adequate corrective measures are not taken by the network administrator.

To enable shutdown prevention for loop detection, follow these steps.

1. Enter global configuration mode.
2. Specify the interface on which you would like to enable the **loop-detection shutdown-disable** command.
3. Enter the **loop-detection shutdown-disable** command.

```
Brocade (config)# interface ethe 1/1/7
Brocade (config-if-e1000-1/1/7)# loop-detection shutdown-disable
```

Periodic log message generation for shutdown prevention

Generates periodic log messages for shutdown prevention.

You can raise a periodic syslog that provides information about loops in the network. When a loop is detected because of a loop detection protocol data unit (PDU), on a loop detection shutdown-disabled interface, the interface will never be put into an error-disabled state, but it will generate a periodic log message indicating that the interface is in the shutdown-disabled mode. The periodic syslog is by default generated at an interval of five minutes. You can change this interval as required.

You can globally specify the interval at which the loop-detection syslog message is generated if the **loop detection shutdown-disable** command is configured on the port. This configuration applies to all the ports that have shutdown prevention for loop detection configured.

During a log interval duration window, a log message will be displayed for the first loop detection PDU received on the interface. This means that there will be only one log message per port in an interval window.

To configure the periodic log message generation for shutdown prevention, follow these steps.

1. Enter global configuration mode.
2. Enter the **loop-detection syslog-interval <num>** command.

The following command will set the syslog-interval to 1 hr.

```
Brocade (config)# loop-detection-syslog-interval 60
```

Syslog for port shutdown prevention

Describes the syslog for port shutdown prevention.

<14>0d01h38m44s:<product type>: port <port-num> detect loop, ignoring shut down event in shutdown-disable mode.

CLI banner configuration

Brocade devices 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 terminal when a Telnet CLI session is established.

For example, to display the message "Welcome to FESX!" when a Telnet CLI session is established.

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

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 4000 characters long, which can consist of multiple lines.

Syntax: `[no] banner motd delimiting-character`

To remove the banner, enter the `no banner motd` command.

NOTE

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

When you access the Web Management Interface, the banner is displayed.



Click the [\[Login\]](#) link to accept and continue the login process..

[\[Login\]](#)

NOTE

If you are using a Web client to view the message of the day, and your banners are very wide, with large borders, you may need to set your PC display resolution to a number greater than the width of your banner. For example, if your banner is 100 characters wide and the display is set to 80 characters, the banner may distort, or wrap, and be difficult to read. If you set your display resolution to 120 characters, the banner will display correctly.

Requiring users to press the Enter key after the message of the day banner

In earlier IronWare software releases, users were required to press the Enter key after the Message of the Day (MOTD) was displayed, prior to logging in to the Brocade device on a console or from a Telnet session.

Now, this requirement is disabled by default. Unless configured, users do not have to press Enter after the MOTD banner is displayed.

For example, if the MOTD "Authorized Access Only" is configured, by default, the following messages are displayed when a user tries to access the Brocade device from a Telnet session.

```
Authorized Access Only ...
Username:
```

The user can then login to the device.

However, if the requirement to press the **Enter key** is enabled, the following messages are displayed when accessing the switch from Telnet.

```
Authorized Access Only ...
Press <Enter> to accept and continue the login process....
```

The user must press the **Enter key** before the login prompt is displayed.

Also, on the console, the following messages are displayed if the requirement to press the **Enter key** is disabled.

```
Press Enter key to login
Authorized Access Only ...
User Access Verification
Please Enter Login Name:
```

However, if the requirement to press the **Enter key** after a MOTD is enabled, the following messages are displayed when accessing the switch on the console.

```
Press Enter key to login
Authorized Access Only ...
Press <Enter> to accept and continue the login process....
```

The user must press the **Enter key** to continue to the login prompt.

To enable the requirement to press the **Enter key** after the MOTD is displayed, enter a command such as the following.

```
Brocade(config)# banner motd require-enter-key
```

Syntax: [no] banner motd require-enter-key

Use the **no** form of the command to disable the requirement.

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

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

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). The delimiting character can be any character except " (double-quotation mark) and cannot appear in the banner text. The text in between the pound signs is the contents of the banner. Banner text can be up to 4000 characters, which can consist of multiple lines.

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

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

Displaying a console message 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.

```
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 209.157.22.63
Incoming Telnet Session!!
```

As with the **banner motd** command, you begin and end the message with a delimiting character; in this example, the delimiting character is \$(dollar sign). The delimiting character can be any character except " (double-quotation mark) and cannot appear in the banner text. The text in between the dollar signs is the contents of the banner. Banner text can be up to 4000 characters, which can consist of multiple lines.

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

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

Operations, Administration, and Maintenance

• OAM Overview.....	85
• Software versions installed and running on a device.....	86
• Software Image file types.....	89
• Flash timeout.....	90
• Software upgrades.....	90
• Boot code synchronization feature.....	90
• Viewing the contents of flash files.....	90
• Using SNMP to upgrade software.....	92
• Software reboot.....	93
• Displaying the boot preference.....	94
• Loading and saving configuration files.....	94
• Loading and saving configuration files with IPv6.....	99
• System reload scheduling.....	103
• Diagnostic error codes and remedies for TFTP transfers.....	104
• Network connectivity testing.....	106
• IEEE 802.3ah EFM-OAM.....	108
• Hitless management on the FSX 800 and FSX 1600.....	116
• Displaying management redundancy information	126
• Layer 3 hitless route purge	126
• DHCP Client-Based Auto-Configuration and Flash image update.....	128
• Energy Efficient Ethernet.....	137
• Histogram information overview.....	137
• External USB Hotplug.....	138

OAM Overview

For easy software image management, all Brocade devices support the download and upload of software images between the flash modules on the devices and a Trivial File Transfer Protocol (TFTP) server on the network.

Brocade devices have two flash memory modules:

- Primary flash - The default local storage device for image files and configuration files.
- Secondary flash - A second flash storage device. You can use the secondary flash to store redundant images for additional booting reliability or to preserve one software image while testing another one.

Only one flash device is active at a time. By default, the primary image will become active upon reload.

You can update the software contained on a flash module using TFTP to copy the update image from a TFTP server onto the flash module. In addition, you can copy software images and configuration files from a flash module to a TFTP server.

NOTE

Brocade devices are TFTP clients but not TFTP servers. You must perform the TFTP transaction from the Brocade device. You cannot "put" a file onto the Brocade device using the interface of your TFTP server.

NOTE

If you are attempting to transfer a file using TFTP but have received an error message, refer to [Diagnostic error codes and remedies for TFTP transfers](#) on page 104.

Software versions installed and running on a device

Use the following methods to display the software versions running on the device and the versions installed in flash memory.

Determining the flash image version running on the device

To determine the flash image version running on a device, enter the **show version** command at any level of the CLI. Some examples are shown below.

Compact devices

To determine the flash image version running on a Compact device, enter the **show version** command at any level of the CLI. The following shows an example output.

```
device#show version
Copyright (c) 1996-2012 Brocade Communications Systems, Inc. All rights reserved.
  UNIT 1: compiled on Mar  2 2012 at 12:38:17 labeled as ICX64S07400
          (10360844 bytes) from Primary ICX64S07400.bin
          SW: Version 07.4.00T311
  Boot-Monitor Image size = 774980, Version:07.4.00T310 (kxz07400)
  HW: Stackable ICX6450-24
=====
UNIT 1: SL 1: ICX6450-24 24-port Management Module
      Serial #: BZSxxxxxxxx
      License: BASE_SOFT_PACKAGE (LID: dbuFJJHiFFi)
      P-ENGINE 0: type DEF0, rev 01
=====
UNIT 1: SL 2: ICX6450-SFP-Plus 4port 40G Module
=====
  800 MHz ARM processor ARMv5TE, 400 MHz bus
  65536 KB flash memory
  512 MB DRAM
  STACKID 1 system uptime is 3 minutes 39 seconds
  The system : started=warm start  reloaded=by "reload"
```

The version information is shown in bold type in this example:

- "03.0.00T53" indicates the flash code version number. The "T53" is used by Brocade for record keeping.
- "labeled as FER03000" indicates the flash code image label. The label indicates the image type and version and is especially useful if you change the image file name.
- "Primary fer03000.bin" indicates the flash code image file name that was loaded.

Displaying flash image version on chassis devices

To determine the flash image version running on a chassis device, enter the **show version** command at any level of the CLI. The following is an example output.

```
device#show version
=====
Active Management CPU [Slot-9]:
  SW: Version 07.4.00T3e3 Copyright (c) 1996-2012 Brocade Communications Systems, Inc. All rights reserved.
  Compiled on Mar 02 2012 at 11:54:29 labeled as SXR07400
  (4585331 bytes) Primary /GA/SXR07400.bin
```

```

BootROM: Version 07.2.00T3e5 (FEv2)
Chassis Serial #: Bxxxxxxxxx
License: SX_V6_HW_ROUTER_IPv6_SOFT_PACKAGE (LID: yGFJGOiFLd)
HW: Chassis FastIron SX 800-PREM6 (PROM-TYPE SX-FIL3U-6-IPV6)
=====
Standby Management CPU [Slot-10]:
SW: Version 07.4.00T3e3 Copyright (c) 1996-2012 Brocade Communications Systems, Inc. All rights reserved.
Compiled on Mar 02 2012 at 11:54:29 labeled as SXR07400
BootROM: Version 07.2.00T3e5 (FEv2)
HW: Chassis FastIron SX 800-PREM6 (PROM-TYPE SX-FIL3U-6-IPV6)
=====
SL 1: SX-FI-8XG 8-port 10G Fiber
Serial #: BQKxxxxxxxxx
P-ASIC 0: type C341, rev 00 subrev 00
=====
SL 2: SX-FI-24GPP 24-port Gig Copper + PoE+
Serial #: BTUxxxxxxxxx
P-ASIC 2: type C300, rev 00 subrev 00
=====
SL 8: SX-FI-48GPP 48-port Gig Copper + PoE+
Serial #: BEVxxxxxxxxx
P-ASIC 14: type C300, rev 00 subrev 00
=====
SL 9: SX-FIZMR6 0-port Management
Serial #: Wxxxxxxxxx
License: SX_V6_HW_ROUTER_IPv6_SOFT_PACKAGE (LID: yGFJGOiFLd)
=====
SL 10: SX-FIZMR6 0-port Management
Serial #: Wxxxxxxxxx
License: SX_V6_HW_ROUTER_IPv6_SOFT_PACKAGE (LID: ŷŷŷŷŷŷŷŷŷŷ)
=====
Active Management Module:
660 MHz Power PC processor 8541 (version 0020/0020) 66 MHz bus
512 KB boot flash memory
16384 KB code flash memory
512 MB DRAM
Standby Management Module:
660 MHz Power PC processor 8541 (version 0020/0020) 66 MHz bus
512 KB boot flash memory
16384 KB code flash memory
512 MB DRAM
The system uptime is 1 minutes 2 seconds
The system : started=warm start reloaded=by "reload"

```

The version information is shown in bold type in this example:

- "03.1.00aT3e3" indicates the flash code version number. The "T3e3" is used by Brocade for record keeping.
- "labeled as SXR03100a" indicates the flash code image label. The label indicates the image type and version and is especially useful if you change the image file name.
- "Primary SXR03100a.bin" indicates the flash code image file name that was loaded.

Displaying the boot image version running on the device

To determine the boot image running on a device, enter the **show flash** command at any level of the CLI. The following shows an example output.

```

device#show flash
Active Management Module (Slot 9):
Compressed Pri Code size = 3613675, Version 03.1.00aT3e3 (sxr03100a.bin)
Compressed Sec Code size = 2250218, Version 03.1.00aT3e1 (sxs03100a.bin)
Compressed BootROM Code size = 524288, Version 03.0.01T3e5
Code Flash Free Space = 9699328
Standby Management Module (Slot 10):
Compressed Pri Code size = 3613675, Version 03.1.00aT3e3 (sxr03100a.bin)
Compressed Sec Code size = 2250218, Version 03.1.00aT3e1 (sxs03100a.bin)

```

```
Compressed BootROM Code size = 524288, Version 03.0.01T3e5  
Code Flash Free Space = 524288
```

The boot code version is shown in bold type.

Displaying the image versions installed in flash memory

Enter the **show flash** command to display the boot and flash images installed on the device. An example of the command output is shown in [Displaying the boot image version running on the device](#) on page 87:

- The "Compressed Pri Code size" line lists the flash code version installed in the primary flash area.
- The "Compressed Sec Code size" line lists the flash code version installed in the secondary flash area.
- The "Boot Monitor Image size" line lists the boot code version installed in flash memory. The device does not have separate primary and secondary flash areas for the boot image. The flash memory module contains only one boot image.

NOTE

To minimize the boot-monitor image size on FastIron devices, the **ping** and **tftp** operations performed in the boot-monitor mode are restricted to copper ports on the FastIron Chassis management modules and to the out-of-band management port on the FastIron stackable switches. The other copper or fiber ports on these devices do not have the ability to ping or tftp from the boot-monitor mode.

Flash image verification

The Flash Image Verification feature allows you to verify boot images based on hash codes, and to generate hash codes where needed. This feature lets you select from three data integrity verification algorithms:

- MD5 - Message Digest algorithm (RFC 1321)
- **SHA1** - US Secure Hash Algorithm (RFC 3174)
- CRC - Cyclic Redundancy Checksum algorithm

Flash image CLI commands

Use the following command syntax to verify the flash image:

Syntax: `verify md5 | sha1 | crc32 ASCII string|primary|secondary[hash code]`

- *md5* - Generates a 16-byte hash code
- *sha1* - Generates a 20-byte hash code
- *crc32* - Generates a 4 byte checksum
- *ascii string* - A valid image filename
- *primary* - The primary boot image (primary.img)
- *secondary* - The secondary boot image (secondary.img)
- *hash code* - The hash code to verify

The following examples show how the **verify** command can be used in a variety of circumstances.

To generate an MD5 hash value for the secondary image, enter the following command.

```
device#verify md5 secondary  
device#.....Done  
Size = 2044830, MD5 01c410d6d153189a4a5d36c955653862
```


To generate a SHA-1 hash value for the secondary image, enter the following command.

```
device#verify sha secondary
device#.....Done
Size = 2044830, SHA1 49d12d26552072337f7f5fcaef4cf4b742a9f525
```

To generate a CRC32 hash value for the secondary image, enter the following command.

```
device#verify crc32 secondary
device#.....Done
Size = 2044830, CRC32 b31fcbc0
```

To verify the hash value of a secondary image with a known value, enter the following commands.

```
device#verify md5 secondary 01c410d6d153189a4a5d36c955653861
device#.....Done
Size = 2044830, MD5 01c410d6d153189a4a5d36c955653862
Verification FAILED.
```

In the previous example, the codes did not match, and verification failed. If verification succeeds, the output will look like this.

```
device#verify md5 secondary 01c410d6d153189a4a5d36c955653861
device#.....Done
Size = 2044830, MD5 01c410d6d153189a4a5d36c955653861
Verification SUCCEEDED.
```

The following examples show this process for SHA-1 and CRC32 algorithms.

```
device#verify sha secondary 49d12d26552072337f7f5fcaef4cf4b742a9f525
device#.....Done
Size = 2044830, sha 49d12d26552072337f7f5fcaef4cf4b742a9f525
Verification SUCCEEDED.
```

and

```
device#verify crc32 secondary b31fcbc0
device#.....Done
Size = 2044830, CRC32 b31fcbc0
Verification SUCCEEDED.
```

Software Image file types

This section lists the boot and flash image file types supported and how to install them on the FastIron family of switches. For information about a specific version of code, refer to the release notes.

NOTE

The boot images are applicable to the listed devices only and are not interchangeable. For example, you cannot load FCX boot or flash images on an FSX device, and vice versa.

TABLE 16 Software image files

Product	Boot image	Flash image
FSX 800	sxzxxxxx.bin	SXLSxxxxx.bin (Layer 2) or
FSX 1600		SXLRxxxxx.bin (full Layer 3)
FCX	grzxxxxx.bin	FCXSxxxxx.bin (Layer 2) or FCXRxxxxx.bin (Layer 3)
ICX 6610		
ICX 6430	kxzxxxxx.bin	ICX64Sxxxxx.bin (Layer 2) or
ICX 6450		ICX64Rxxxxx.bin (Layer 3 - ICX 6450 only)
ICX 6650	fxzxxxxx.bin	ICXRxxxxx.bin

TABLE 16 Software image files (continued)

Product	Boot image	Flash image
		ICXSxxxx.bin
ICX 7250	spzxxxx.bin	SPSxxxx.bin (Layer 2) or
ICX 7450		SPRxxxx.bin (Layer 3)
ICX 7750	swzxxxx.bin	SWSxxxx.bin (Layer 2) or
		SWRxxxx.bin (Layer 3)

Flash timeout

The operations that require access to the flash device are expected to be completed within the default flash timeout value of 12 minutes.

If the operations exceed the timeout value, the flash device will be locked and further flash operations cannot be processed. To facilitate prolonged flash operations without the device being locked, you can manually configure the flash timeout for a longer duration using the **flash-timeout** command. You can configure the flash timeout to a value from 12 through 60 minutes. The new timeout value is applicable for all flash operations and will be effective from the next flash operation.

Software upgrades

For instructions about upgrading the software, refer to the *FastIron Ethernet Switch Software Upgrade Guide*.

Boot code synchronization feature

The Brocade device supports automatic synchronization of the boot image in the active and redundant management modules. When the new boot image is copied into the active module, it is automatically synchronized with the redundant management module.

NOTE

There is currently no option for manual synchronization of the boot image.

To activate the boot synchronization process, enter the following command.

```
device#copy tftp flash 10.20.65.194 /GA/SXZ07200.bin bootrom
```

The system responds with the following message.

```
device#Load to buffer (8192 bytes per dot)
.....Write to boot flash.....
TFTP to Flash Done.
device#Synchronizing with standby module...
Boot image synchronization done.
```

Viewing the contents of flash files

The **copy flash console** command can be used to display the contents of a configuration file, backup file, or renamed file stored in flash memory. The file contents are displayed on the console when the command is entered at the CLI.

To display a list of files stored in flash memory, do one of the following:

- For devices other than FCX and ICX, enter the **dir** command at the monitor mode. To enter monitor mode from any level of the CLI, press the **Shift** and **Control+Y** keys simultaneously then press the **M** key. Enter the **dir** command to display a list of the files stored in flash memory. To exit monitor mode and return to the CLI, press **Control+Z**.
- For FCX devices, enter the **show dir** command at any level of the CLI, or enter the **dir** command at the monitor mode.
- For ICX devices, enter the **show files** command at the device configuration prompt.

The following shows an example command output.

```
device#show dir
133 [38f4] boot-parameter
    0 [ffff] bootrom
3802772 [0000] primary
4867691 [0000] secondary
    163 [dd8e] stacking.boot
    1773 [0d2d] startup-config
    1808 [acfa] startup-config.backup
8674340 bytes 7 File(s)
56492032 bytes free
```

Syntax: show dir

The following example shows the output of the **show files** command.

```
device#show files
Type      Size      Name
-----
F          28203908 primary
F          27949956 secondary
F           641 startup-config.txt
F           391 stacking.boot
F          76942 debug.boot
F           638 startup-config.backup
F           0 startup-config.no

56232476 bytes 7 File(s) in FI root

1771020288 bytes free in FI root
1771020288 bytes free in /
```

Syntax: show files

To display the contents of a flash configuration file, enter a command such as the following from the User EXEC or Privileged EXEC mode of the CLI:

```
device#copy flash console startup-config.backup
ver 07.0.00b1T7f1 !
stack unit 1
  module 1 fcx-24-port-management-module
  module 2 fcx-cx4-2-port-16g-module
  module 3 fcx-xfp-2-port-10g-module
  priority 80
  stack-port 1/2/1 1/2/2
stack unit 2
  module 1 fcx-48-poe-port-management-module
  module 2 fcx-cx4-2-port-16g-module
  module 3 fcx-xfp-2-port-10g-module
  stack-port 2/2/1 2/2/2
stack enable
!
!
!
!
vlan 1 name DEFAULT-VLAN by port
  no spanning-tree
  metro-rings 1
  metro-ring 1
```

```

master
ring-interfaces ethernet 1/1/2 ethernet 1/1/3
enable
!
vlan 10 by port
mac-vlan-permit ethe 1/1/5 to 1/1/6 ethe 2/1/5 to 2/1/6 no spanning-tree !
vlan 20 by port
untagged ethe 1/1/7 to 1/1/8
no spanning-tree
pvlan type primary
pvlan mapping 40 ethe 1/1/8
pvlan mapping 30 ethe 1/1/7
!
vlan 30 by port
untagged ethe 1/1/9 to 1/1/10
no spanning-tree
pvlan type community
!
...
some lines omitted for brevity...

```

Syntax: `copy flash console filename`

For *filename*, enter the name of a file stored in flash memory.

Using SNMP to upgrade software

You can use a third-party SNMP management application such as HP OpenView to upgrade software on a Brocade device.

NOTE

The syntax shown in this section assumes that you have installed HP OpenView in the "/usr" directory.

NOTE

Brocade recommends that you make a backup copy of the startup-config file before you upgrade the software. If you need to run an older release, you will need to use the backup copy of the startup-config file.

1. Configure a read-write community string on the Brocade device, if one is not already configured. To configure a read-write community string, enter the following command from the global CONFIG level of the CLI: `snmp-server community string ro | rw` where *string* is the community string and can be up to 32 characters long.
2. On the Brocade device, enter the following command from the global CONFIG level of the CLI.

no snmp-server pw-check

This command disables password checking for SNMP set requests. If a third-party SNMP management application does not add a password to the password field when it sends SNMP set requests to a Brocade device, by default the Brocade device rejects the request.

- From the command prompt in the UNIX shell, enter the following command.

```
/usr/OV/bin/snmpset -c rw-community-string brcd-ip-addr 1.3.6.1.4.1.1991.1.1.2.1.5.0 ipaddress tftp-ip-addr 1.3.6.1.4.1.1991.1.1.2.1.6.0 octetstringascii file-name 1.3.6.1.4.1.1991.1.1.2.1.7.0 integer command-integer
```

where

rw-community-string is a read-write community string configured on the Brocade device.

brcd-ip-addr is the IP address of the Brocade device.

tftp-ip-addr is the TFTP server IP address.

file-name is the image file name.

command-integer is one of the following.

20 - Download the flash code into the primary flash area.

22 - Download the flash code into the secondary flash area.

Software reboot

You can use boot commands to immediately initiate software boots from a software image stored in primary or secondary flash on a Brocade device or from a BootP or TFTP server. You can test new versions of code on a Brocade device or choose the preferred boot source from the console boot prompt without requiring a system reset.

NOTE

It is very important that you verify a successful TFTP transfer of the boot code before you reset the system. If the boot code is not transferred successfully but you try to reset the system, the system will not have the boot code with which to successfully boot.

By default, the Brocade device first attempts to boot from the image stored in its primary flash, then its secondary flash, and then from a TFTP server. You can modify this booting sequence at the global CONFIG level of the CLI using the **boot system** command.

NOTE

FSX device with FastIron 08.0.00a, ICX 6430, and ICX 6450 devices support only one configured system boot preference.

To initiate an immediate boot from the CLI, enter one of the **boot system** commands.

NOTE

When using the **boot system tftp** command, the IP address of the device and the TFTP server should be in the same subnet.

Software boot configuration notes

- In FastIron X Series devices, the **boot system tftp** command is supported on ports e 1 through e 12 only.
- If you are booting the device from a TFTP server through a fiber connection, use the following command: **boot system tftp ip-address filename fiber-port** .
- The **boot system tftp** command is not supported in a stacking environment.

Displaying the boot preference

Use the **show boot-preference** command to display the boot sequence in the startup config and running config files. The boot sequence displayed is also identified as either user-configured or the default.

The following example shows the default boot sequence preference.

```
device#show boot-preference
Boot system preference (Configured):
  Use Default
Boot system preference(Default):
  Boot system flash primary
  Boot system flash secondary
```

The following example shows a user-configured boot sequence preference.

```
Brocade#show boot-preference
Boot system preference(Configured):
  Boot system tftp 10.1.1.1 FCXR08000.bin
  Boot system flash primary
Boot system preference(Default):
  Boot system flash primary
  Boot system flash secondary
```

Syntax: show boot-preference

The results of the **show run** command for the configured example above appear as follows.

```
Brocade#show run
Current configuration:
!
ver 08.0.00T7f3
!
stack unit 1
  module 1 fcx-24-poe-port-management-module
  module 2 fcx-cx4-2-port-16g-module
  priority 128
  stack-port 1/2/1 1/2/2
stack unit 2
  module 1 fcx-48-port-management-module
  module 2 fcx-cx4-2-port-16g-module
  stack-port 2/2/1 2/2/2
stack enable
stack mac 748e.f80e.dcc0
!
boot sys tf 10.1.1.1 FCXR08000.bin
boot sys fl pri
ip route 0.0.0.0/0 10.37.234.129
!
end
```

Loading and saving configuration files

For easy configuration management, all Brocade devices support both the download and upload of configuration files between the devices and a TFTP server on the network.

You can upload either the startup configuration file or the running configuration file to the TFTP server for backup and use in booting the system:

- Startup configuration file - This file contains the configuration information that is currently saved in flash. To display this file, enter the **show configuration** command at any CLI prompt.

- Running configuration file - This file contains the configuration active in the system RAM but not yet saved to flash. These changes could represent a short-term requirement or general configuration change. To display this file, enter the **show running-config** or **write terminal** command at any CLI prompt.

Each device can have one startup configuration file and one running configuration file. The startup configuration file is shared by both flash modules. The running configuration file resides in DRAM.

When you load the startup-config file, the CLI parses the file three times.

1. During the first pass, the parser searches for **system-max** commands. A **system-max** command changes the size of statically configured memory.
2. During the second pass, the parser implements the **system-max** commands if present and also implements trunk configuration commands (**trunk** command) if present.
3. During the third pass, the parser implements the remaining commands.

Replacing the startup configuration with the running configuration

After you make configuration changes to the active system, you can save those changes by writing them to flash memory. When you write configuration changes to flash memory, you replace the startup configuration with the running configuration.

To replace the startup configuration with the running configuration, enter the following command at any Enable or CONFIG command prompt.

```
device# write memory
```

NOTE

To return the unit to the default startup configuration, use the **delete startup-config** command.

Replacing the running configuration with the startup configuration

If you want to back out of the changes you have made to the running configuration and return to the startup configuration, enter the following command at the Privileged EXEC level of the CLI.

```
device# reload
```

Logging changes to the startup-config file

You can configure a Brocade device to generate a Syslog message when the startup-config file is changed. The trap is enabled by default.

The following Syslog message is generated when the startup-config file is changed.

```
startup-config was changed
```

If the startup-config file was modified by a valid user, the following Syslog message is generated.

```
startup-config was changed by
username
```

To disable or re-enable Syslog messages when the startup-config file is changed, use the following command.

Syntax:[no] logging enable config-changed

Copying a configuration file to or from a TFTP server

To copy the startup-config or running-config file to or from a TFTP server, use the following method.

NOTE

For details about the **copy** command used with IPv6, refer to [Using the IPv6 copy command](#) on page 99.

NOTE

You can name the configuration file when you copy it to a TFTP server. However, when you copy a configuration file from the server to a Brocade device, the file is always copied as "startup-config" or "running-config", depending on which type of file you saved to the server.

To initiate transfers of configuration files to or from a TFTP server using the CLI, enter one of the following commands:

- **copy startup-config tftp tftp-ip-addr filename** - Use this command to upload a copy of the startup configuration file from the Layer 2 Switch or Layer 3 Switch to a TFTP server.
- **copy running-config tftp tftp-ip-addr filename** - Use this command to upload a copy of the running configuration file from the Layer 2 Switch or Layer 3 Switch to a TFTP server.
- **copy tftp startup-config tftp-ip-addr filename** - Use this command to download a copy of the startup configuration file from a TFTP server to a Layer 2 Switch or Layer 3 Switch.

NOTE

It is recommended to use a script or the **copy running-config tftp** command for extensive configuration. You should not copy-paste configuration with more than 2000 characters into CLI.

Dynamic configuration loading

You can load dynamic configuration commands (commands that do not require a reload to take effect) from a file on a TFTP server into the running-config on the Brocade device. You can make configuration changes off-line, then load the changes directly into the device running-config, without reloading the software.

Dynamic configuration usage considerations

- Use this feature only to load configuration information that does not require a software reload to take effect. For example, you cannot use this feature to change statically configured memory (**system-max** command) or to enter trunk group configuration information into the running-config.
- Do not use this feature if you have deleted a trunk group but have not yet placed the changes into effect by saving the configuration and then reloading. When you delete a trunk group, the command to configure the trunk group is removed from the device running-config, but the trunk group remains active. To finish deleting a trunk group, save the configuration (to the startup-config file), then reload the software. After you reload the software, then you can load the configuration from the file.
- Do not load port configuration information for secondary ports in a trunk group. Since all ports in a trunk group use the port configuration settings of the primary port in the group, the software cannot implement the changes to the secondary port.

Preparing the configuration file

A configuration file that you create must follow the same syntax rules as the startup-config file the device creates.

- The configuration file is a script containing CLI configuration commands. The CLI reacts to each command entered from the file in the same way the CLI reacts to the command if you enter it. For example, if the command results in an error message or a change to the CLI configuration level, the software responds by displaying the message or changing the CLI level.

- The software retains the running-config that is currently on the device, and changes the running-config only by adding new commands from the configuration file. If the running config already contains a command that is also in the configuration file you are loading, the CLI rejects the new command as a duplicate and displays an error message. For example, if the running-config already contains a command that configures ACL 1, the software rejects ACL 1 in the configuration file, and displays a message that ACL 1 is already configured.
- The file can contain global CONFIG commands or configuration commands for interfaces, routing protocols, and so on. You cannot enter User EXEC or Privileged EXEC commands.
- The default CLI configuration level in a configuration file is the global CONFIG level. Thus, the first command in the file must be a global CONFIG command or "!". The ! (exclamation point) character means "return to the global CONFIG level".

NOTE

You can enter text following "!" as a comment. However, the "!" is not a comment marker. It returns the CLI to the global configuration level.

NOTE

If you copy-and-paste a configuration into a management session, the CLI ignores the "!" instead of changing the CLI to the global CONFIG level. As a result, you might get different results if you copy-and-paste a configuration instead of loading the configuration using TFTP.

- Make sure you enter each command at the correct CLI level. Since some commands have identical forms at both the global CONFIG level and individual configuration levels, if the CLI response to the configuration file results in the CLI entering a configuration level you did not intend, then you can get unexpected results.

For example, if a trunk group is active on the device, and the configuration file contains a command to disable STP on one of the secondary ports in the trunk group, the CLI rejects the commands to enter the interface configuration level for the port and moves on to the next command in the file you are loading. If the next command is a spanning-tree command whose syntax is valid at the global CONFIG level as well as the interface configuration level, then the software applies the command globally. Here is an example.

The configuration file contains these commands.

```
interface ethernet
 2
no spanning-tree
```

The CLI responds like this.

```
device(config)#interface ethernet 2
Error - cannot configure secondary ports of a trunk
device(config)#no spanning-tree
device(config)#
```

- If the file contains commands that must be entered in a specific order, the commands must appear in the file in the required order. For example, if you want to use the file to replace an IP address on an interface, you must first remove the old address using "no" in front of the **ip address** command, then add the new address. Otherwise, the CLI displays an error message and does not implement the command. Here is an example.

The configuration file contains these commands.

```
interface ethernet 11
ip address 10.10.10.69/24
```

The running-config already has a command to add an address to port 11, so the CLI responds like this.

```
device(config)#interface ethernet 11
device(config-if-e1000-11)#ip add 10.10.10.69/24
Error: can only assign one primary ip address per subnet
device(config-if-e1000-11)#
```

To successfully replace the address, enter commands into the file as follows.

```
interface ethernet
 11
no ip address 10.20.20.69/24
ip address 10.10.10.69/24
```

This time, the CLI accepts the command, and no error message is displayed.

```
device(config)#interface ethernet 11
device(config-if-e1000-11)#no ip add 10.20.20.69/24
device(config-if-e1000-11)#ip add 10.10.10.69/24
device(config-if-e1000-11)
```

- Always use the **end** command at the end of the file. The **end** command must appear on the last line of the file, by itself.

Loading the configuration information into the running-config

To load the file from a TFTP server, use the following command:

```
copy tftp running-config ip-addr filename
```

NOTE

In the current FastIron release, the **copy tftp running-config** command merges only the access-lists and mac-filters configuration from the configuration file on the TFTP server to the running configuration on the device.

NOTE

If you are loading a configuration file that uses a truncated form of the CLI command **access-list**, the software will not go into batch mode.

For example, the following command line *will initiate* batch mode.

```
access-list 131 permit host pc1 host pc2
```

The following command line *will not* initiate batch mode.

```
acc 131 permit host pc1 host pc2
```

Maximum file sizes for startup-config file and running-config

Each Brocade device has a maximum allowable size for the running-config and the startup-config file. If you use TFTP to load additional information into a device running-config or startup-config file, it is possible to exceed the maximum allowable size. If this occurs, you will not be able to save the configuration changes.

The maximum size for the running-config and the startup-config file is 640K each.

To determine the size of a running-config or startup-config file, copy it to a TFTP server, then use the directory services on the server to list the size of the copied file. To copy the running-config or startup-config file to a TFTP server, use the following commands:

- Command to copy the running-config to a TFTP server:
 - **copy running-config tftp ip-addr filename**
- Command to copy the startup-config file to a TFTP server:
 - **copy startup-config tftp ip-addr filename**

Loading and saving configuration files with IPv6

This section describes the IPv6 **copy** command.

Using the IPv6 copy command

The **copy** command for IPv6 allows you to do the following:

- Copy a file from a specified source to an IPv6 TFTP server
- Copy a file from an IPv6 TFTP server to a specified destination

Copying a file to an IPv6 TFTP server

You can copy a file from the following sources to an IPv6 TFTP server:

- Flash memory
- Running configuration
- Startup configuration

Copying a file from flash memory

For example, to copy the primary or secondary boot image from the device flash memory to an IPv6 TFTP server, enter a command such as the following.

```
device#copy flash tftp 2001:DB8:e0ff:7837::3 test.img secondary
```

This command copies the secondary boot image named test.img from flash memory to a TFTP server with the IPv6 address of 2001:DB8:e0ff:7837::3.

Syntax: `copy flash tftp ipv6-address source-file-name primary | secondary`

The *ipv6-address* parameter specifies the address of the TFTP server. You must specify this address in hexadecimal using 16-bit values between colons as documented in RFC 2373.

The *source-file-name* parameter specifies the name of the file you want to copy to the IPv6 TFTP server.

The **primary** keyword specifies the primary boot image, while the **secondary** keyword specifies the secondary boot image.

Copying a file from the running or startup configuration

For example, to copy the running configuration to an IPv6 TFTP server, enter a command such as the following.

```
device#copy running-config tftp 2001:DB8:e0ff:7837::3 newrun.cfg
```

This command copies the running configuration to a TFTP server with the IPv6 address of 2001:DB8:e0ff:7837::3 and names the file on the TFTP server newrun.cfg.

Syntax: `copy running-config | startup-config tftp ipv6-address destination-file-name`

Specify the **running-config** keyword to copy the running configuration file to the specified IPv6 TFTP server.

Specify the **startup-config** keyword to copy the startup configuration file to the specified IPv6 TFTP server.

The tftp *ipv6-address* parameter specifies the address of the TFTP server. You must specify this address in hexadecimal using 16-bit values between colons as documented in RFC 2373.

The *destination-file-name* parameter specifies the name of the file that is copied to the IPv6 TFTP server.

Copying a file from an IPv6 TFTP server

You can copy a file from an IPv6 TFTP server to the following destinations:

- Flash memory
- Running configuration
- Startup configuration

Copying a file to flash memory

For example, to copy a boot image from an IPv6 TFTP server to the primary or secondary storage location in the device flash memory, enter a command such as the following.

```
device#copy tftp flash 2001:DB8:e0ff:7837::3 test.img secondary
```

This command copies a boot image named test.img from an IPv6 TFTP server with the IPv6 address of 2001:DB8:e0ff:7837::3 to the secondary storage location in the device flash memory.

Syntax: `copy tftp flash ipv6-address source-file-name primary | secondary`

The *ipv6-address* parameter specifies the address of the TFTP server. You must specify this address in hexadecimal using 16-bit values between colons as documented in RFC 2373.

The *source-file-name* parameter specifies the name of the file you want to copy from the IPv6 TFTP server.

The **primary** keyword specifies the primary storage location in the device flash memory, while the **secondary** keyword specifies the secondary storage location in the device flash memory.

Copying a file to the running or startup configuration

For example, to copy a configuration file from an IPv6 TFTP server to the running or startup configuration, enter a command such as the following.

```
device#copy tftp running-config 2001:DB8:e0ff:7837::3 newrun.cfg overwrite
```

This command copies the newrun.cfg file from the IPv6 TFTP server and overwrites the running configuration file with the contents of newrun.cfg.

NOTE

To activate this configuration, you must reload (reset) the device.

Syntax: `copy tftp running-config | startup-config ipv6-address source-file-name [overwrite]`

Specify the **running-config** keyword to copy the running configuration from the specified IPv6 TFTP server.

The *ipv6-address* parameter specifies the address of the TFTP server. You must specify this address in hexadecimal using 16-bit values between colons as documented in RFC 2373.

The *source-file-name* parameter specifies the name of the file that is copied from the IPv6 TFTP server.

The **overwrite** keyword specifies that the device should overwrite the current configuration file with the copied file. If you do not specify this parameter, the device copies the file into the current running or startup configuration but does not overwrite the current configuration.

IPv6 copy command

The **copy** command for IPv6 allows you to do the following:

- Copy a primary or secondary boot image from flash memory to an IPv6 TFTP server.
- Copy the running configuration to an IPv6 TFTP server.
- Copy the startup configuration to an IPv6 TFTP server
- Upload various files from an IPv6 TFTP server.

Copying a primary or secondary boot Image from flash memory to an IPv6 TFTP server

For example, to copy the primary or secondary boot image from the device flash memory to an IPv6 TFTP server, enter a command such as the following.

```
device#copy flash primary tftp 2001:DB8:e0ff:7837::3 primary.img
```

This command copies the primary boot image named primary.img from flash memory to a TFTP server with the IPv6 address of 2001:DB8:e0ff:7837::3.

Syntax: **copy flash primary** | **secondary tftp** *ipv6-address source-file-name*

The **primary** keyword specifies the primary boot image, while the **secondary** keyword specifies the secondary boot image.

The *tftp ipv6-address* parameter specifies the address of the TFTP server. You must specify this address in hexadecimal using 16-bit values between colons as documented in RFC 2373.

The *source-file-name* parameter specifies the name of the file you want to copy from flash memory.

Copying the running or startup configuration to an IPv6 TFTP server

For example, to copy a device running or startup configuration to an IPv6 TFTP server, enter a command such as the following.

```
device#copy running-config tftp 2001:DB8:e0ff:7837::3 bakrun.cfg
```

This command copies a device running configuration to a TFTP server with the IPv6 address of 2001:DB8:e0ff:7837::3 and names the destination file bakrun.cfg.

Syntax: **copy running-config** | **startup-config tftp** *ipv6-address destination-file-name*

Specify the **running-config** keyword to copy the device running configuration or the **startup-config** keyword to copy the device startup configuration.

The *tftp ipv6-address* parameter specifies the address of the TFTP server. You must specify this address in hexadecimal using 16-bit values between colons as documented in RFC 2373.

The *destination-file-name* parameter specifies the name of the running configuration that is copied to the IPv6 TFTP server.

IPv6 TFTP server file upload

You can upload the following files from an IPv6 TFTP server:

- Primary boot image.
- Secondary boot image.
- Running configuration.
- Startup configuration.

Uploading a primary or secondary boot image from an IPv6 TFTP server

For example, to upload a primary or secondary boot image from an IPv6 TFTP server to a device flash memory, enter a command such as the following.

```
device#copy tftp 2001:DB8:e0ff:7837::3 primary.img flash primary
```

This command uploads the primary boot image named `primary.img` from a TFTP server with the IPv6 address of `2001:DB8:e0ff:7837::3` to the device primary storage location in flash memory.

Syntax: `copy tftp ipv6-address source-file-name flash primary | secondary`

The `tftp ipv6-address` parameter specifies the address of the TFTP server. You must specify this address in hexadecimal using 16-bit values between colons as documented in RFC 2373.

The `source-file-name` parameter specifies the name of the file you want to copy from the TFTP server.

The **primary** keyword specifies the primary location in flash memory, while the **secondary** keyword specifies the secondary location in flash memory.

Uploading a running or startup configuration from an IPv6 TFTP server

For example to upload a running or startup configuration from an IPv6 TFTP server to a device, enter a command such as the following.

```
device#copy tftp 2001:DB8:e0ff:7837::3 newrun.cfg running-config
```

This command uploads a file named `newrun.cfg` from a TFTP server with the IPv6 address of `2001:DB8:e0ff:7837::3` to the device.

Syntax: `copy tftp ipv6-address source-file-name running-config|startup-config`

The `tftp ipv6-address` parameter specifies the address of the TFTP server. You must specify this address in hexadecimal using 16-bit values between colons as documented in RFC 2373.

The `source-file-name` parameter specifies the name of the file you want to copy from the TFTP server.

Specify the **running-config** keyword to upload the specified file from the IPv6 TFTP server to the device. The device copies the specified file into the current running configuration but does not overwrite the current configuration.

Specify the **startup-config** keyword to upload the specified file from the IPv6 TFTP server to the device. The the device copies the specified file into the current startup configuration but does not overwrite the current configuration.

Using SNMP to save and load configuration information

You can use a third-party SNMP management application such as HP OpenView to save and load a configuration on a Brocade device. To save and load configuration information using HP OpenView, use the following procedure.

NOTE

The syntax shown in this section assumes that you have installed HP OpenView in the `/usr` directory.

1. Configure a read-write community string on the Brocade device, if one is not already configured. To configure a read-write community string, enter the following command from the global CONFIG level of the CLI.

```
snmp-server community string ro|rw
```

where *string* is the community string and can be up to 32 characters long.

- On the Brocade device, enter the following command from the global CONFIG level of the CLI.

```
no snmp-server pw-check
```

This command disables password checking for SNMP set requests. If a third-party SNMP management application does not add a password to the password field when it sends SNMP set requests to a Brocade device, by default the Brocade device rejects the request.

- From the command prompt in the UNIX shell, enter the following command.

```
/usr/OV/bin/snmpset -c rw-community-string device-ip-addr
```

```
1.3.6.1.4.1.1991.1.1.2.1.5.0
```

```
a tftp-ip-addr 1.3.6.1.4.1.1991.1.1.2.1.8.0 s config-file-name
```

```
1.3.6.1.4.1.1991.1.1.2.1.9.0 integer command-integer
```

where

rw-community-string is a read-write community string configured on the Brocade device.

dry-ip-addr is the IP address of the Brocade device.

tftp-ip-addr is the TFTP server IP address.

config-file-name is the configuration file name.

command-integer is one of the following:

20 - Upload the startup-config file from the flash memory of the Brocade device to the TFTP server.

21 - Download a startup-config file from a TFTP server to the flash memory of the Brocade device.

22 - Upload the running-config from the flash memory of the Brocade device to the TFTP server.

23 - Download a configuration file from a TFTP server into the running-config of the Brocade device.

NOTE

Option **23** adds configuration information to the running-config on the device, and does not replace commands. If you want to replace configuration information in the device, use "no" forms of the configuration commands to remove the configuration information, then use configuration commands to create the configuration information you want.

Follow the guidelines in [Dynamic configuration loading](#) on page 96.

Erasing image and configuration files

To erase software images or configuration files, use the commands described below. These commands are valid at the Privileged EXEC level of the CLI:

- erase flash primary** erases the image stored in primary flash of the system.
- erase flash secondary** erases the image stored in secondary flash of the system.
- erase startup-config** erases the configuration stored in the startup configuration file; however, the running configuration remains intact until system reboot.

System reload scheduling

In addition to reloading the system manually, you can configure the Brocade device to reload itself at a specific time or after a specific amount of time has passed.

NOTE

The scheduled reload feature requires the system clock. Refer to [Network Time Protocol Version 4 \(NTPv4\)](#) on page 32.

Reloading at a specific time

To schedule a system reload for a specific time, use the **reload at** command. For example, to schedule a system reload from the primary flash module for 6:00:00 AM, April 1, 2003, enter the following command at the global CONFIG level of the CLI.

```
device#reload at 06:00:00 04-01-03
```

Syntax: **reload at** *hh:mm:ss mm-dd-yy* [**primary** | **secondary**]

hh:mm:ss is the hours, minutes, and seconds.

mm-dd-yy is the month, day, and year.

primary / *secondary* specifies whether the reload is to occur from the primary code flash module or the secondary code flash module. The default is **primary**.

Reloading after a specific amount of time

To schedule a system reload to occur after a specific amount of time has passed on the system clock, use **reload after** command. For example, to schedule a system reload from the secondary flash one day and 12 hours later, enter the following command at the global CONFIG level of the CLI.

```
device#reload after 01:12:00 secondary
```

Syntax: **reload after** *dd:hh:mm* [**primary** | **secondary**]

dd:hh:mm is the number of days, hours, and minutes.

primary / *secondary* specifies whether the reload is to occur from the primary code flash module or the secondary code flash module.

Displaying the amount of time remaining before a scheduled reload

To display how much time is remaining before a scheduled system reload, enter the following command from any level of the CLI.

```
device#show reload
```

Canceling a scheduled reload

To cancel a scheduled system reload using the CLI, enter the following command at the global CONFIG level of the CLI.

```
device#reload cancel
```

Diagnostic error codes and remedies for TFTP transfers

This section describes the error messages associated with TFTP transfer of configuration files, software images or flash images to or from a Brocade device.

Error code	Message	Explanation and action
1	Flash read preparation failed.	A flash error occurred during the download. Retry the download. If it fails again, contact customer support.
2	Flash read failed.	
3	Flash write preparation failed.	
4	Flash write failed.	
5	TFTP session timeout.	TFTP failed because of a time out. Check IP connectivity and make sure the TFTP server is running.
6	TFTP out of buffer space.	The file is larger than the amount of room on the device or TFTP server. If you are copying an image file to flash, first copy the other image to your TFTP server, then delete it from flash. (Use the erase flash ... CLI command at the Privileged EXEC level to erase the image in the flash.) If you are copying a configuration file to flash, edit the file to remove unnecessary information, then try again.
7	TFTP busy, only one TFTP session can be active.	Another TFTP transfer is active on another CLI session, or Web management session, or network management system. Wait, then retry the transfer.
8	File type check failed.	You accidentally attempted to copy the incorrect image code into the system. For example, you might have tried to copy a Chassis image into a Compact device. Retry the transfer using the correct image.
16	TFTP remote - general error.	The TFTP configuration has an error. The specific error message describes the error. Correct the error, then retry the transfer.
17	TFTP remote - no such file.	
18	TFTP remote - access violation.	
19	TFTP remote - disk full.	
20	TFTP remote - illegal operation.	
21	TFTP remote - unknown transfer ID.	
22	TFTP remote - file already exists.	
23	TFTP remote - no such user.	

This section describes the error messages associated with the TFTP transfer of PoE firmware file to a Brocade device.

Message	Explanation and action
Firmware TFTP timeout.	TFTP failed because of a time out. Check IP connectivity and make sure the TFTP server is running.
Firmware is not valid for this platform.	Each PoE firmware file delivered by Brocade is meant to be used on the specific platform only. If the file is used on a platform for which it is not meant, then this error message will display. Download the correct file, then retry the transfer.
Firmware is not valid for the IEEE 802.3at (PoE-Plus) controller type.	Each PoE firmware file delivered by Brocade is meant to be used on the specific platform only. If the file is used on a platform for which it is not meant, then this error message will display.

Message	Explanation and action
	Download the correct file, then retry the transfer.
Firmware is not valid for the IEEE 802.3af PoE controller type.	Each PoE firmware file delivered by Brocade is meant to be used on the specific platform and the specific PoE controller on the specified module. If the file is used for a platform for which it is meant, but the PoE controller is not same then this error message will display. Download the correct file, then retry the transfer.
Firmware type cannot be detected from the firmware content.	
TFTP File not Valid for PoE Controller Type.	
Firmware tftp remote file access failed.	The TFTP server needs read access on the PoE firmware file. Check the permissions on the file, then try again.

Network connectivity testing

After you install the network cables, you can test network connectivity to other devices by pinging those devices. You also can observe the LEDs related to network connection and perform trace routes.

For more information about observing LEDs, refer to the Brocade FastIron X Series Chassis Hardware Installation Guide and the Brocade FastIron Compact Switch Hardware Installation Guide.

Pinging an IPv4 address

NOTE

This section describes the *IPv4 ping* command. For details about *IPv6 ping*, refer to the *FastIron Ethernet Layer 3 Routing Configuration Guide*.

To verify that a Brocade device can reach another device through the network, enter a command such as the following at any level of the CLI on the Brocade device:

```
device> ping 10.33.4.7
```

Syntax: ping *ip-addr* | *hostname* [**source** *ip-addr*] [**count** *num*] [**timeout** *msec*] [**tll** *num*] [**sizebyte**][**quiet**][**numeric**][**no-fragment**][**verify**][**data1-to-4 byte hex**][**brief**[**max-print-per-sec number**]]

NOTE

If the device is a Brocade Layer 2 Switch or Layer 3 Switch, you can use the host name only if you have already enabled the Domain Name Server (DNS) resolver feature on the device from which you are sending the ping. Refer to "IP Configuration" chapter in the *FastIron Ethernet Switch Layer 3 Routing Configuration Guide*.

The required parameter is the IP address or host name of the device.

The *source ip-addr* 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 Brocade 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 *tll 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 - 10000. The default is 16.

The *no-fragment* parameter turns on the "don't 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 only displays 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 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. Instead, if you do exceed the range for a numeric value, the software rounds the value to the nearest valid value.

The *brief* parameter 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.

NOTE

The number of ! characters displayed may not correspond to the number of successful replies by the **ping** command. Similarly, the number of . characters displayed may not correspond to the number of server timeouts that occurred while waiting for a reply. The "success" or "timeout" results are shown in the display as "Success rate is XX percent (X/Y)".

The optional *max-print-per-sec number* parameter specifies the maximum number of target responses the Brocade device can display per second while in brief mode. You can specify from 0 - 2047. The default is 511.

NOTE

If you address the ping to the IP broadcast address and network address, the device lists the first four responses to the ping.

NOTE

On 48GC modules in non-jumbo mode, the maximum size of ping packets is 1486 bytes and the maximum frame size of tagged traffic is no larger than 1581 bytes.

Tracing an IPv4 route

NOTE

This section describes the **IPv4tracert** command. For details about **IPv6tracert**, refer to the *FastIron Ethernet Switch Layer 3 Routing Configuration Guide*.

Use the **tracert** command to determine the path through which a Brocade device can reach another device. Enter the command at any level of the CLI.

The CLI displays trace route information for each hop as soon as the information is received. Tracert requests display all responses to a given TTL. In addition, if there are multiple equal-cost routes to the destination, the Brocade device displays up to three responses by default.

```
device> tracert 10.33.4.7
```

Syntax: **tracert** *host-ip-addr* [**maxttl** *value*] [**minttl** *value*] [**numeric**] [**timeout** *value*] [**source-ip** *ip-addr*]

Possible and default values are as follows.

minttl value - Minimum TTL (hops) value: Possible values are 1 through 255. Default value is 1 second.

maxttl value - Maximum TTL (hops) value: Possible values are 1 through 255. Default value is 30 seconds.

timeout value - Possible values are 1 through 120. Default value is 2 seconds.

numeric - Changes the display to list the devices by their IP addresses instead of their names.

source-ip ip-addr - Specifies an IP address to be used as the origin for the traceroute.

IEEE 802.3ah EFM-OAM

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.

Using the Ethernet in the First Mile solution, you 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 IEEE 802.3ah EFM 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.

When OAM is present, two connected OAM sub-layers exchange protocol data units (OAMPDUs). OAMPDUs 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 OAMPDU frames from other frames.

OAM functionality is designed to provide reliable service assurance mechanisms for both provider and customer networks.

Network deployment use case

The data-link layer OAM is targeted at last-mile applications, and service providers can use it for demarcation point OAM services.

Ethernet last-mile applications require robust infrastructure that is both passive and active. 802.3ah OAM aims to solve validation and testing problems in such an infrastructure.

Using the Ethernet demarcation, service providers can additionally manage the remote device without utilizing an IP layer. This can be done by using link-layer SNMP counters, request and reply, loopback testing, and other techniques.

EFM-OAM protocol

The functionality of the EFM-OAM can be summarized under the following categories:

- **Discovery:** Discovery is the mechanism to detect the presence of an OAM sub-layer on the remote device. During the discovery process, information about OAM entities, capabilities, and configurations are exchanged.
- **Remote fault detection:** Provides a mechanism for an OAM entity to convey error conditions to its peer by way of a flag in the OAMPDUs.
- **Remote loopback:** This mechanism is used to troubleshoot networks and to isolate problem segments in a large network by sending test segments.

Discovery

Discovery is the first phase of EFM-OAM. At this phase, EFM-OAM identifies network devices along with their OAM capabilities. The Discovery process relies on the Information OAMPDUs. During discovery, the following information is advertised through the TLVs within periodic information OAMPDUs:

- OAM capabilities: Advertises the capabilities of the local OAM entity. Using this information, a peer can determine what functions are supported and accessible (for example, loopback capability).
- OAM mode: The OAM mode is conveyed to the remote OAM entity. The mode can be either active or passive, and can also be used to determine a device's functionality.
- OAMPDU configuration: This configuration includes the maximum OAMPDU size to delivery. In combination with the limited rate of 10 frames per second, this information can be used to limit the bandwidth allocated to OAM traffic.

Timers

Two configurable timers control the protocol, one determining the rate at which OAMPDUs are to be sent, and the second controlling the rate at which OAMPDUs are to be received to maintain the Discovery procedure from resetting.

- The timer should generate PDUs in the range of 1 through 10 PDUs per second. The default value is 1 PDU per second.
- The Hold timer assumes the peer is dead if no packet is received for a period of 1 through 10 seconds. The default value is 5 seconds.

Flags

Included in every OAMPDU is a flags field, which contains, besides other information, the status of the discovery process. There are three possible values for the status:

- Discovering: Discovery is in progress.
- Stable: Discovery is completed. Once aware of this, the remote OAM entity can start sending any type of OAMPDU.
- Unsatisfied: When there are mismatches in the OAM configuration that prevent OAM from completing the discovery, the discovery process is considered unsatisfied and cannot continue.

Process overview

The discovery process allows local Data Terminating Entity (DTE) to detect OAM on a remote DTE. Once OAM support is detected, both ends of the link exchange state and configuration information (such as mode, PDU size, loopback support, and so on). If both DTEs are satisfied with the settings, OAM is enabled on the link. However, the loss of a link or a failure to receive OAMPDUs for five seconds may cause the discovery process the start over again.

DTEs may be in either active or passive mode. Active mode DTEs instigate OAM communications and can issue queries and commands to a remote device. Passive mode DTEs generally wait for the peer device to instigate OAM communications and respond to, but do not instigate, commands and queries. Rules of what DTEs in active or passive mode can do are discussed in the following sections.

Rules for active mode

A DTE in active mode:

- Initiates the OAM Discovery process
- Sends information PDUs
- May send event notification PDUs

- May send variable request or response PDUs
- May send loopback control PDUs

Exceptions

- A DTE in active mode does not respond to variable request PDUs from DTEs in passive mode
- A DTE in active mode does not react to loopback control PDUs from DTEs in passive mode

Rules for passive mode

A DTE in passive mode:

- Waits for the remote device to initiate the Discovery process
- Sends information PDUs
- May send event notification PDUs
- May respond to variable request PDUs
- May react to received loopback control PDUs
- Is not permitted to send variable request or loopback control OAMPDUs

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:

- Dying gasp: This condition is detected when the receiver goes down. The dying gasp condition is considered as unrecoverable. The conditions for a dying gasp condition include:
 - Reload command (Warm reboot)
 - Boot system flash pri/sec command (Warm reboot)
 - Failure on the box (Cold reboot)
- Critical event: On any critical event, the DTE will set the critical event bit in the information OAMPDU. The device will generate critical event in the following cases:
 - When the temperature of the box breaches the warning/shutdown threshold
 - Fan failure

The battleshort mode allows you to prevent the shutdown of ICX 7450 and ICX 7750 when the temperature of the box breaches the warning or shutdown threshold. This is intended to be used in emergency conditions to allow the switches to function in a hostile environment as long as possible.

To enable the battleshort mode, execute the **ignore-temp-shutdown** command from global configuration mode. This command can also be configured at a unit level. By default, the battleshort mode is disabled.

Remote loopback

An OAM entity can put its remote entity into loopback mode using a loopback control OAMPDU. This helps you ensure quality of links during installation or when troubleshooting. In loopback mode, each frame received is transmitted back on that same port except for OAMPDUs and pause frames. The periodic exchange of OAMPDUs must continue while in the loopback state to maintain the OAM session. The loopback command is acknowledged by responding with an information OAMPDU with the loopback state indicated in the state field.

NOTE

Brocade recommends to ensure that any higher layer protocol running over the local and remote loopback ports does not block the interfaces in the VLAN on which loopback traffic testing is being performed.

NOTE

Ethernet loopback and EFM-OAM remote loopback cannot be configured on the same interface.

NOTE

If EEE is enabled globally, port ceases to be in the remote loopback mode.

EFM-OAM error disable recovery

The error disable recovery feature enables the device to recover the EFM-OAM interface from the error-disabled state caused by reception of a critical event from the remote device. Enter the **errdisable recovery cause loam-critical-event** command to enable automatic recovery of ports from error-disabled state.

The ports will recover automatically from the error-disabled state upon the expiry of the error disable recovery timeout value.

Configuring EFM-OAM

The EFM-OAM configuration includes the following procedural steps to enable EFM-OAM on an interface or multiple interfaces for advanced monitoring and maintenance of Ethernet network.

1. Enter the **configure terminal** command to enter global configuration mode.

```
device# configure terminal
```

2. Enter the **link-oam** command to enable the EFM-OAM protocol and enter EFM-OAM protocol configuration mode.

```
device(config)# link-oam
device(config-link-oam)#
```

3. Enter the **timeout** command to configure the time in seconds for which the local Data Terminal Equipment (DTE) waits to receive OAM Protocol Data Units (OAM-PDUs) from the remote entity.

```
device(config-link-oam)# timeout 5
```

4. Enter the **pdu-rate** command to configure the number of PDUs to be transmitted per second by the DTE.

```
device(config-link-oam)# pdu-rate 2
```

- Enter the **ethernet** command to enable EFM-OAM on an interface.

EFM-OAM can be enabled on more than one interface. You can also specify a range of interfaces to enable EFM-OAM on multiple interfaces.

You can set the operational mode of EFM-OAM as Active or Passive.

- Enter the **ethernet stackid/slot/port active** command to set the EFM-OAM operational mode as active on an interface.

```
device(config-link-oam)# ethernet 1/1/3 active
device(config-link-oam)# ethernet 1/1/4 active
```

- Enter the **ethernet stackid/slot/port to stackid/slot/port active** command to set the EFM-OAM operational mode as active on a range of interfaces.

```
device(config-link-oam)# ethernet 1/1/5 to 1/1/8 active
```

- Enter the **ethernet stackid/slot/port passive** command to set the EFM-OAM operational mode as passive on an interface.

```
device(config-link-oam)# ethernet 2/1/1 passive
```

- Enter the **ethernet stackid/slot/port to stackid/slot/port passive** command to set the EFM-OAM operational mode as passive on a range of interfaces.

```
device(config-link-oam)# ethernet 2/1/1 to 2/1/8 passive
```

- (Optional) Enter the **ethernet stackid/slot/port allow-loopback** command to enable the interface to respond to a loopback request from the remote device.

```
device(config-link-oam)# ethernet 1/1/3 allow-loopback
```

- (Optional) Enter the **ethernet stackid/slot/port remote-failure** command to set the device for the remote-failure action to be taken upon the reception of critical event information on the interface.

```
device(config-link-oam)# ethernet 1/1/3 remote-failure critical-event action block-interface
```

- (Optional) Enter the **remote-loopback ethernet stackid/slot/port** command to start or stop the remote loopback procedure on a remote device.

```
device(config-link-oam)# remote-loopback ethernet 2/1/1 start
device(config-link-oam)# remote-loopback ethernet 2/1/1 stop
```

The following output shows an example of EFM-OAM configuration.

```
device# configure terminal
device(config)# link-oam
device(config-link-oam)# timeout 5
device(config-link-oam)# pdu-rate 2
device(config-link-oam)# ethernet 1/1/3 active
device(config-link-oam)# ethernet 1/1/3 allow-loopback
device(config-link-oam)# remote-loopback ethernet 2/1/1 start
device(config-link-oam)# ethernet 1/1/3 remote-failure critical-event action block-interface
```

Displaying OAM information

The following sample output of the **show link-oam info** command displays the OAM information on all OAM-enabled ports.

```
device (config)# show link-oam info
Ethernet Link Status      OAM Status      Mode      Local Stable      Remote Stable
1/1/1      up              up              active      satisfied         satisfied
```


1/1/2	up	up	passive	satisfied	satisfied
1/1/3	up	up	active	satisfied	satisfied
1/1/4	up	init	passive	unsatisfied	unsatisfied
1/1/5	down	down	passive	unsatisfied	unsatisfied
1/1/6	down	down	passive	unsatisfied	unsatisfied
1/1/7	down	down	passive	unsatisfied	unsatisfied

The following sample output of the **show link-oam info detail** command displays detailed OAM information on all OAM-enabled ports.

```
device(config)# show link-oam info detail
OAM information for Ethernet port: 10/1/1
+link-oam mode:      passive
+link status:       down
+oam status:        down
Local information
  multiplexer action: forward
  parse action:      forward
  stable:            unsatisfied
  state:             linkFault
  loopback state:   disabled
  dying-gasp:       false
  critical-event:   false
  link-fault:       true
Remote information
  multiplexer action: forward
  parse action:      forward
  stable:            unsatisfied
  loopback support: disabled
  dying-gasp:       false
  critical-event:   true
  link-fault:       false
```

```
OAM information for Ethernet port: 10/1/3
+link-oam mode:      active
+link status:       up
+oam status:        down
Local information
  multiplexer action: forward
  parse action:      forward
  stable:            unsatisfied
  state:             activeSend
  loopback state:   disabled
  dying-gasp:       false
  critical-event:   false
  link-fault:       false
Remote information
  multiplexer action: forward
  parse action:      forward
  stable:            unsatisfied
  loopback support: disabled
  dying-gasp:       false
  critical-event:   false
  link-fault:       false
```

```
OAM information for Ethernet port: 10/1/4
+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:     true
link-fault:         false

```

The following sample output of the **show link-oam info detail ethernet** command displays detailed OAM information on a specific Ethernet port.

```

device(config)# show link-oam info detail ethernet 1/1/3
OAM information for Ethernet port: 1/1/3
+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

```

Displaying OAM statistics

The following sample output of the **show link-oam statistics** command displays the OAM statistics on all OAM-enabled ports.

```

device(config)# show link-oam statistics
Ethernet Tx Pdus      Rx Pdus
10/1/1  377908        377967
10/1/3   400          44
10/1/4   400          385
10/1/5   400          385
10/1/6   400          385

```

The following sample output of the **show link-oam statistics detail** command displays detailed OAM statistics on all OAM-enabled ports.

```

device(config)# show link-oam statistics detail
OAM statistics for Ethernet port: 10/1/1
Tx statistics
  information OAMPDUs:          377908
  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:          377967
  loopback control OAMPDUs:    0
  loopback control OAMPDUs dropped: 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:       377395
        dying-gasp records:           0
        discarded TLVs:               0
        unrecognized TLVs:            0

OAM statistics for Ethernet port: 10/1/3
  Tx statistics
    information OAMPDUs:              427
    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:              44
    loopback control OAMPDUs:         0
    loopback control OAMPDUs dropped:  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: 10/1/4
  Tx statistics
    information OAMPDUs:              428
    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:              413
    loopback control OAMPDUs:         0
    loopback control OAMPDUs dropped:  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:            350
    dying-gasp records:               0
    discarded TLVs:                   0
    unrecognized TLVs:                0

```

The following sample output of the **show link-oam statistics detail ethernet** command displays detailed OAM statistics on a specific Ethernet port.

```

device(config)# show link-oam statistics detail ethernet 1/1/3
OAM statistics for Ethernet port: 1/1/3
  Tx statistics
    information OAMPDUs:              122474
    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:               94691
loopback control OAMPDUs:          0
loopback control OAMPDUs dropped:  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

```

EFM-OAM syslog messages

When EFM-OAM is enabled on an interface, the syslog messages in the following table are generated when the link goes up or down, or when loopback mode is entered or cleared on an interface.

TABLE 17 EFM-OAM syslog messages

Event	Syslog output
Port 1 is LOAM logically Up	Link-OAM: Logical link on interface Ethernet 1/1/1 is up.
Port 1 is LOAM logically Down	Link-OAM: Logical link on interface Ethernet 1/1/1 is down.
Port 1 entered remote Loopback mode	Link-OAM: Link entered remote loopback on ethernet 1/1/1
Port 1 cleared remote Loopback mode	Link-OAM: Link cleared remtote loopback on ethernet 1/1/1
Port 1 entered local Loopback mode	Link-OAM: Link entered local loopback on ethernet 1/1/1
Port 1 cleared local Loopback mode	Link-OAM: Link cleared local loopback on ethernet 1/1/1
Dying gasp event on port 1	Link-OAM: Link received dying-gasp event on ethernet 1/1/1
Critical event on port 1	Link-OAM: Link received critical event on ethernet 1/1/1

Hitless management on the FSX 800 and FSX 1600

Hitless management is supported on the FSX 800 and FSX 1600 chassis with dual management modules. It is a high-availability feature set that ensures no loss of data traffic during the following events:

- Management module failure or role change
- Software failure
- Addition or removal of modules
- Operating system upgrade

During such events, the standby management module takes over the active role and the system continues to forward traffic seamlessly, as if no failure or topology change has occurred. In software releases that do not support hitless management, events such as these could cause a system reboot, resulting in an impact to data traffic.

The following Hitless management features are supported:

Hitless Switchover - A manually controlled (CLI-driven) switchover of the active and standby management modules without any packet loss to the services and protocols that are supported by Hitless management. A switchover is activated by the CLI command **switch-over-active-role** .

Hitless Failover - An automatic, forced switchover of the active and standby management modules because of a failure or abnormal termination of the active management module. In the event of a failover, the active management module abruptly leaves and the standby management module immediately assumes the active role. Like a switchover, a failover occurs without any packet loss to hitless-supported services and protocols. Unlike a switchover, a failover generally happens without warning.

Hitless Operating System (OS) Upgrade - An operating system upgrade and controlled switchover without any packet loss to the services and protocols that are supported by Hitless management. The services and protocols supported by Hitless management are listed in this section. Hitless failover and hitless switchover are disabled by default.

Benefits of hitless management

The benefits of Hitless management include the following:

- The standby management module (the module that takes over the active role) and all interface modules in the chassis are not reset
- Existing data traffic flows continue uninterrupted with no traffic loss
- Port link states remain UP for the duration of the hitless management event
- System configurations applied through Console/SNMP/HTTP interfaces remain intact
- Hitless switchover can be used by a system administrator, for example, to perform maintenance on a management module that has been functioning as the active management module. Some advantages of a hitless switchover over a hitless software reload are:
 - A manual switchover is quicker, since the standby module does not have to reboot.
 - Switched traffic through the Ethernet interfaces on the standby management module is not interrupted.

NOTE

All traffic going through Ethernet interfaces (if present) on the management modules will be interrupted during a hitless OS upgrade. This is because both management modules must be reloaded with the new image. This applies to hitless OS upgrade only. It does not apply to hitless switchover or failover, which does not interrupt traffic going through Ethernet interfaces on the standby management module (the module that takes over the active role).

Supported protocols and services for hitless management events

The following table lists the services and protocols that are supported by Hitless management, and also highlights the impact of *Hitless management events* (switchover, failover, and OS upgrade) to the system's major functions. The services and protocols that are not listed may be disrupted, but will resume normal operation once the new active management module is back up and running.

TABLE 18 Hitless-supported services and protocols - FSX 800 and FSX 1600

Traffic type	Supported protocols and services	Impact
Layer 2 switched traffic, including unicast and multicast + System-level + Layer 4	<ul style="list-style-type: none"> • 802.1p and 802.1Q • 802.3ad - LACP • 802.3af - PoE • 802.3at - PoE+ • DSCP honoring and Diffserv • Dual-mode VLAN • IGMP v1, v2, and v3 snooping • IPv4 ACLs • IPv6 ACLs • Layer 2 switching (VLAN and 802.1Q-in-Q) • MLD v1 and v2 snooping • MRP • Multiple spanning tree (MSTP) • Physical port/link state • PIM SM snooping • Port mirroring and monitoring • Port trunking • Rapid spanning tree (RSTP) • Spanning tree (STP) • ToS-based QoS • Policy Based Routing • Traffic policies • UDLD • VSRP 	<p>Layer 2 switched traffic is not impacted during a Hitless management event. All existing switched traffic flows continue uninterrupted.</p> <p>New switched flows are not learned by the FastIron switch during the switchover process and are flooded to the VLAN members in hardware. After the new active management module becomes operational, new switched flows are learned and forwarded accordingly. The Layer 2 control protocol states are not interrupted during the switchover process.</p> <p>Configured ACLs, PBR or GRE & IPv6 to IPv4 Tunnels will operate in a hitless manner.</p>
Layer 3 IPv4 routed traffic	<ul style="list-style-type: none"> • BGP4 • IPv4 unicast forwarding • OSPFv2 • OSPFv2 with ECMP • Static routes • IPv4 PIM (IPv4 non-stop multicast routing needs to be enabled for IPv4 PIM to be hitless.) • VRRP • VRRP-E • GRE • IPv6 to IPv4 Tunnels 	<p>Layer 3 routed traffic for supported protocols is not impacted during a Hitless management event.</p> <p>Other Layer 3 protocols that are not supported will be interrupted during the switchover or failover.</p> <p>If BGP4 graceful restart or OSPF graceful restart is enabled, it will be gracefully restarted and traffic will converge to normalcy after the new active module becomes operational.</p> <p>Configured ACLs, PBR or GRE & IPv6 to IPv4 Tunnels will operate in a hitless manner.</p>
Layer 3 IPv6 routed traffic	<ul style="list-style-type: none"> • BGP4+ • IPv6 unicast forwarding • OSPFv3 • OSPFv3 with ECMP • Static routes • VRRP • VRRP-E 	<p>Layer 3 routed traffic for supported protocols is not impacted during a Hitless management event. Traffic will converge to normalcy after the new active module becomes operational.</p> <p>Other Layer 3 protocols that are not supported will be interrupted during the switchover or failover.</p> <p>If BGP4+ graceful restart or OSPF graceful restart / OSPFv3 NSR is enabled, it will be gracefully restarted and traffic will converge to normalcy after the new active module becomes operational.</p> <p>Configured ACLs will operate in a hitless manner.</p>

TABLE 18 Hitless-supported services and protocols - FSX 800 and FSX 1600 (continued)

Traffic type	Supported protocols and services	Impact
Management traffic	N/A	All existing management sessions (SNMP, TELNET, HTTP, HTTPS, FTP, TFTP, SSH etc.), are interrupted during the switchover or failover process. All such sessions are terminated and can be re-established after the new Active Controller takes over.
Security	<ul style="list-style-type: none"> 802.1X, including use with dynamic ACLs and VLANs IPv4 ACLs IPv6 ACLs DHCP snooping Dynamic ARP inspection EAP with RADIUS IP source guard Multi-device port authentication, including use with dynamic ACLs and VLANs 	<p>Supported security protocols and services are not impacted during a switchover or failover.</p> <p>NOTE If 802.1X and multi-device port authentication are enabled together on the same port, both will be impacted during a switchover or failover. Hitless support for these features applies to ports with 802.1X <i>only</i> or multi-device port authentication <i>only</i>.</p> <p>Configured ACLs will operate in a hitless manner, meaning the system will continue to permit and deny traffic during the switchover or failover process.</p>
Other services to Management	<ul style="list-style-type: none"> AAA DHCP sFlow SNMP v1, v2, and v3 SNMP traps NTPv4 Traceroute 	<p>Supported protocols and services are not impacted during a switchover or failover.</p> <p>DNS lookups will continue after a switchover or failover. This information is not synchronized.</p> <p>Ping traffic will be minimally impacted.</p>

Hitless management configuration notes and feature limitations

The following limitations apply to hitless management support.

- All traffic going through Ethernet interfaces (if present) on the management modules will be interrupted during a hitless OS upgrade. This is because both management modules must be reloaded with the new image. This applies to hitless OS upgrade only. It does not apply to hitless switchover or failover, which does not interrupt traffic going through Ethernet interfaces on the standby management module (the module that takes over the active role).
- Static and dynamic multi-slot trunks will flap during a hitless switchover if any of the trunk port members reside on the management module.
- Layer 3 multicast traffic is not supported by Hitless management.

Hitless reload or switchover requirements and limitations

The section describes the design limitation on devices with the following configuration:

- 0-port management modules
- One or more third generation line cards

For hitless reload or switch-over-active-role to succeed, the following requirements and limitations must be met:

- Hitless reload is not allowed from Switch code when the primary image is in Router code and boot primary is a part of the configuration.
- The standby management module must be up and in an "OK {Enabled}" state.

- A configuration requiring a reload must not be pending.
- A hitless-reload must not have already been issued on the previous active management module.
- POE firmware must not be in progress.
- The SXR running configuration must not be classified as too large (greater than 512KB).
- A TFTP session must not be in progress.
- An image sync session must not be in progress.
- The current active management card cannot have a memory utilization of greater than 90% of available memory.
- A line card hotswap must not be in progress.

If any of these conditions are not met, an appropriate error message is printed to the console and hitless-reload or switch-over will not succeed.

With following steps, after switchover, the new standby goes into continuous reload state:

1. SXL box is running with build "x"
2. Perform **copy tftp** of build "x+1" and wait for both active and standby to sync.
3. Execute **switch-over-active-role**.

With above step, the new active comes up but the new standby tries to load the primary image "x+1" and due to this there is image sync issue and new standby goes to continuous reload state without recovery. Hence, it is a limitation that after **copy tftp** operation to primary, **switch-over-active-role** operation should be avoided.

What happens during a Hitless switchover or failover

This section describes the internal events that enable a controlled or forced switchover (failover) to take place in a hitless manner, as well as the events that occur during the switchover.

Separate data and control planes

The FSX 800 and FSX 1600 management modules have separate data and control planes. The *data plane* forwards traffic between the switch fabric modules and all of the Interface modules in the chassis. The *control plane* carries traffic that is destined for the CPU of the active management module. Control plane traffic includes the following:

- Management traffic
- Control protocol traffic
- In some cases, the first packet of a data flow

During a controlled or forced switchover, the data plane is not affected. Traffic in the forwarding plane will continue to run without interruption while the standby management module takes over operation of the system. However, traffic in the control plane will be minimally impacted.

Real-time synchronization between management modules

Hitless management requires that the active and standby management modules are fully synchronized at any given point in time. This is accomplished by *baseline* and *dynamicsynchronization* of the modules.

When a standby management module is inserted and becomes operational in the FSX 800 or FSX 1600 chassis, the standby module sends a baseline synchronization request to the active management module. The request prompts the active management module to copy the current state of its CPU to the standby CPU, including:

- Start-up and run-time configuration (CLI)
- Layer 2 protocols - Layer 2 protocols such as STP, RSTP, MRP, and VSRP run concurrently on both the active and standby management modules.
- Hardware Abstraction Layer (HAL) - This includes the prefix-based routing table, next hop information for outgoing interfaces, and tunnel information.
- Layer 3 IP forwarding information - This includes the routing table, IP cache table, and ARP table, as well as static and connected routes.
- If NSR is enabled, OSPFv2 and OSPFv3 information is copied to the standby.

As baseline synchronization is performed, the console of the active management module displays the progress of the synchronization.

```
ACTIVE: Detected Stdbby heart-beat
ACTIVE: Standby is ready for baseline synchronization.
ACTIVE: Baseline SYNC is completed. Protocol Sync is in progress.
ACTIVE: State synchronization is complete.
```

The first message indicates that the active management module has detected the standby management module. The second message indicates that the standby module has been hot-inserted and is ready for baseline synchronization. The third message is seen when baseline synchronization is completed, and the fourth message is seen when protocol synchronization is completed.

The console of the standby management module also displays the progress of the synchronization.

```
STBY: Baseline SYNC is completed. Protocol Sync is in progress.
STBY: State synchronization is complete.
```

The first message indicates that baseline synchronization is completed, and the second message indicates that protocol synchronization is completed.

When control protocols are synchronized and protocol synchronization timers expire, the standby management module will be in *hot-standby* mode, meaning the standby module is ready to take over as the active management module. In the event of a switchover, the standby module will pick up where the active module left off, without interrupting data traffic.

After baseline synchronization, any new events that occur on the active CPU will be dynamically synchronized on the standby CPU.

Examples of such events include:

- CLI/HTTP/SNMP configurations
- CPU receive packets
- Link events
- Interrupts
- Layer 2 and Layer 3 forwarding table updates
- Dynamic user authentication updates such as 802.1X or multi-device port authentication
- Routing protocols OSPFv2 and OSPFv3 updates if NSR is enabled.

Dynamic events are synchronized in such a way that if the active CPU fails before fully executing an event, the standby CPU (newly active CPU) will execute the event after the failover. Also, if the active CPU aborts the event, the standby CPU will abort the event as well.

NOTE

Since both the standby and active management modules run the same code, a command that brings down the active management module will most likely bring down the standby management module. Because all configuration commands are synchronized from active to standby management module in real time, both management modules will reload at almost the same time. This in turn will cause the system to reset all interface modules (similar to the behavior when the **reboot** command is executed) and will cause packet loss associated with a system reboot.

NOTE

If the new active management module becomes out-of-sync with an interface module, information on the interface module can be overwritten in some cases, which can cause an interruption of traffic forwarding.

How a Hitless switchover or failover impacts system functions

For a description of the feature's impact to major system functions, refer to [Supported protocols and services for hitless management events](#) on page 117.

Enabling hitless failover on the FSX 800 and FSX 1600

Hitless failover is disabled by default. When disabled, the following limitations are in effect:

- If a failover occurs, the system will reload. The following message will display on the console prior to a reload.

```
STBY:- - - - Active Hitless Failover is disabled. Re-setting the system - -
```

- Manual switchover (CLI command **switch-over-active-role**) is not allowed. If this command is entered, the following message will display on the console:

```
Switch-over is not allowed. Reason: hitless-failover not configured.
```

NOTE

Hitless OS upgrade is *not* impacted by this option and is supported whether or not hitless failover is enabled.

NOTE

Synchronization between the active management module and standby management module will occur whether or not hitless failover is enabled.

To enable hitless failover, enter the following command at the Global CONFIG level of the CLI:

```
device(config)#hitless-failover enable
```

The command takes effect immediately. Manual switchover is allowed, and in the event of a failover, the standby management module will take over the active role without reloading the system.

Syntax: [no] hitless-failoverenable

Use the **no** form of the command to disable hitless failover once it has been enabled.

Executing a hitless switchover on the FSX 800 and FSX 1600

Hitless failover must be enabled before a hitless switchover can be executed.

To switch over to the standby module (and thus make it the active module), enter the following command.

```
device# switch-over-active-role
```

Once you enter this command, the system will prompt you as follows.

```
Are you sure? (enter 'y' or 'n'): y
Running Config data has been changed. Do you want to continue
the switch-over without saving the running config? (enter 'y' or 'n'): n
Please save the running config and try switch-over again
```

Syntax: switch-over-activerole

If this command is entered when hitless failover is disabled, the following message will appear on the console:

Switch-over is not allowed. Reason: hitless-failover not configured.

A management slot which is in active management preference will always attempt to be active on the next reboot.

To reset the preference, enter the command such as the following:

```
Brocade(config)# set-active-mgmt mgmt0/mgmt1
```

Syntax: set-active-management *management slot numbers*

NOTE

The default active management preference is set to mgmt0 (slot 9).

Hitless OS upgrade on the FSX 800 and FSX 1600

Hitless Operating System (OS) Upgrade enables an operating system upgrade and switchover without any packet loss to the services and protocols that are supported by Hitless management.

What happens during a Hitless OS upgrade

The following steps describe the internal events that occur during a hitless OS upgrade.

1. The standby management module resets and reloads with the new software image in its flash memory.
2. The Ethernet interfaces (if present) on the standby module become operational and start carrying data traffic.
3. The active management module synchronizes the standby management module with all the information required to take over the active role.
4. The Layer 2 and Layer 3 control protocols on the standby management module converge. This process takes approximately 70 seconds.
5. The standby management module takes over the active role.
6. The old active management module resets and reloads with the same software image running on the newly active management module.
7. The FastIron switch is now operating with the new software image. The management module that was initially configured as the standby management module is now the active management module and the management module that was initially configured as the active management module is now the standby.

NOTE

The events described above occur internally and do not create or affect the external network topology.

Hitless OS upgrade considerations

Consider the following when using the hitless OS upgrade feature:

- Hitless OS upgrade allows for upgrading the software in a system between two releases of the OS that support this functionality and have compatible data structures. A hitless O/S downgrade may also be supported if the current and target code releases have compatible data structures. From time to time it may be necessary, when enhancing the software or adding new features, to change or add data structures that may cause some releases to be incompatible. In such cases, an upgrade or downgrade will not be hitless, and the software will use the regular Brocade upgrade process - relying on fast reboot.
- For a description of how this feature impacts major system functions, refer to [Supported protocols and services for hitless management events](#) on page 117.
- You must have both active and standby management modules installed to use this feature.
- Hitless OS upgrade is supported in software release FSX 05.0.00 or higher, with boot image FSX 05.0.00 or higher. In general, it is supported with patch upgrades, for example, when upgrading from release 07.0.01a to 07.0.01b. It is not supported during major release upgrades, for example when upgrading from release 07.0.00 to 07.1.00.
- This feature can be used to upgrade an image to a higher or lower compatible version of the software. However, if hitless upgrade to a particular software version is not supported, the software upgrade must be performed through a fast reload of the system.
- Hitless OS upgrade between different types of software images is not supported. For example, hitless OS upgrade is supported when upgrading the Layer 2 image to another Layer 2 image. It is not supported when upgrading the Layer 2 image to Layer 3 image, and so on.
- Hitless OS upgrade should be performed locally, since remote connectivity will be lost during the upgrade. During a reload, HTTP, SSH, Telnet, SNMP, and ping sessions will be dropped.
- The active management module switches from the initial active management module to the standby management module during the hitless upgrade process. Therefore, a connection to the console interface on both management modules is required.
- Upon being reset, any traffic going through the ports on the management module will be interrupted. Once the management module is up and running, it will be able to send and receive packets, even before the hitless upgrade process is complete.
- The running configuration is not allowed to be changed any time during the hitless upgrade process.
- System-max configuration changes require a system reload. System-max configuration changes do not take effect by the hitless upgrade. Even if a system-max parameter is changed and saved in the startup configuration, the FastIron switch will revert to the default system-max value upon a hitless software upgrade. The new system-max value will only take effect after a regular system reload.
- Other commands requiring a software reload, such as CAM mode changes, also do not take effect upon hitless upgrade and require a system reload before being placed in effect.

Hitless OS upgrade configuration steps

The following is a summary of the configuration steps for a hitless OS software upgrade.

1. Copy the software image that supports hitless software upgrade from a TFTP server to the FastIron switch. Refer to [Loading the software onto the switch](#) on page 125.
2. Install the software image in flash memory on the active and standby management modules.
3. Enter the **hitless-reload** command on the active management module. The command triggers the events described in the section [What happens during a Hitless OS upgrade](#) on page 123.

Loading the software onto the switch

Hitless OS upgrade loads from the primary and secondary images on the FSX 800 and FSX 1600 Management modules. If you will be using the **hitless-reload** command to perform the hitless upgrade, you must first copy the software image that supports hitless software upgrade onto the flash memory of the active and standby management modules. For instructions, refer to the release notes.

Performing a hitless upgrade

After loading the software image onto the flash memory of the active and standby management modules, you can begin the process of performing a hitless OS upgrade using the **hitless-reload** command. For example,

```
device#hitless-reload primary
```

Syntax: **hitless-reload**primary | secondary

The *primary* parameter specifies that the management module will be reloaded with the **primary** image.

The *secondary* parameter specifies that the management module will be reloaded with the **secondary** image.

NOTE

The **hitless-reload** command is accepted only when the running configuration and startup configuration files match. If the configuration file has changed, you must first save the file (**write mem**) before executing a hitless reload. Otherwise, the following message will display on the console. Error: Running config and start-up config differs. Please reload the system or save the configuration before attempting hitless reload.

Syslog message for Hitless management events

The following Syslog message is generated as a result of a switchover or hitless OS upgrade.

```
SWITCHOVER COMPLETED - by admin - Mgmt Module in slot
slotnum
is now Active
```

The following Syslog message is generated as a result of a failover.

```
SWITCHOVER COMPLETED - by active CPU failure - Mgmt Module in slot
slotnum
is now Active
```

Displaying diagnostic information

Use the following commands to display diagnostic information for a hitless switchover or failover.

```
device#show ipc
Version 6, Grp 0, Recv: stk-p0: 840918, p1: 0, sum: 840918
Message types have callbacks:
 1:Reliable IPC message 2:Reliable IPC atomic 4:fragmentation,jumbo
20:SYNC dynamic change 22:SYNC download reply 24:SYNC download spec i
25:SYNC restart download 26:SYNC verification 27:SYNC disable/enable
29:SYNC mgmt hello 35:IPC Ready Msg 36:IPC Msg for Sync Fra
38:SYNC reliable
Send message types:
 [1]=815798, [21]=1, [35]=1, [38]=24442,
Recv message types:
 [1]=816446,0, [20]=2,0 [22]=1,0
 [29]=25,0, [38]=24442,0,
Statistics:
send pkt num : 840242, recv pkt num : 840918
send msg num : 840242, recv msg num : 840918,
send frag pkt num : 0, recv frag pkt num : 0,
```

```

pkt buf alloc : 832113,
Reliable-mail  send  success  receive  time us
target ID      0      0      0      0
target MAC     0      0      0      0
There is 0 current jumbo IPC session
Possible errors:
***recv msg no callback 2, last msg_type=20, from stack0, e1/9

```

Syntax:show ipc

```

device#show ipc_stat
Total available Hsync channel space = 1048580
Total available Appl channel space = 524292
Total number of application msgs in dyn queue = 0
Total number of hsync msgs in dyn queue = 0
Total number of rel sync msgs in dyn queue = 0
Total number of rx pkt msgs in standby dynamic queue
Total number of rx pkt msgs in active dyn queue = 0
Total number of rx pkts relayed = 0
Total number of rx pkts received = 5686578
Total number of dyn-sync messages received so far = 3
Total number of rel-sync pending complete = 0
Total number of L3 baseline-sync packets = 655
Total number of packet drops in sync = 0
Is image_sync_in_progress? = 0
Total num of rx dyn queue drops = 0
Total num of jumbo corrupts = 0
Total number of messages in IP send queue = 0

```

Syntax: showipc_stat

Displaying management redundancy information

Enter the following command at any level of the CLI, to view the redundancy parameter settings and statistics.

```

Brocade(config)# show redundancy
=== MP Redundancy Settings ===
Configured Active Slot = 9
Running-Config Sync Period = (upon "write mem")
=== MP Redundancy Statistics ===
Current Active Session:
Active mgmt slot = 9, Standby mgmt slot = 10 (Absent)
Switchover cause = No Switchover
Start Time      = Jan  1 00:00:09
Sxr Sys Hitless Enable Status = 0
Total number of Switchover/Failovers = 0
L3 slib baseline sync status: 0 [complete]

```

Layer 3 hitless route purge

Layer 3 traffic is forwarded seamlessly during a failover, switchover, or OS upgrade when hitless management is enabled.

Some protocols support non-stop routing. On enabling non-stop routing, after switchover the management module quickly re-converge the protocol database. Whereas, some protocols support graceful restart, in which the protocol state is re-established with the help of neighboring devices. Once all the protocols converge the routes which were removed from the network during the convergence period, the routes are deleted from the devices. You can set the route purge timer per VRF instance. Configure the timer to set the duration for which the routes should be preserved after switchover. Once this period elapses, the route purging starts, if by then all other protocols have finished non-stop routing or graceful restart.

When switchover occurs, the route purge timer starts. If non-stop routing or graceful restart is also configured, the route validation and purging starts only when they are complete and the purge timer has elapsed. If for some reason more delay is expected in learning the routes, you can configure a larger period for the purge timer.

Setting the IPv4 hitless purge timer on the default VRF

To configure the purge timer, enter the **ip hitless-route-purge-timer** command in global configuration mode.

Example for setting IPv4 hitless purge timer on the default VRF

The following example shows how to set the IPv4 hitless purge timer on the default VRF:

```
Brocade(config)# ip hitless-route-purge-timer 60
```

Setting the IPv4 hitless purge timer on the non-default VRF

1. Enter the VRF configuration mode using the **vrf** command.
2. Configure route distinguisher using the **rd** command.
3. Enter IPv4 address family configuration mode using the **address-family ipv4** command.
4. Configure the router purge timer using the **ip hitless-route-purge-timer** command.

Example for setting the IPv4 hitless purge timer on the non-default VRF

The following example shows how to set the IPv4 purge timer on the non-default VRF:

```
Brocade(config)# vrf blue
Brocade(config-vrf-blue)# rd 10:10
Brocade(config-vrf-blue)# address-family ipv4
Brocade(config-vrf-blue-ipv4)# ip hitless-route-purge-timer 60
```

Setting the IPv6 hitless purge timer on the default VRF

To configure the purge timer, enter the **ipv6 hitless-route-purge-timer** command in global configuration mode.

Example for setting the IPv6 hitless purge timer on the default VRF

The following example shows how to set the IPv6 hitless purge timer on the default VRF:

```
Brocade(config)# ipv6 hitless-route-purge-timer 60
```

Setting the IPv4 hitless purge timer on the non-default VRF

Before you begin: Enable IPv6 unicast routing using the **ipv6 unicast-routing** command in global configuration mode.

1. Enter the VRF configuration mode using the **vrf** command.
2. Configure route distinguisher using the **rd** command.
3. Enter the IPv6 address family configuration mode using the **address-family ipv6** command.
4. Configure the router purge timer using the **ipv6 hitless-route-purge-timer** command.

Example for setting the IPv6 hitless purge timer on the non-default VRF

The following example shows how to set the IPv6 purge timer on the non-default VRF:

```
Brocade(config)# vrf blue
Brocade(config-vrf-blue)# rd 10:10
Brocade(config-vrf-blue)# address-family ipv6
Brocade(config-vrf-blue-ipv4)# ipv6 hitless-route-purge-timer 60
```

DHCP Client-Based Auto-Configuration and Flash image update

NOTE

The DHCP Client-Based Auto-Configuration and Flash image update are platform independent and have no differences in behavior or configuration across platforms (FSX, FCX, and ICX).

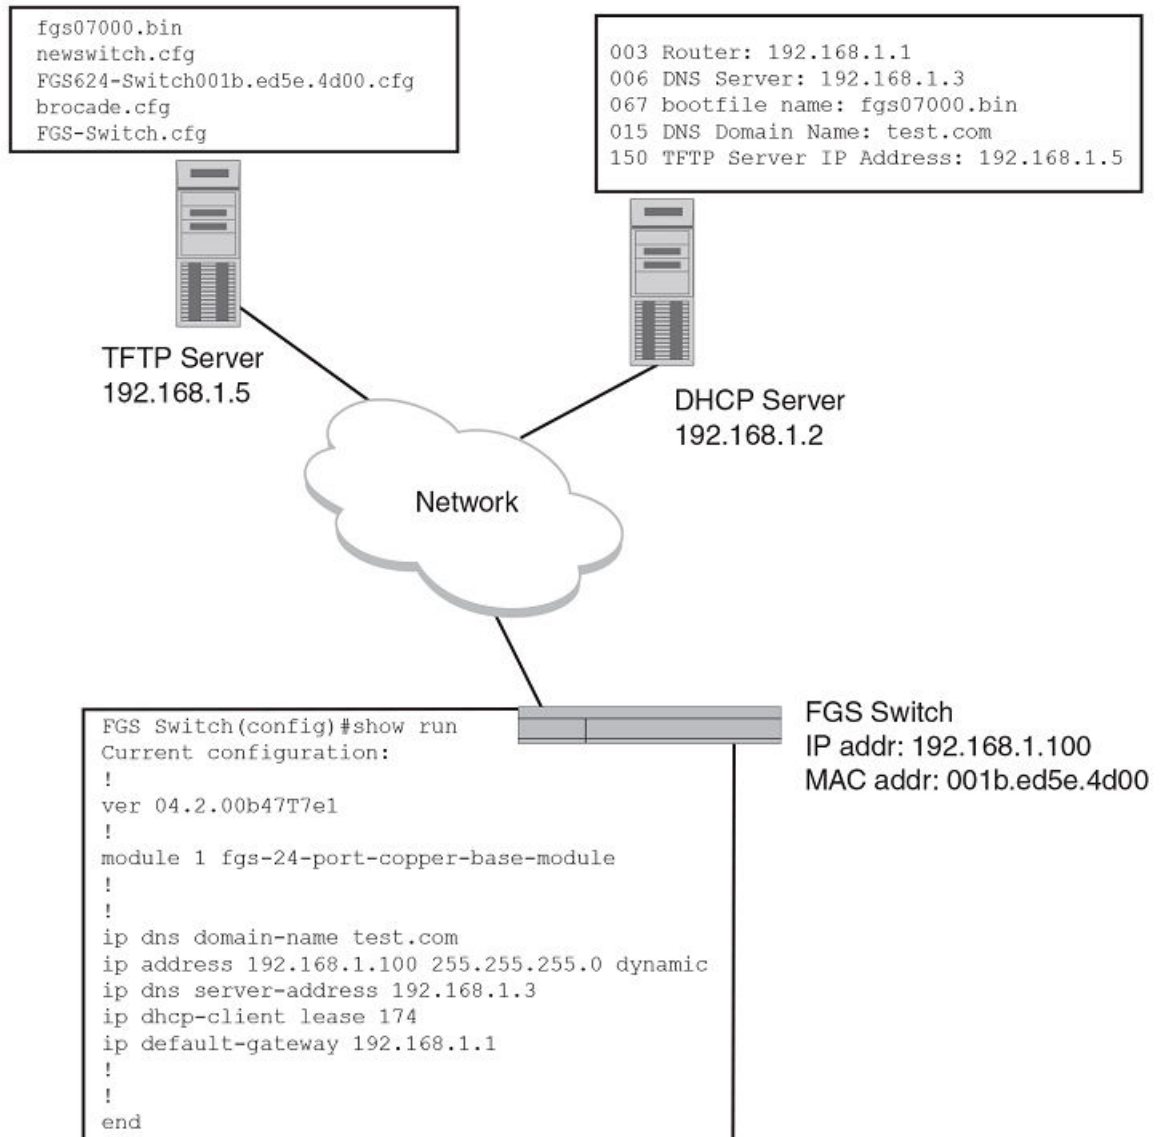
DHCP Client-Based Auto-Configuration allows Layer 2 and Layer 3 devices to automatically obtain leased IP addresses through a DHCP server, negotiate address lease renewal, and obtain flash image and configuration files.

DHCP Client-Based Auto-Configuration occurs as follows.

1. The IP address validation and lease negotiation enables the DHCP client (a Brocade Layer 2 or Layer 3 device) to automatically obtain and configure an IP address, as follows:
 - One lease is granted for each Layer 2 device. If the device is configured with a static IP address, the DHCP Auto-Configuration feature is automatically disabled.
 - For a Layer 3 device, one leased address is granted (per device) to the interface that first receives a response from the DHCP server.
2. If **auto update** is enabled, the TFTP flash image is downloaded and updated. The device compares the file name of the requested flash image with the image stored in flash. If the file names are different, then the device will download the new image from a TFTP server, write the downloaded image to flash, and then reload the device or stack.

- In the final step, TFTP configuration download and update, the device downloads a configuration file from a TFTP server and saves it as the running configuration.

FIGURE 2 DHCP Client-Based Auto-Configuration



Configuration notes and feature limitations for DHCP Client-Based Auto-Configuration

- For Layer 2 devices, this feature is available for default VLANs and management VLANs. This feature is not supported on virtual interfaces (VEs), trunked ports, or LACP ports.
- Although the DHCP server may provide multiple addresses, only one IP address is installed at a time.
- This feature is not supported together with DHCP snooping.

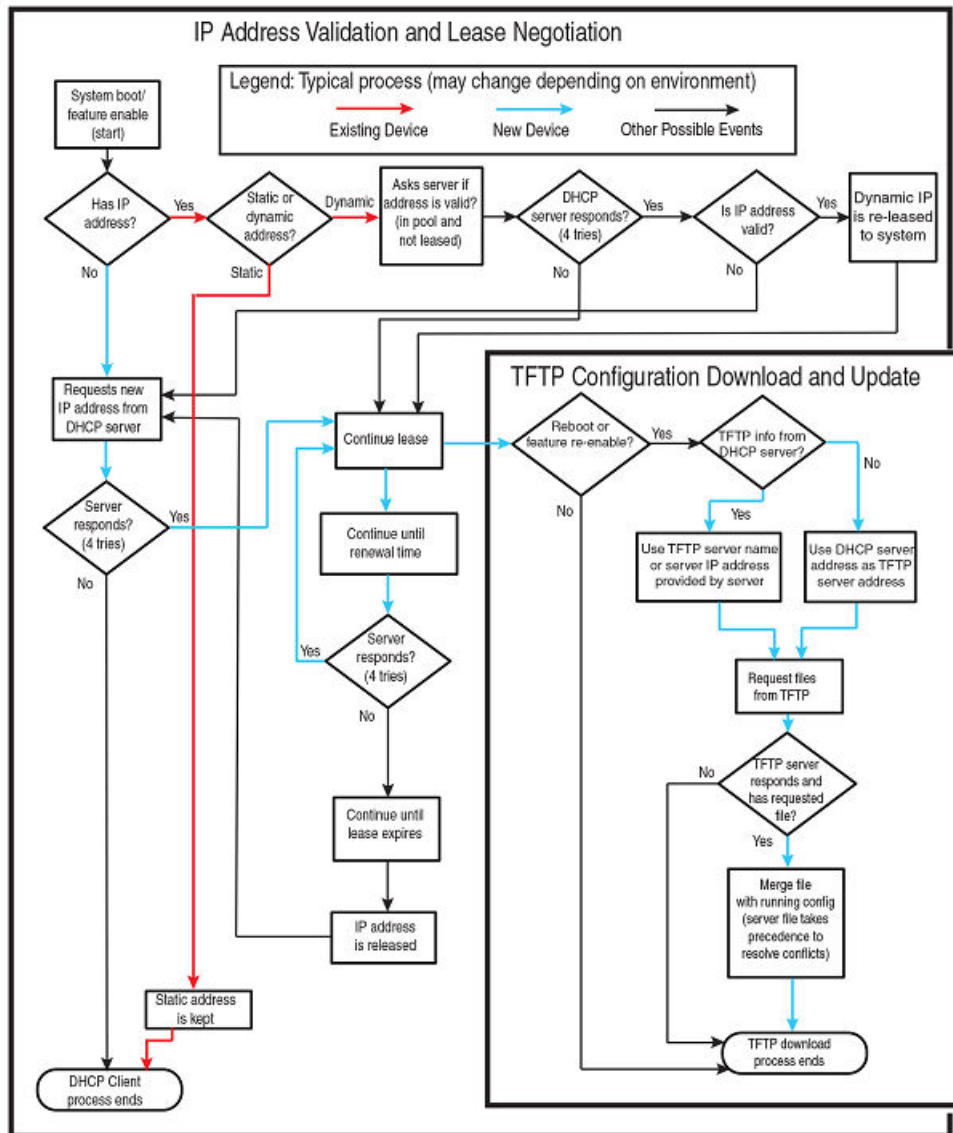
The following configuration rules apply to flash image update:

- To enable flash image update (**ip dhcp-client auto-update enable** command), also enable auto-configuration (**ip dhcp-client enable** command).
- The image file name to be updated must have the extension **.bin**.
- The DHCP option 067 bootfile name will be used for image update if it has the extension **.bin**.
- The DHCP option 067 bootfile name will be used for configuration download if it does not have the extension **.bin**.
- If the DHCP option 067 bootfile name is not configured or does not have the extension **.bin**, then the auto-update image will not occur.
- While updating the image using option 67, the image types (Layer 2 or Layer 3) should match. For example, if the DHCP client with a Layer 3 image downloads an updated Layer 2 image using option 067, the configuration download will fail.

How DHCP Client-Based Auto-Configuration and Flash image update works

Auto-Configuration and Auto-update are enabled by default. To disable this feature, refer to [Disabling or re-enabling Auto-Configuration](#) on page 134 and [Disabling or re-enabling Auto-Update](#) on page 134, respectively.

FIGURE 3 DHCP Client-Based Auto-Configuration steps



Validate the IP address and lease negotiation

1. At boot-up, the device automatically checks its configuration for an IP address.
2. If the device does not have a static IP address, it requests the lease of an address from the DHCP server:
 - If the server responds, it leases an IP address to the device for the specified lease period.
 - If the server does not respond (after four tries) the DHCP Client process is ended.

3. If the device has a dynamic address, the device asks the DHCP server to validate that address. If the server does not respond, the device will continue to use the existing address until the lease expires. If the server responds, and the IP address is outside of the DHCP address pool or has been leased to another device, it is automatically rejected, and the device receives a new IP address from the server. If the existing address is valid, the lease continues.

NOTE

The lease time interval is configured on the DHCP server, not on the client device. The `ip dhcp-client lease` command is set by the system, and is non-operational to a user.

4. If the existing address is **static**, the device keeps it and the DHCP Client process is ended.
5. For a leased IP address, when the lease interval reaches the renewal point, the device requests a renewal from the DHCP server:
 - If the device is able to contact the DHCP server at the renewal point in the lease, the DHCP server extends the lease. This process can continue indefinitely.
 - If the device is unable to reach the DHCP server after four attempts, it continues to use the existing IP address until the lease expires. When the lease expires, the dynamic IP address is removed and the device contacts the DHCP server for a new address. If the device is still unable to contact the DHCP server after four attempts, the process is ended.

TFTP Flash image download and update**NOTE**

This process only occurs when the client device reboots, or when DHCP-client has been disabled and then re-enabled.

Once a lease is obtained from the server, the device compares the file name of the requested flash image with the image stored in flash. In a stacking configuration, the device compares the file name with the image stored in the Active controller only.

- If the .bin file names match, then the DHCP client skips the flash image download. If auto configuration is enabled, the DHCP client proceeds with downloading the configuration files.
- If the .bin file names are different, then the DHCP client downloads the new image from a TFTP server and then writes the downloaded image to flash. In a stacking configuration, the device copies the flash image to flash in all stack member units.

The code determines which flash (primary or secondary) to use based on how the device is booted. In a stacking configuration, the member units use the same flash as the Active controller. Once the flash is updated with the newer flash image, the device is reloaded, and any member units in a stacking configuration are reloaded as well. If auto configuration is enabled, the DHCP client then proceeds to download the configuration files.

NOTE

In a stacking environment, the DHCP client flash image download waits five minutes for all member units to join and update. Then the DHCP client downloads the new image from the TFTP server using the TFTP server IP address (option 150), if it is available. If the TFTP server IP address is not available, the DHCP client requests the TFTP file from the DHCP server.

TFTP configuration download and update

NOTE

This process only occurs when the client device reboots, or when Auto-Configuration has been disabled and then re-enabled.

- When the device reboots, or the Auto-Configuration feature has been disabled and then re-enabled, the device uses information from the DHCP server to contact the TFTP server to update the running-configuration file:
 - If the DHCP server provides a TFTP server name or IP address, the device uses this information to request files from the TFTP server.
 - If the DHCP server does not provide a TFTP server name or IP address, the device requests the configuration files from the DHCP server.

- The device requests the configuration files from the TFTP server by asking for file names in the following order:

- bootfile name provided by the DHCP server (if configured)
- hostnameMAC-config.cfg, for example:

```
FCX001p-Switch0000.005e.4d00-config.cfg
```

- hostnameMAC.cfg, for example:

```
FCX002p-Switch0000.005e.4d00.cfg
```

- brocade.cfg (applies to all devices), for example:

```
brocade.cfg
```

- <fcx | icx>-<switch | router>.cfg** (applies to Layer 2 devices), for example:

```
fcx-switch.cfg
(FCX Layer 2)
icx-switch.cfg
(ICX Layer 2)
```

If the device is successful in contacting the TFTP server and the server has the configuration file, the files are merged. If there is a conflict, the server file takes precedence.

If the device is unable to contact the TFTP server or if the files are not found on the server, the TFTP part of the configuration download process ends.

Supported options for DHCP servers

DHCP Client supports the following options:

- 001 - subnetmask
- 003 - router ip
- 015 - domain name
- 006 - domain name server
- 012 - hostname (optional)
- 066 - TFTP server name (only used for Client-Based Auto Configuration)
- 067 - bootfile name
- 150 - TFTP server IP address (private option, datatype = IP Address)

Configuration notes for DHCP servers

- When using DHCP on a router, if you have a DHCP address for one interface, and you want to connect to the DHCP server from another interface, you must disable DHCP on the first interface, then enable DHCP on the second interface.
- When DHCP is disabled, and then re-enabled, or if the system is rebooted, the TFTP process requires approximately three minutes to run in the background before file images can be downloaded manually.
- Once a port is assigned a leased IP address, it is bound by the terms of the lease regardless of the link state of the port.

Disabling or re-enabling Auto-Configuration

For a switch, you can disable or enable this feature using the following commands.

```
device(config)# ip dhcp-client enable
device(config)# no ip dhcp-client enable
```

For a router, you can disable or enable this feature using the following commands.

```
device(config-if-e1000-0/1/1)# ip dhcp-client enable
device(config-if-e1000-0/1/1)# no ip dhcp-client enable
```

Syntax: [no] ip dhcp-client enable

Disabling or re-enabling Auto-Update

Auto-update is enabled by default. To disable it, use the following command.

```
device(config)# no ip dhcp-client auto-update enabled
```

To re-enable auto-update after it has been disabled, use the following command.

```
device(config)# ip dhcp-client auto-update enabled
```

Syntax:[no] ip dhcp-client auto-update enabled

Configurable DHCP address acquisition attempts

This DHCP enhancement allows a DHCP client to make configurable DHCP address acquisition attempts at lower rates without moving the client to a stopped state.

The Brocade implementation of this enhancement follows RFC 2131. Normally, a DHCP client acquires dynamic IP addresses from the DHCP server in two modes:

- **Boot mode** - When the system is initially booted, the DHCP client tries to acquire dynamic IP addresses from the server when the DHCP server is reachable; otherwise, it disables the DHCP client automatically.
- **Run mode** - When the administrator enables the DHCP client at the interface or global level to get dynamic IP addresses, the DHCP client tries a maximum of four times (in an exponential manner) to acquire dynamic IP addresses from the DHCP server; otherwise, it stops the DHCP client automatically.

With the configurable DHCP address acquisition attempts enhancement, the DHCP client sends DHCP discover messages periodically in run mode, based on two configured time intervals: the discovery interval and the continuous mode max duration interval.

Configuring DHCP address acquisition attempts

The client attempts to obtain an IP address during start up. If an IP address is not obtained on the first attempt, the client makes up to three further attempts to obtain an address. If after four attempts the client is unable to obtain an IP address the following occurs:

- If DHCP continuous discovery mode is enabled, the client continues to request an IP address at regular intervals. This ensures that if initial DHCP requests fail, due to an outage between the client and server for example, the client continues to send DHCP requests to the server.
- If DHCP continuous discovery mode is not enabled, no further attempts are made to obtain an IP address and the client is disabled.

DHCP continuous discovery mode

DHCP continuous discovery mode is not enabled by default. To enable this mode, you must configure both of these commands:

- `ip dhcp-client continuous-mode max-duration interval`
- `ip dhcp-client discover-interval interval`

Provided both these commands are configured, the DHCP client discovery process starts automatically when the system boots up and runs continuously.

For example, to configure the client to retry address acquisition every 20 minutes and to continue repeating address acquisition for up to 2 hours, if an address is not obtained, complete the following steps.

1. Enter global configuration mode (on a switch) or interface configuration mode (on a router).
2. Enter the `ip dhcp-client discover-interval` command.

```
device(config)# ip dhcp-client discover-interval 20
```

3. Enter the `ip dhcp-client continuous-mode max-duration` command.

```
device(config)# ip dhcp-client continuous-mode max-duration 2
```

Displaying DHCP configuration information

The following example shows output from the `show ip` command for Layer 2 devices.

```
device(config)# show ip
Switch IP address: 10.44.16.116
Subnet mask: 255.255.255.0
Default router address: 10.44.16.1
TFTP server address: 10.44.16.41
Configuration filename: foundry.cfg
Image filename: None
```

The following example shows output from the `show ip address` command for a Layer 2 device.

```
device(config)# show ip address
IP Address      Type      Lease Time      Interface
10.44.16.116    Dynamic   174              0/1/1
```

The following example shows output from the `show ip address` command for a Layer 3 device.

```
device(config)# show ip address
IP Address      Type      Lease Time      Interface
10.44.3.233     Dynamic   672651          0/1/2
10.0.0.1        Static    N/A              0/1/15
```

The following example shows a Layer 2 device configuration as a result of the `show run` command.

```
device(config)# show run
Current configuration:
```

```

!
ver 08.0.00a
!
module 1 fcx-24-port-base-module
!
!ip dns domain-list englab.brocade.com
ip dns domain-list companynet.com
ip dns server-address 10.31.2.10
ip route 0.0.0.0/0 10.25.224.1
!ipv6 rguard policy pl
!ipv6 dns server-address 200::1 8000::60 7000::61
!!
end

```

The following example shows a Layer 3 device configuration as a result of the **show run** command.

```

device(config)# show run
Current configuration:
!
ver 08.0.00a
!
module 1 fcx-24-port-management-module
module 2 fcx-2-port-10g-module
module 3 fcx-1-port-10g-module
!
vlan 1 name DEFAULT-VLAN by port
!
ip dns server-address 10.44.3.111
interface ethernet 0/1/2
 ip address 10.44.3.233 255.255.255.0 dynamic
 ip dhcp-client lease 691109
!
interface ethernet 0/1/15
 ip address 10.0.0.1 255.0.0.0
 ip helper-address 1 10.44.3.111
!
end

```

NOTE

The **ip dhcp-client lease** entry in the previous example applies to FastIron X Series devices only.

DHCP log messages

The following DHCP notification messages are sent to the log file.

```

2d01h48m21s:I: DHCPC: existing ip address found, no further action needed by DHCPC
2d01h48m21s:I: DHCPC: Starting DHCP Client service
2d01h48m21s:I: DHCPC: Stopped DHCP Client service
2d01h48m21s:I: DHCPC: FCX24P Switch running-configuration changed
2d01h48m21s:I: DHCPC: sending TFTP request for bootfile name fgs-switch.cfg
2d01h48m21s:I: DHCPC: TFTP unable to download running-configuration
2d01h48m21s:I: DHCPC: Found static IP Address 10.1.1.1 subnet mask 255.255.255.0 on port 0/1/5
2d01h48m21s:I: DHCPC: Client service found no DHCP server(s) on 3 possible subnet
2d01h48m21s:I: DHCPC: changing 0/1/3 protocol from stopped to running

```


Energy Efficient Ethernet

Energy Efficient Ethernet (EEE) regulates and saves power consumed by the active hardware components in the switch and conserves power during idle time.

EEE allows Brocade devices to conform to green computing standards. This functionality is achieved by moving the data ports to a low-power state when their function is not necessary or when they are in a passive, no traffic condition. The EEE feature in switching platforms reduces overall energy consumption, cooling, noise, and operating costs for energy and cooling. Lower power consumption also means lower heat dissipation and increased system stability, less energy usage, thereby reducing costs and impact on the environment.

EEE is a set of enhancements to the Ethernet specification to address power consumption during periods of low data activity. EEE is specified in IEEE Std 802.3az-2010 which is an amendment to the IEEE Std 802.3-2008 specification. The optional EEE capability combines the IEEE 802.3 Media Access Control (MAC) sublayer with a family of physical layers defined to support operation in the Low Power Idle (LPI) mode. When the LPI mode is enabled, systems on both sides of the link can save power during periods of low link utilization. LPI signaling allows the LPI client to indicate to the PHY, and to the link partner, that a break in the data stream is expected. The LPI client can then use this information to enter power-saving modes that require additional time to resume normal operation. LPI signaling also informs the LPI client when the link partner sends such an indication.

Port support for Energy Efficient Ethernet

- On ICX 7450 devices EEE is supported on 1G copper ports and 10G copper module ports.
- On ICX 7250 devices EEE is supported on 1G copper ports.
- You may notice port flap on the port when EEE is enabled.
- EEE is not supported on 1G fiber ports (ICX7450-48F), 4x10F module ports, and 1x40Q module ports.

Enabling Energy Efficient Ethernet

Energy Efficient Ethernet (EEE) is supported on select Brocade devices and can be enabled globally or per port.

Follow these steps to enable EEE globally or per port.

1. Enter global configuration mode.
2. Enter the **eee** command. The following example shows enabling EEE globally.

```
device(config)# eee
EEE Feature Enabled
```

3. To enable EEE at the interface level, enter the **eee** command in interface configuration mode.

```
device(config-if-e1000-1/1/1)# eee
EEE Feature Enabled on port 1/1/1
```

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 (task usage is divided into different interval levels called *buckets*). For example, the task run time is divided into buckets: bucket 1 (0-50 ms), bucket 2 (50-100 ms), bucket 3 (100-150 ms), and so on. The CPU histogram collects the task CPU usage in each bucket. This includes how many times a task run time or hold time falls in each bucket, and the maximum run time and total run time for each bucket. CPU histogram information is measured for the hold-time and wait-time of the task.

- Hold time - The time that the task is holding the CPU without yield.
- Wait time - The time that the task is waiting for execution.

External USB Hotplug

External USB Hotplug support allows you to copy images, cores, logs, and configurations between the external USB and the internal eUSB.

Brocade device images are stored in the raw partition. Cores, logs and configurations are stored in the ext4 filesystem partition. The introduction of the External USB Hotplug gives you the option to easily copy device images, cores, logs, and configurations between the external USB and the internal flash.

External USB Hotplug considerations

- Only USB drives of up to 128 GB of any vendor type are supported.
- USB 3.0 is not supported.
- You can copy files of less than 2 GB only.
- Make sure the external USB is formatted as a "FAT" filesystem before attempting to use it. Formatting can be done on a PC or on the Brocade device with the **format disk0** command.
- You should not insert a USB-based disk drive, nor should you insert a USB hub to connect multiple USB disks.
- copy TFTP/SCP to disk0 and disk0 to TFTP/SCP commands are not supported.
- Only an administrator can execute operations on an external USB, similar to TFTP.
- You cannot access the active unit's local external USB from a member unit and vice versa.
- Boot from an external USB is not supported.
- You must run the **unmount disk0** command before unplugging the external USB. The external USB can be mounted using the **mount disk0** command.
- The USB drive is only functional on the active member in a stacked environment.

Using External USB Hotplug

Plug in the External USB to begin using the External USB Hotplug commands. Use the **show files disk0** command to check if the external USB is mounted and ready to use.

You can use the commands in the following table as part of the External USB Hotplug functionality.

TABLE 19 External USB Hotplug commands

Command	Description
show files disk0	Displays the files in the external USB drive.
format disk0	Formats the external USB.
mount disk0	Mounts the filesystem in the external USB drive.
unmount disk0	Unmounts the filesystem of the external USB drive. This command is required to safely plug out the USB, so that files are not lost or corrupted.
copy flash disk0 primary \ secondary	Copies the image binary stored in the primary or secondary partition of the flash to a destination file in the external USB.
copy flash disk0 file	Copies any file from a source file in the system flash to an external USB destination file.
copy disk0 license	Copies the license file present in the external USB drive to the system.
copy disk0 running-config	Copies the configuration file present on the external USB drive to the system's running configuration.
copy disk0 startup-config	Copies the configuration file present on the external USB drive to the system's startup configuration file.

Refer to the *FastIron Command Reference Guide* for details on using the External USB Hotplug commands.

IPv6

- [Static IPv6 route configuration.....141](#)
- [IPv6 over IPv4 tunnels.....143](#)

Static IPv6 route configuration

NOTE

Static IPv6 route configuration is supported only with the IPv6 Layer 3 license on FSX devices and the full Layer 3 image on other devices.

You can configure a static IPv6 route to be redistributed into a routing protocol, but you cannot redistribute routes learned by a routing protocol into the static IPv6 routing table.

NOTE

The maximum IPv6 static routes supported on an ICX 6450 device is 1070.

Before configuring a static IPv6 route, you must enable the forwarding of IPv6 traffic on the Layer 3 switch using the **ipv6 unicast-routing** command and enable IPv6 on at least one interface by configuring an IPv6 address or explicitly enabling IPv6 on that interface. For more information on performing these configuration tasks, refer to "Configuring IPv4 and IPv6 protocol stacks" section in the *FastIron Ethernet Switch Administration Guide*.

Configuring a static IPv6 route

To configure a static IPv6 route for a destination network with the prefix 2001:DB8::0/32, a next-hop gateway with the global address 2001:DB8:0:ee44::1, and an administrative distance of 110, enter the following command.

```
device(config)#ipv6 route 2001:DB8::0/32 2001:DB8:2343:0:ee44::1 distance 110
```

Syntax: **ipv6 route** *dest-ipv6-prefix / prefix-length next-hop-ipv6-address* [**metric**] [**distance number**]

To configure a static IPv6 route for a destination network with the prefix 2001:DB8::0/32 and a next-hop gateway with the link-local address fe80::1 that the Layer 3 switch can access through Ethernet interface 1/3/1, enter the following command.

```
device(config)#ipv6 route 2001:DB8::0/32 ethernet 1/3/1 fe80::1
```

Syntax: **ipv6 route** *dest-ipv6-prefix / prefix-length* [**ethernet unit/slot/port | ve num**] *next-hop-ipv6-address* [**metric**] [**distance number**]

To configure a static IPv6 route for a destination network with the prefix 2001:DB8::0/32 and a next-hop gateway that the Layer 3 switch can access through tunnel 1, enter the following command.

```
device(config)#ipv6 route 2001:DB8::0/32 tunnel 1
```

Syntax: **ipv6 route** *dest-ipv6-prefix / prefix-length interface port* [**metric**] [**distance number**]

The following table describes the parameters associated with this command and indicates the status of each parameter.

TABLE 20 Static IPv6 route parameters

Parameter	Configuration details	Status
The IPv6 prefix and prefix length of the route's destination network.	<p>You must specify the <i>dest-ipv6-prefix</i> parameter in hexadecimal using 16-bit values between colons as documented in RFC 2373.</p> <p>You must specify the <i>prefix-length</i> parameter as a decimal value. A slash mark (/) must follow the <i>ipv6-prefix</i> parameter and precede the <i>prefix-length</i> parameter.</p>	Mandatory for all static IPv6 routes.
<p>The route's next-hop gateway, which can be one of the following:</p> <ul style="list-style-type: none"> The IPv6 address of a next-hop gateway. A tunnel interface. 	<p>You can specify the next-hop gateway as one of the following types of IPv6 addresses:</p> <ul style="list-style-type: none"> A global address. A link-local address. <p>If you specify a global address, you do not need to specify any additional parameters for the next-hop gateway.</p> <p>If you specify a link-local address, you must also specify the interface through which to access the address. You can specify one of the following interfaces:</p> <ul style="list-style-type: none"> An Ethernet interface. A tunnel interface. A virtual interface (VE). <p>If you specify an Ethernet interface, also specify the port number associated with the interface. If you specify a VE or tunnel interface, also specify the VE or tunnel number.</p> <p>You can also specify the next-hop gateway as a tunnel interface. If you specify a tunnel interface, also specify the tunnel number.</p>	Mandatory for all static IPv6 routes.
The route's metric.	You can specify a value from 1 - 16.	Optional for all static IPv6 routes. (The default metric is 1.)
The route's administrative distance.	You must specify the distance keyword and any numerical value.	Optional for all static IPv6 routes. (The default administrative distance is 1.)

A metric is a value that the Layer 3 switch uses when comparing this route to other static routes in the IPv6 static route table that have the same destination. The metric applies only to routes that the Layer 3 switch has already placed in the IPv6 static route table.

The administrative distance is a value that the Layer 3 switch uses to compare this route with routes from other route sources that have the same destination. (The Layer 3 switch performs this comparison before placing a route in the IPv6 route table.) This parameter does not apply to routes that are already in the IPv6 route table. In general, a low administrative distance indicates a preferred route. By default, static routes take precedence over routes learned by routing protocols. If you want a dynamic route to be chosen over a static route, you can configure the static route with a higher administrative distance than the dynamic route.

Configuring a static route in a non-default VRF or User VRF

To configure a static IPv6 route for a destination network with the prefix 2001:DB8::0/32, a next-hop gateway with the global address 2001:DB8:0:ee44::1, in the non-default VRF named "blue", enter the following at the general configuration prompt.

```
device(config)# ipv6 route vrf blue 2001:DB8::0/32 2001:DB8:0:ee44::1
```

Syntax: [no] **ipv6 route vrf** *vrf-name* *dest-ipv6-prefix/prefix-length* *next-hop-ipv6-address*

The *dest-ip-addr* is the route's destination. The *dest-mask* is the network mask for the route's destination IPv6 address.

The *vrf-name* is the name of the VRF that contains the next-hop router (gateway) for the route.

The *next-hop-ip-addr* is the IPv6 address of the next-hop router (gateway) for the route.

NOTE

The *vrf* needs to be a valid VRF to be used in this command.

IPv6 over IPv4 tunnels

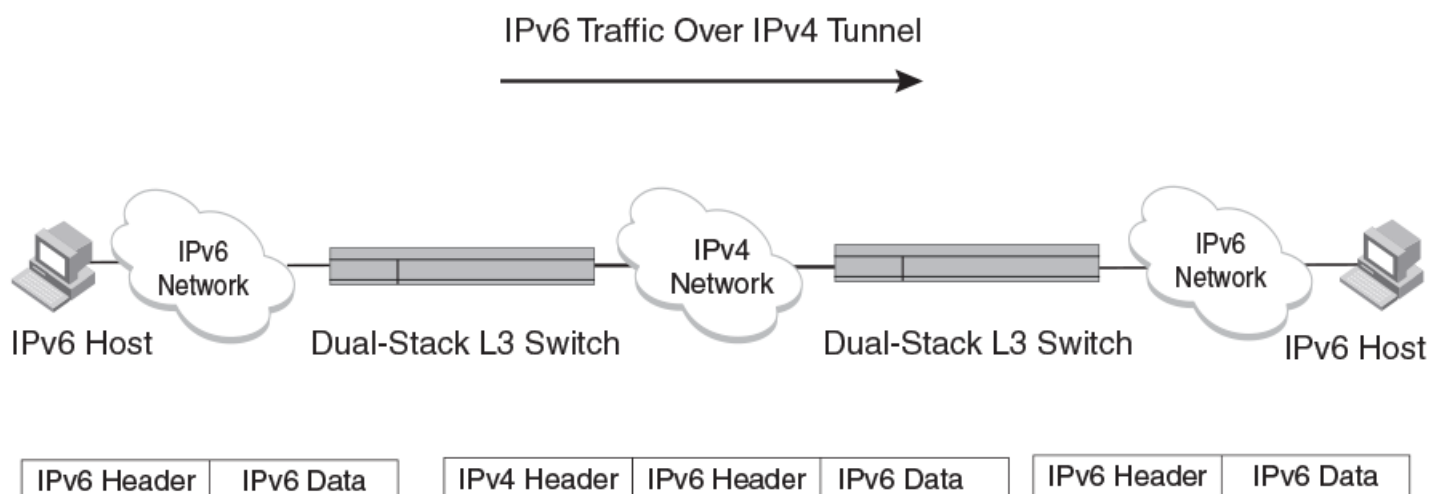
NOTE

This feature is supported only with the IPv6 Layer 3 license on FSX devices and the full Layer 3 image on other devices.

To enable communication between isolated IPv6 domains using the IPv4 infrastructure, you can manually configure IPv6 over IPv4 tunnels that provide static point-point connectivity.

As shown in the following illustration, these tunnels encapsulate an IPv6 packet within an IPv4 packet.

FIGURE 4 IPv6 over an IPv4 tunnel



In general, a manually configured tunnel establishes a permanent link between switches in IPv6 domains. A manually configured tunnel has explicitly configured IPv4 addresses for the tunnel source and destination.

This tunneling mechanism requires that the Layer 3 switch at each end of the tunnel run both IPv4 and IPv6 protocol stacks. The Layer 3 switches running both protocol stacks, or dual-stack routers, can interoperate directly with both IPv4 and IPv6 end systems and routers. Refer to "Configuring IPv4 and IPv6 protocol stacks" section in the *FastIron Ethernet Switch Administration Guide*.

IPv6 over IPv4 tunnel configuration notes

- The local tunnel configuration must include both source and destination addresses.
- The remote side of the tunnel must have the opposite source/destination pair.
- A tunnel interface supports static and dynamic IPv6 configuration settings and routing protocols.

- Duplicate Address Detection (DAD) is not currently supported with IPv6 tunnels. Make sure tunnel endpoints do not have duplicate IP addresses.
- Neighbor Discovery (ND) is not supported with IPv6 tunnels.
- If a tunnel source port is a multi-homed IPv4 source, the tunnel will use the first IPv4 address only. For proper tunnel operation, use the **ip address** option.

Configuring a manual IPv6 tunnel

You can use a manually configured tunnel to connect two isolated IPv6 domains. You should deploy this point-to-point tunneling mechanism if you need a permanent and stable connection.

To configure a manual IPv6 tunnel, enter commands such as the following on a Layer 3 Switch running both IPv4 and IPv6 protocol stacks on each end of the tunnel.

```
device(config)#interface tunnel 1
device(config-tnif-1)#tunnel source ethernet 1/3/1
device(config-tnif-1)#tunnel destination 10.162.100.1
device(config-tnif-1)#tunnel mode ipv6ip
device(config-tnif-1)#ipv6 enable
```

This example creates tunnel interface 1 and assigns a link local IPv6 address with an automatically computed EUI-64 interface ID to it. The IPv4 address assigned to Ethernet interface 1/3/1 is used as the tunnel source, while the IPv4 address 10.168.100.1 is configured as the tunnel destination. The tunnel mode is specified as a manual IPv6 tunnel. Finally, the tunnel is enabled. Note that instead of entering **ipv6 enable**, you could specify an IPv6 address, for example, **ipv6 address 2001:DB8:384d:34::/64 eui-64**, which would also enable the tunnel.

Syntax: **[no] interfacetunnel** *number*

For the *number* parameter, specify a value between 1-8.

Syntax: **[no] tunnelsource** *ipv4-address* | **ethernet** *port* | **loopback** *number* | **ve** *number*

The tunnel source can be an IP address or an interface.

For *ipv4-address*, use 8-bit values in dotted decimal notation.

The **ethernet** | **loopback** | **ve** parameter specifies an interface as the tunnel source. If you specify an Ethernet interface, also specify the port number associated with the interface. If you specify a loopback, VE, or interface, also specify the loopback, VE, or number, respectively.

Syntax: **[no] tunneldestination** *ipv4-address*

Specify the *ipv4-address* parameter using 8-bit values in dotted decimal notation.

Syntax: **[no] tunnelmode** **ipv6ip**

ipv6ip indicates that this is an IPv6 manual tunnel.

Syntax: **ipv6 enable**

The **ipv6 enable** command enables the tunnel. Alternatively, you could specify an IPv6 address, which would also enable the tunnel.

Syntax: **ipv6 address** *ipv6-prefix / prefix-length* [**eui-64**]

The **ipv6 address** command enables the tunnel. Alternatively, you could enter **ipv6 enable**, which would also enable the tunnel.

Specify the *ipv6-prefix* parameter in hexadecimal format using 16-bit values between colons as documented in RFC 2373.

Specify the *prefix-length* parameter as a decimal value. A slash mark (/) must follow the *ipv6-prefix* parameter and precede the *prefix-length* parameter. The **eu1-64** keyword configures the global address with an EUI-64 interface ID in the low-order 64 bits. The interface ID is automatically constructed in IEEE EUI-64 format using the interface's MAC address.

Clearing IPv6 tunnel statistics

You can clear statistics (reset all fields to zero) for all IPv6 tunnels or for a specific tunnel interface.

For example, to clear statistics for tunnel 1, enter the following command at the Privileged EXEC level or any of the Config levels of the CLI.

```
device#clear ipv6 tunnel 1
```

To clear statistics for all IPv6 tunnels, enter the following command.

```
device#clear ipv6 tunnel
```

Syntax: `clear ipv6 tunnel [number]`

The *number* parameter specifies the tunnel number.

Displaying IPv6 tunnel information

Use the commands in this section to display the configuration, status, and counters associated with IPv6 tunnels.

Displaying a summary of tunnel information

To display a summary of tunnel information, enter the following command at any level of the CLI.

```
device#show ipv6 tunnel
IP6 Tunnels
 Tunnel  Mode           Packet Received  Packet Sent
  1      configured      0                0
  2      configured      0                22419
```

Syntax: `show ipv6tunnel`

This display shows the following information.

TABLE 21 IPv6 tunnel summary information

Field	Description
Tunnel	The tunnel interface number.
Mode	The tunnel mode. Possible modes include the following: <ul style="list-style-type: none"> configured - Indicates a manually configured tunnel.
Packet Received	The number of packets received by a tunnel interface. Note that this is the number of packets received by the CPU. It does not include the number of packets processed in hardware.
Packet Sent	The number of packets sent by a tunnel interface. Note that this is the number of packets sent by the CPU. It does not include the number of packets processed in hardware.

Displaying tunnel interface information

To display status and configuration information for tunnel interface 1, enter the following command at any level of the CLI.

```
device#show interfaces tunnel 1
Tunnell is up, line protocol is up
Hardware is Tunnel
Tunnel source ve 30
Tunnel destination is 10.2.2.10
Tunnel mode ipv6ip
No port name
MTU 1480 bytes, encapsulation IPV4
```

Syntax: `show interfacestunnel number`

The *number* parameter indicates the tunnel interface number for which you want to display information.

TABLE 22 IPv6 tunnel interface information

Field	Description
Tunnel interface status	The status of the tunnel interface can be one of the following: <ul style="list-style-type: none"> • <i>up</i> - The tunnel mode is set and the tunnel interface is enabled. • <i>down</i> - The tunnel mode is not set. • <i>administratively down</i> - The tunnel interface was disabled with the disable command.
Line protocol status	The status of the line protocol can be one of the following: <ul style="list-style-type: none"> • <i>up</i> - IPv4 connectivity is established. • <i>down</i> - The line protocol is not functioning and is down.
Hardware is tunnel	The interface is a tunnel interface.
Tunnel source	The tunnel source can be one of the following: <ul style="list-style-type: none"> • An IPv4 address • The IPv4 address associated with an interface/port.
Tunnel destination	The tunnel destination can be an IPv4 address.
Tunnel mode	The tunnel mode can be the following: <ul style="list-style-type: none"> • <i>ipv6ip</i> - indicates a manually configured tunnel
Port name	The port name configured for the tunnel interface.
MTU	The setting of the IPv6 maximum transmission unit (MTU).

Displaying interface level IPv6 settings

To display Interface level IPv6 settings for tunnel interface 1, enter the following command at any level of the CLI.

```
device#show ipv6 inter tunnel 1
Interface Tunnel 1 is up, line protocol is up
IPv6 is enabled, link-local address is fe80::3:4:2 [Preferred]
Global unicast address(es):
 1001::1 [Preferred], subnet is 1001::/64
 1011::1 [Preferred], subnet is 1011::/64
Joined group address(es):
 ff02::1:ff04:2
 ff02::5
 ff02::1:ff00:1
 ff02::2
 ff02::1
MTU is 1480 bytes
ICMP redirects are enabled
No Inbound Access List Set
No Outbound Access List Set
OSPF enabled
```

The display command above reflects the following configuration.

```
device#show running-config interface tunnel 1
!
interface tunnel 1
port-name ManualTunnell
tunnel mode ipv6ip
tunnel source loopback 1
tunnel destination 10.1.1.1
ipv6 address 1011::1/64
ipv6 address 1001::1/64
ipv6 ospf area 0
```

TABLE 23 Interface level IPv6 tunnel information

Field	Description
Interface Tunnel status	<p>The status of the tunnel interface can be one of the following:</p> <ul style="list-style-type: none"> • up - IPv4 connectivity is established. • down - The tunnel mode is not set. • administratively down - The tunnel interface was disabled with the disable command.
Line protocol status	<p>The status of the line protocol can be one of the following:</p> <ul style="list-style-type: none"> • up - IPv6 is enabled through the ipv6 enable or ipv6 address command. • down - The line protocol is not functioning and is down.

SNMP Access

• SNMP overview.....	149
• SNMP community strings.....	149
• User-based security model.....	152
• Defining SNMP views.....	155
• SNMP version 3 traps.....	156
• Displaying SNMP Information.....	159
• SNMP v3 configuration examples.....	161

SNMP overview

SNMP is a set of protocols for managing complex networks. SNMP sends messages, called protocol data units (PDUs), to different parts of a network. SNMP-compliant devices, called agents, store data about themselves in Management Information Bases (MIBs) and return this data to the SNMP requesters.

"Security Access" chapter in the *FastIron Ethernet Switch Security Configuration Guide* introduced a few methods used to secure SNMP access. They included the following:

- Using ACLs to restrict SNMP access
- Restricting SNMP access to a specific IP address
- Restricting SNMP access to a specific VLAN
- Disabling SNMP access

This section presents additional methods for securing SNMP access to Brocade devices.

Restricting SNMP access using ACL, VLAN, or a specific IP address constitute the first level of defense when the packet arrives at a Brocade device. The next level uses one of the following methods:

- Community string match In SNMP versions 1 and 2
- User-based model in SNMP version 3

SNMP views are incorporated in community strings and the user-based model.

SNMP community strings

SNMP versions 1 and 2 use community strings to restrict SNMP access.

- To access a read-only management session using the Web Management Interface, enter the default username and password which are "get" and "public" respectively in the Web.
- To access a read-write management session using the Web Management Interface, configure a read-write community string using the CLI. Then log on using "set" as the user name and the read-write community string you configure as the password.

You can configure as many additional read-only and read-write community strings as you need. The number of strings you can configure depends on the memory on the device. There is no practical limit.

The Web Management Interface supports only one read-write session at a time. When a read-write session is open on the Web Management Interface, subsequent sessions are read-only, even if the session login is "set" with a valid read-write password.

NOTE

As an alternative to the SNMP community strings, you can secure Web management access using local user accounts or ACLs.

Encryption of SNMP community strings

The software automatically encrypts SNMP community strings. Users with read-only access or who do not have access to management functions in the CLI cannot display the strings. For users with read-write access, the strings are encrypted in the CLI but are shown in the clear in the Web Management Interface.

Encryption is enabled by default. You can disable encryption for individual strings or trap receivers if desired. Refer to the next section for information about encryption.

Adding an SNMP community string

You can assign SNMP community strings, and indicate if the string is encrypted or clear. By default, the string is encrypted.

To add an encrypted community string, enter commands such as the following.

```
device(config)# snmp-server community private rw
device(config)# write memory
```

Syntax: `snmp-server community [0 | 1] string ro | rw [view viewname] [standard-ACL-name | standard-ACL-id]`

The *string* parameter specifies the community string name. The string can be up to 32 characters long.

The **ro** | **rw** parameter specifies whether the string is **read-only (ro)** or **read-write (rw)**.

NOTE

If you issue a **no snmp-server community public ro** command and then enter a **write memory** command to save that configuration, the read-only "public" community string is removed and will have no SNMP access. If for some reason the device is brought down and then brought up, the "no snmp-server community public ro" command is restored in the system and the read-only "public" community string has no SNMP access.

The **0** | **1** parameter affects encryption for display of the string in the running-config and the startup-config file. Encryption is enabled by default. When encryption is enabled, the community string is encrypted in the CLI regardless of the access level you are using. In the Web Management Interface, the community string is encrypted at the read-only access level but is visible at the read-write access level.

The encryption option can be omitted (the default) or can be one of the following:

- **0** - Disables encryption for the community string you specify with the command. The community string is shown as clear text in the running-config and the startup-config file. Use this option if you do not want the display of the community string to be encrypted.
- **1** - Assumes that the community string you enter is encrypted, and decrypts the value before using it.

NOTE

If you want the software to assume that the value you enter is the clear-text form, and to encrypt display of that form, do not enter **0** or **1**. Instead, omit the encryption option and allow the software to use the default behavior.

NOTE

If you specify encryption option **1**, the software assumes that you are entering the encrypted form of the community string. In this case, the software decrypts the community string you enter before using the value for authentication. If you accidentally enter option **1** followed by the clear-text version of the community string, authentication will fail because the value used by the software will not match the value you intended to use.

The command in the example above adds the read-write SNMP community string "private". When you save the new community string to the startup-config file (using the **write memory** command), the software adds the following command to the file.

```
snmp-server community 1
encrypted-string
rw
```

To add a non-encrypted community string, you must explicitly specify that you do not want the software to encrypt the string. Here is an example.

```
device(config)#snmp-server community 0 private rw
device(config)#write memory
```

The command in this example adds the string "private" in the clear, which means the string is displayed in the clear. When you save the new community string to the startup-config file, the software adds the following command to the file.

```
snmp-server community 0 private rw
```

The *view viewname* parameter is optional. It allows you to associate a view to the members of this community string. Enter up to 32 alphanumeric characters. If no view is specified, access to the full MIB is granted. The view that you want must exist before you can associate it to a community string. Here is an example of how to use the view parameter in the community string command.

```
device(config)#snmp-s community myread ro view sysview
```

The command in this example associates the view "sysview" to the community string named "myread". The community string has read-only access to "sysview". For information on how to create views, refer to [SNMP v3 configuration examples](#) on page 161.

The *standard-ACL-name / standard-ACL-id* parameter is optional. It allows you to specify which ACL group will be used to filter incoming SNMP packets. You can enter either the ACL name or its ID. Here are some examples.

```
device(config)#snmp-s community myread ro view sysview 2
device(config)#snmp-s community myread ro view sysview myACL
```

The command in the first example indicates that ACL group 2 will filter incoming SNMP packets; whereas, the command in the second example uses the ACL group called "myACL" to filter incoming packets.

NOTE

To make configuration changes, including changes involving SNMP community strings, you must first configure a read-write community string using the CLI. Alternatively, you must configure another authentication method and log on to the CLI using a valid password for that method.

Displaying the SNMP community strings

To display the configured community strings, enter the following command at any CLI level.

```
device#show snmp server
Contact: Marshall
Location: Copy Center
Community(ro): public
Community(rw): private
Traps
    Cold start: Enable
    Link up: Enable
    Link down: Enable
    Authentication: Enable
    Locked address violation: Enable
    Power supply failure: Enable
    Fan failure: Enable
    Temperature warning: Enable
    STP new root: Enable
    STP topology change: Enable
    ospf: Enable
```

```

Total Trap-Receiver Entries: 4
Trap-Receiver IP Address      Community
1                             10.95.6.211
2                             10.95.5.21

```

Syntax: `show snmp server`

NOTE

If display of the strings is encrypted, the strings are not displayed. Encryption is enabled by default.

User-based security model

SNMP version 3 (RFC 2570 through 2575) introduces a User-Based Security model (RFC 2574) for authentication and privacy services.

SNMP version 1 and version 2 use community strings to authenticate SNMP access to management modules. This method can still be used for authentication. In SNMP version 3, the User-Based Security model of SNMP can be used to secure against the following threats:

- Modification of information
- Masquerading the identity of an authorized entity
- Message stream modification
- Disclosure of information

SNMP version 3 also supports View-Based Access Control Mechanism (RFC 2575) to control access at the PDU level. It defines mechanisms for determining whether or not access to a managed object in a local MIB by a remote principal should be allowed. For more information, refer to [SNMP v3 configuration examples](#) on page 161.)

Configuring your NMS

In order to use the SNMP version 3 features.

1. Make sure that your Network Manager System (NMS) supports SNMP version 3.
2. Configure your NMS agent with the necessary users.
3. Configure the SNMP version 3 features in Brocade devices.

Configuring SNMP version 3 on Brocade devices

Follow the steps given below to configure SNMP version 3 on Brocade devices.

1. Enter an engine ID for the management module using the `snmp-server engineid` command if you will not use the default engine ID. Refer to [Defining the engine id](#) on page 153.
2. Create views that will be assigned to SNMP user groups using the `snmp-server view` command. refer to [SNMP v3 configuration examples](#) on page 161 for details.
3. Create ACL groups that will be assigned to SNMP user groups using the `access-list` command.
4. Create user groups using the `snmp-server group` command. Refer to [Defining an SNMP group](#) on page 153.
5. Create user accounts and associate these accounts to user groups using the `snmp-server user` command. Refer to [Defining an SNMP user account](#) on page 154.

If SNMP version 3 is not configured, then community strings by default are used to authenticate access.

Defining the engine id

A default engine ID is generated during system start up. To determine what the default engine ID of the device is, enter the **show snmp engineid** command and find the following line:

```
Local SNMP Engine ID: 800007c70300e05290ab60
```

See the section [Displaying the Engine ID](#) on page 159 for details.

The default engine ID guarantees the uniqueness of the engine ID for SNMP version 3. If you want to change the default engine ID, enter the **snmp-server engineid local** command.

```
device(config)#snmp-server engineid local 800007c70300e05290ab60
```

Syntax: **[no] snmp-server engineid local** *hex-string*

The *local* parameter indicates that engine ID to be entered is the ID of this device, representing an SNMP management entity.

NOTE

Each user localized key depends on the SNMP server engine ID, so all users need to be reconfigured whenever the SNMP server engine ID changes.

NOTE

Since the current implementation of SNMP version 3 does not support Notification, remote engine IDs cannot be configured at this time.

The *hex-string* variable consists of 11 octets, entered as hexadecimal values. There are two hexadecimal characters in each octet. There should be an even number of hexadecimal characters in an engine ID.

The default engine ID has a maximum of 11 octets:

- Octets 1 through 4 represent the agent's SNMP management private enterprise number as assigned by the Internet Assigned Numbers Authority (IANA). The most significant bit of Octet 1 is "1". For example, "000007c7" is the ID for Brocade Communications, Inc. in hexadecimal. With Octet 1 always equal to "1", the first four octets in the default engine ID is always "800007c7" (which is 1991 in decimal).
- Octet 5 is always 03 in hexadecimal and indicates that the next set of values represent a MAC address.
- Octets 6 through 11 form the MAC address of the lowest port in the management module.

NOTE

Engine ID must be a unique number among the various SNMP engines in the management domain. Using the default engine ID ensures the uniqueness of the numbers.

Defining an SNMP group

SNMP groups map SNMP users to SNMP views. For each SNMP group, you can configure a read view, a write view, or both. Users who are mapped to a group will use its views for access control.

To configure an SNMP user group, enter a command such as the following.

```
device(config)#snmp-server group admin v3 auth read all write all
```

Syntax: **[no] snmp-server group** *groupname* { **v1** | **v2c** | **v3** { **auth** | **noauth** | **priv** } } [**access** { *standard-ACL-id* | **ipv6** *ipv6-ACL-name* }] [**read** *viewname*] [**write** *viewname*]

NOTE

This command is not used for SNMP version 1 and SNMP version 2. In these versions, groups and group views are created internally using community strings. (refer to [SNMP community strings](#) on page 149.) When a community string is created, two groups are created, based on the community string name. One group is for SNMP version 1 packets, while the other is for SNMP version 2 packets.

The *group groupname* parameter defines the name of the SNMP group to be created.

The **v1** , **v2c** , or **v3** parameter indicates which version of SNMP is used. In most cases, you will be using v3, since groups are automatically created in SNMP versions 1 and 2 from community strings.

The **auth** | **noauth** parameter determines whether or not authentication will be required to access the supported views. If **auth** is selected, then only authenticated packets are allowed to access the view specified for the user group. Selecting **noauth** means that no authentication is required to access the specified view. Selecting **priv** means that an authentication password will be required from the users.

The *access standard-ACL-id* parameter is optional. It allows incoming SNMP packets to be filtered based on the standard ACL attached to the group.

The **ipv6** *ipv6-ACL-name* option configures IPv6 ACL for SNMP group and allows incoming SNMP packets to be filtered based on the IPv6 ACL attached to the group.

The *read viewname* / *write viewname* parameter is optional. It indicates that users who belong to this group have either read or write access to the MIB.

The *viewname* variable is the name of the view to which the SNMP group members have access. If no view is specified, then the group has no access to the MIB.

The value of *viewname* is defined using the **snmp-server view** command. The SNMP agent comes with the "all" default view, which provides access to the entire MIB; however, it must be specified when creating the group. The "all" view also allows SNMP version 3 to be backwards compatible with SNMP version 1 and version 2.

NOTE

If you will be using a view other than the "all" view, that view must be configured before creating the user group. Refer to the section [SNMP v3 configuration examples](#) on page 161, especially for details on the include | exclude parameters.

Defining an SNMP user account

The **snmp-server user** command does the following:

- Creates an SNMP user.
- Defines the group to which the user will be associated.
- Defines the type of authentication to be used for SNMP access by this user.
- Specifies one of the following encryption types used to encrypt the privacy password:
 - Data Encryption Standard (DES) - A symmetric-key algorithm that uses a 56-bit key.
 - Advanced Encryption Standard (AES) - The 128-bit encryption standard adopted by the U.S. government. This standard is a symmetric cipher algorithm chosen by the National Institute of Standards and Technology (NIST) as the replacement for DES.

Here is an example of how to create an SNMP User account.

```
device(config)#snmp-s user bob admin v3 access 2 auth md5 bobmd5 priv des bobdes
```

The CLI for creating SNMP version 3 users has been updated as follows.

Syntax: `no snmp-server user name groupname v3 [[access standard-ACL-id] [encrypted] [auth md5 md5-password | sha sha-password] [priv [encrypted] des des-password-key | aes aes-password-key]]`

The *name* parameter defines the SNMP user name or security name used to access the management module.

The *groupname* parameter identifies the SNMP group to which this user is associated or mapped. All users must be mapped to an SNMP group. Groups are defined using the **snmp-server group** command.

NOTE

The SNMP group to which the user account will be mapped should be configured before creating the user accounts; otherwise, the group will be created without any views. Also, ACL groups must be configured before configuring user accounts.

The **v3** parameter is required.

The *access standard-ACL-id* parameter is optional. It indicates that incoming SNMP packets are filtered based on the ACL attached to the user account.

NOTE

The ACL specified in a user account overrides the ACL assigned to the group to which the user is mapped. If no ACL is entered for the user account, then the ACL configured for the group will be used to filter packets.

The *encrypted* parameter means that the MD5 or SHA password will be a digest value. MD5 has 16 octets in the digest. SHA has 20. The digest string has to be entered as a hexadecimal string. In this case, the agent need not generate any explicit digest. If the *encrypted* parameter is not used, the user is expected to enter the authentication password string for MD5 or SHA. The agent will convert the password string to a digest, as described in RFC 2574.

The **auth md5 | sha** parameter is optional. It defines the type of encryption that the user must have to be authenticated. Choose between MD5 or SHA encryption. MD5 and SHA are two authentication protocols used in SNMP version 3.

The **md5-password** and **sha-password** define the password the user must use to be authenticated. These password must have a minimum of 8 characters. If the encrypted parameter is used, then the digest has 16 octets for MD5 or 20 octets for SHA.

NOTE

Once a password string is entered, the generated configuration displays the digest (for security reasons), not the actual password.

The *priv [encrypted]* parameter is optional after you enter the md5 or sha password. The **priv** parameter specifies the encryption type (DES or AES) used to encrypt the privacy password. If the **encrypted** keyword is used, do the following:

- If DES is the privacy protocol to be used, enter **des** followed by a 16-octet DES key in hexadecimal format for the *des-password-key*. If you include the encrypted keyword, enter a password string of at least 8 characters.
- If AES is the privacy protocol to be used, enter **aes** followed by the AES password key. For a small password key, enter 12 characters. For a big password key, enter 16 characters. If you include the encrypted keyword, enter a password string containing 32 hexadecimal characters.

Defining SNMP views

SNMP views are named groups of MIB objects that can be associated with user accounts to allow limited access for viewing and modification of SNMP statistics and system configuration. SNMP views can also be used with other commands that take SNMP views as an argument. SNMP views reference MIB objects using object names, numbers, wildcards, or a combination of the three. The numbers represent the hierarchical location of the object in the MIB tree. You can reference individual objects in the MIB tree or a subset of objects from the MIB tree.

To configure the number of SNMP views available on the Brocade device, enter the following command.

```
device(config)#system-max view 15
```

Syntax: `system-maxview` *number-of-views*

This command specifies the maximum number of SNMPv2 and v3 views that can be configured on a device. The number of views can be from 10 - 65536. The default is 10 views.

To add an SNMP view, enter one of the following commands.

```
device(config)#snmp-server view Maynes system included
device(config)#snmp-server view Maynes system.2 excluded
device(config)#snmp-server view Maynes 2.3.*.6 included
device(config)#write mem
```

NOTE

The `snmp-server view` command supports the MIB objects as defined in RFC 1445.

Syntax: `[no] snmp-serverview` *name mib_tree* `included | excluded`

The *name* parameter can be any alphanumeric name you choose to identify the view. The names cannot contain spaces.

The *mib_tree* parameter is the name of the MIB object or family. MIB objects and MIB sub-trees can be identified by a name or by the numbers called Object Identifiers (OIDs) that represent the position of the object or sub-tree in the MIB hierarchy. You can use a wildcard (*) in the numbers to specify a sub-tree family.

The `included | excluded` parameter specifies whether the MIB objects identified by the *mib_family* parameter are included in the view or excluded from the view.

NOTE

All MIB objects are automatically excluded from any view unless they are explicitly included; therefore, when creating views using the `snmp-server view` command, indicate which portion of the MIB you want users to access.

For example, you may want to assign the view called "admin" a community string or user group. The "admin" view will allow access to the Brocade MIBs objects that begin with the 1.3.6.1.4.1.1991 object identifier. Enter the following command.

```
device(config)#snmp-server view admin 1.3.6.1.4.1.1991 included
```

You can exclude portions of the MIB within an inclusion scope. For example, if you want to exclude the snAgentSys objects, which begin with 1.3.6.1.4.1.1991.1.1.2 object identifier from the admin view, enter a second command such as the following.

```
device(config)#snmp-server view admin 1.3.6.1.4.1.1991.1.1.2 excluded
```

NOTE

Note that the exclusion is within the scope of the inclusion.

To delete a view, use the `no` parameter before the command.

SNMP version 3 traps

Brocade devices support SNMP notifications in SMlv2 format. This allows notifications to be encrypted and sent to the target hosts in a secure manner.

Defining an SNMP group and specifying which view is notified of traps

The SNMP group command allows configuration of a viewname for notification purpose, similar to the read and write view. The default viewname is "all", which allows access to the entire MIB.

To configure an SNMP user group, first configure SNMPv3 views using the **snmp-server view** command. Refer to [SNMP v3 configuration examples](#) on page 161. Then enter a command such as the following.

```
device(config)#snmp-server group admin v3 auth read all write all
notify all
```

Syntax: `[no] snmp-server group groupname { v1 | v2c | v3 { auth | noauth | priv } } [access { standard-ACL-id | ipv6 ipv6-ACL-name }] [notify viewname] [read viewname] [write viewname]`

The *group groupname* parameter defines the name of the SNMP group to be created.

The **v1** , **v2c** , or **v3** parameter indicates which version of SNMP to use. In most cases, you will use v3, since groups are automatically created in SNMP versions 1 and 2 from community strings.

The **auth** | **noauth** parameter determines whether or not authentication will be required to access the supported views. If **auth** is selected, then only authenticated packets are allowed to access the view specified for the user group. Selecting **noauth** means that no authentication is required to access the specified view. Selecting **priv** means that an authentication password will be required from the users.

The *access standard-ACL-id* allows incoming SNMP packets to be filtered based on the standard ACL attached to the group.

The **ipv6** *ipv6-ACL-name* option configures IPv6 ACL for SNMP group and allows incoming SNMP packets to be filtered based on the IPv6 ACL attached to the group.

The *read viewname* / *write viewname* parameter is optional. It indicates that users who belong to this group have either read or write access to the MIB.

The **notify** view allows administrators to restrict the scope of varbind objects that will be part of the notification. All of the varbinds need to be in the included view for the notification to be created.

The *viewname* variable is the name of the view to which the SNMP group members have access. If no view is specified, then the group has no access to the MIB.

Defining the UDP port for SNMP v3 traps

The SNMP host command enhancements allow configuration of notifications in SMIv2 format, with or without encryption, in addition to the previously supported SMIv1 trap format.

You can define a port that receives the SNMP v3 traps by entering a command such as the following.

```
device(config)#snmp-server host 192.168.4.11 version v3 auth security-name port 4/1
```

Syntax: `[no] snmp-server host ip-addr | ipv6-addr version [v1 | v2c community-string | v3 auth | noauth | priv security-name] [port trap-UDP-port-number]`

The *ip-addr* parameter specifies the IP address of the host that will receive the trap.

For *version* , indicate one of the following

For SNMP version 1, enter **v1** and the name of the community string (*community-string*). This string is encrypted within the system.

NOTE

If the configured version is v2c, then the notification is sent out in SMIv2 format, using the community string, but in cleartext mode. To send the SMIv2 notification in SNMPv3 packet format, configure v3 with auth or privacy parameters, or both, by specifying a security name. The actual authorization and privacy values are obtained from the security name.

For SNMP version 2c, enter **v2** and the name of the community string. This string is encrypted within the system.

For SNMP version 3, enter one of the following depending on the authorization required for the host:

- *v3 auth security-name* : Allow only authenticated packets.
- *v3 no auth security-name* : Allow all packets.
- *v3 priv security-name* : A password is required

For *port trap-UDP-port-number* , specify the UDP port number on the host that will receive the trap.

Trap MIB changes

To support the SNMP V3 trap feature, the Brocade Enterprise Trap MIB was rewritten in SMIv2 format, as follows:

- The MIB name was changed from FOUNDRY-SN-TRAP-MIB to FOUNDRY-SN-NOTIFICATION-MIB
- Individual notifications were changed to NOTIFICATION-TYPE instead of TRAP-TYPE.
- As per the SMIv2 format, each notification has an OID associated with it. The root node of the notification is snTraps (OID enterprise.foundry.0). For example, OID for snTrapRunningConfigChanged is {snTraps.73}. Earlier, each trap had a trap ID associated with it, as per the SMIv1 format.

Backward compatibility with SMIv1 trap format

The Brocade device will continue to support creation of traps in SMIv1 format, as before. To allow the device to send notifications in SMIv2 format, configure the device as described above. The default mode is still the original SMIv1 format.

Specifying an IPv6 host as an SNMP trap receiver

You can specify an IPv6 host as a trap receiver to ensure that all SNMP traps sent by the device will go to the same SNMP trap receiver or set of receivers, typically one or more host devices on the network. To do so, enter a command such as the following.

```
device(config)#snmp-server host ipv6 2001:DB8:89::13
```

Syntax: `snmp-serverhost ipv6 ipv6-address`

The *ipv6-address* must be in hexadecimal format using 16-bit values between colons as documented in RFC 2373.

SNMP v3 over IPv6

Some FastIron devices support IPv6 for SNMP version 3.

Restricting SNMP Access to an IPv6 Node

You can restrict SNMP access so that the Brocade device can only be accessed by the IPv6 host address that you specify. To do so, enter a command such as the following .

```
device(config)#snmp-client ipv6 2001:DB8:89::23
```

Syntax: `snmp-clientipv6 ipv6-address`

The *ipv6-address* must be in hexadecimal format using 16-bit values between colons as documented in RFC 2373.

Specifying an IPv6 host as an SNMP trap receiver

You can specify an IPv6 host as a trap receiver to ensure that all SNMP traps sent by the Brocade device will go to the same SNMP trap receiver or set of receivers, typically one or more host devices on the network. To do so, enter the **snmp-server host ipv6** command .

```
device(config)#snmp-server host ipv6 2001:DB8:89::13
```

Syntax: **snmp-server host ipv6** *ipv6-address*

The *ipv6-address* must be in hexadecimal format using 16-bit values between colons as documented in RFC 2373.

Viewing IPv6 SNMP server addresses

Many of the existing **show** commands display IPv6 addresses for IPv6 SNMP servers. The following example shows output for the **show snmp server** command.

```
device#show snmp server
Contact:
Location:
Community(ro): .....
Traps
    Warm/Cold start: Enable
    Link up: Enable
    Link down: Enable
    Authentication: Enable
    Locked address violation: Enable
    Power supply failure: Enable
    Fan failure: Enable
    Temperature warning: Enable
    STP new root: Enable
    STP topology change: Enable
    vsrp: Enable
Total Trap-Receiver Entries: 4
Trap-Receiver IP-Address      Port-Number Community
1      10.147.201.100          162      .....
2      2001:DB8::200          162      .....
3      10.147.202.100          162      .....
4      2001:DB8::200          162      .....
```

Displaying SNMP Information

This section lists the commands for viewing SNMP-related information.

Displaying the Engine ID

To display the engine ID of a management module, enter a command such as the following.

```
device#show snmp engineid
Local SNMP Engine ID: 800007c70300e05290ab60
Engine Boots: 3
Engine time: 5
```

Syntax: **show snmp engineid**

The engine ID identifies the source or destination of the packet.

The engine boots represents the number of times that the SNMP engine reinitialized itself with the same engine ID. If the engineID is modified, the boot count is reset to 0.

The engine time represents the current time with the SNMP agent.

Displaying SNMP groups

To display the definition of an SNMP group, enter a command such as the following.

```
device#show snmp group
groupname = exceptifgrp
security model = v3
security level = authNoPriv
ACL id = 0
IPv6 ACL name: ipv6acl
readview = exceptif
writeview =
none
```

Syntax: show snmp group

The value for security level can be one of the following.

Security level	Authentication
none	If the security model shows v1 or v2, then security level is blank. User names are not used to authenticate users; community strings are used instead.
noauthNoPriv	Displays if the security model shows v3 and user authentication is by user name only.
authNoPriv	Displays if the security model shows v3 and user authentication is by user name and the MD5 or SHA algorithm.

Displaying user information

To display the definition of an SNMP user account, enter a command such as the following.

```
device#show snmp user
username = bob
ACL id = 2
group = admin
security model = v3
group ACL id = 0
authtype = md5
authkey = 3aca18d90b8d172760e2dd2e8f59b7fe
privtype = des, privkey = 1088359afb3701730173a6332d406eec
engine ID= 800007c70300e052ab0000
```

Syntax: show snmp user

Interpreting varbinds in report packets

If an SNMP version 3 request packet is to be rejected by an SNMP agent, the agent sends a report packet that contains one or more varbinds. The varbinds contain additional information, showing the cause of failures. An SNMP manager application decodes the description from the varbind. The following table presents a list of varbinds supported by the SNMP agent.

Varbind object Identifier	Description
1.3.6.1.6.3.11.2.1.3.0	Unknown packet data unit.
1.3.6.1.6.3.12.1.5.0	The value of the varbind shows the engine ID that needs to be used in the snmp-server engineid command
1.3.6.1.6.3.15.1.1.1.0	Unsupported security level.
1.3.6.1.6.3.15.1.1.2.0	Not in time packet.
1.3.6.1.6.3.15.1.1.3.0	Unknown user name. This varbind may also be generated: <ul style="list-style-type: none"> • If the configured ACL for this user filters out this packet. • If the group associated with the user is unknown.
1.3.6.1.6.3.15.1.1.4.0	Unknown engine ID. The value of this varbind would be the correct authoritative engineID that should be used.
1.3.6.1.6.3.15.1.1.5.0	Wrong digest.
1.3.6.1.6.3.15.1.1.6.0	Decryption error.

SNMP v3 configuration examples

The following sections present examples of how to configure SNMP v3.

Example 1

```
device(config)#snmp-s group admingrp v3 priv read all write all notify all
device(config)#snmp-s user adminuser admingrp v3 auth md5
auth password
  priv
privacy password
device(config)#snmp-s host
dest-ip
  version v3 privacy adminuser
```

Example 2

```
device(config)#snmp-server view internet internet included
device(config)#snmp-server view system system included
device(config)#snmp-server community ..... ro
device(config)#snmp-server community ..... rw
device(config)#snmp-server contact isc-operations
device(config)#snmp-server location sdh-pillbox
device(config)#snmp-server host 128.91.255.32 .....
device(config)#snmp-server group ops v3 priv read internet write system
device(config)#snmp-server group admin v3 priv read internet write internet
device(config)#snmp-server group restricted v3 priv read internet
device(config)#snmp-server user ops ops v3 encrypted auth md5 ab8e9cd6d46e7a270b8c9549d92a069 priv
encrypted des 0e1b153303b6188089411447dbc32de
device(config)#snmp-server user admin admin v3 encrypted auth md5 0d8a2123f91bfbdb8695fef16a6f4207b priv
encrypted des 18e0cf359fce4fcd60df19c2b6515448
device(config)#snmp-server user restricted restricted v3 encrypted auth md5
261fd8f56a3ad51c8bcecle4609f54dc priv encrypted des d32e66152f89de9b2e0cb17a65595f43
```


Foundry Discovery Protocol (FDP) and Cisco Discovery Protocol (CDP) Packets

- FDP overview.....163
- CDP packets.....167

FDP overview

The Foundry Discovery Protocol (FDP) enables Brocade devices to advertise themselves to other Brocade devices on the network. When you enable FDP on a Brocade device, the device periodically advertises information including the following:

- Hostname (device ID)
- Product platform and capability
- Software version
- VLAN and Layer 3 protocol address information for the port sending the update. IP, IPX, and AppleTalk Layer 3 information is supported.

A Brocade device running FDP sends FDP updates on Layer 2 to MAC address 00-00-00-CC-CC-CC. Other Brocade devices listening on that address receive the updates and can display the information in the updates. Brocade devices can send and receive FDP updates on ethernet interfaces.

FDP is disabled by default.

NOTE

If FDP is not enabled on a Brocade device that receives an FDP update or the device is running a software release that does not support FDP, the update passes through the device at Layer 2.

FDP configuration

The following sections describe how to enable Foundry Discovery Protocol (FDP) and how to change the FDP update and hold timers.

Enabling FDP globally

To enable a Brocade device to globally send FDP packets, enter the following command at the global CONFIG level of the CLI.

```
device(config)# fdprun
```

Syntax: [no] fdprun

The feature is disabled by default.

Enabling FDP at the interface level

By default, FDP is enabled at the interface level after FDP is enabled on the device. When FDP is enabled globally, you can disable and re-enable FDP on individual ports.

Disable FDP by entering commands such as the following:

```
device(config)# int e 1/2/1
device(config-if-1/2/1)# no fdp enable
```

Enable or re-enable FDP by entering commands such as the following:

```
device(config-if-1/2/1)# fdp enable
```

Syntax: [no] fdp enable

Specifying the IP management address to advertise

When FDP is enabled, by default, the Brocade device advertises one IPv4 address and one IPv6 address to its FDP neighbors. You can configure the device to advertise only the IPv4 management address or only the IPv6 management address.

You can set the configuration globally on a Layer 2 switch, or on an interface on a Layer 3 switch.

For example, to configure a Layer 2 switch to advertise the IPv4 address, enter the following command at the Global CONFIG level of the CLI:

```
device(config)# fdp advertise ipv4
```

To configure a Layer 3 switch to advertise the IPv6 address, enter the following command at the Interface level of the CLI:

```
device(config-if-1/2/1)# fdp advertise ipv6
```

Syntax: fdp advertise ipv4 | ipv6

Changing the FDP update timer

By default, a Brocade device enabled for FDP sends an FDP update every 60 seconds. You can change the update timer to a value from 5 - 900 seconds.

To change the FDP update timer, enter a command such as the following at the global CONFIG level of the CLI.

```
device(config)# fdp timer 120
```

Syntax: [no] fdp timer secs

The secs parameter specifies the number of seconds between updates and can be from 5 - 900 seconds. The default is 60 seconds.

Changing the FDP hold time

By default, a Brocade device that receives an FDP update holds the information until one of the following events occurs:

- The device receives a new update.
- 180 seconds have passed since receipt of the last update. This is the hold time.

Once either of these events occurs, the device discards the update.

To change the FDP hold time, enter the **fdp holdtime** command at the global CONFIG level of the CLI.

```
device(config)# fdp holdtime 360
```

Syntax: [no] fdp holdtime secs

The `secs` parameter specifies the number of seconds a Brocade device that receives an FDP update can hold the update before discarding it. You can specify from 10 - 255 seconds. The default is 180 seconds.

Displaying FDP information

You can display the following Foundry Discovery Protocol (FDP) information:

- FDP entries for Brocade neighbors
- Individual FDP entries
- FDP information for an interface on the device you are managing
- FDP packet statistics

NOTE

If the Brocade device has intercepted CDP updates, then the CDP information is also displayed.

Displaying neighbor information

To display a summary list of all the Brocade neighbors that have sent FDP updates to this Brocade device, enter the **show fdp neighbors** command.

```
device# show fdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater
(*) indicates a CDP device
-----
Device ID      Local Int      Holdtm  Capability Platform      Port ID
-----
FastIronB     Eth 1/2/9     178    Router    FastIron Rou  Eth 1/2/9
```

Syntax: `show fdp neighbors [ethernet port] [detail]`

The `port` variable specifies the display of information for updates received on the specified port.

The `detail` parameter specifies the display of detailed information for each device.

The **show fdp neighbors** command, without optional parameters, displays the following information.

TABLE 24 Summary FDP and CDP neighbor information

This line...	Displays...
Device ID	The hostname of the neighbor.
Local Int	The interface on which this Brocade device received an FDP or CDP update for the neighbor.
Holdtm	The maximum number of seconds this device can keep the information received in the update before discarding it.
Capability	The role that the neighbor is capable of playing in the network.
Platform	The product platform of the neighbor device.
Port ID	The interface through which the neighbor sent the update.

To display detailed information, enter the **show fdp neighbors detail** command.

```
deviceA# show fdp neighbors detail
Device ID: FastIronB configured as default VLAN1, tag-type8100
Entry address(es):
  IP address: 192.168.0.13
  IPv6 address (Global): c:a:f:e:c:a:f:e
Platform: FastIron Router, Capabilities: Router
Interface: Eth 1/2/9
```

```

Port ID (outgoing port): Eth 1/2/9 is TAGGED in following VLAN(s):
 9 10 11
Holdtime : 176 seconds
Version :
Foundry, Inc. Router, IronWare Version 07.6.01b1T53 Compiled on Aug 29
2002 at 10:35:21 labeled as B2R07601b1

```

The **show fdp neighbors detail** command displays the following information.

TABLE 25 Detailed FDP and CDP neighbor information

Parameter	Definition
Device ID	The hostname of the neighbor. In addition, this line lists the VLAN memberships and other VLAN information for the neighbor port that sent the update to this device.
Entry address(es)	The Layer 3 protocol addresses configured on the neighbor port that sent the update to this device. If the neighbor is a Layer 2 Switch, this field lists the management IP address.
Platform	The product platform of the neighbor.
Capabilities	The role that the neighbor is capable of playing in the network.
Interface	The interface on which this device received an FDP or CDP update for the neighbor.
Port ID	The interface through which the neighbor sent the update.
Holdtime	The maximum number of seconds this device can keep the information received in the update before discarding it.
Version	The software version running on the neighbor.

Displaying FDP entries

To display the detailed neighbor information for a specific device, enter the **show fdp entry** command.

```

deviceA# show fdp entry FastIronB
Device ID: FastIronB configured as default VLAN1, tag-type8100
Entry address(es):
Platform: FastIron Router, Capabilities: Router
Interface: Eth 1/2/9
Port ID (outgoing port): Eth 1/2/9 is TAGGED in following VLAN(s):
 9 10 11
Holdtime : 176 seconds
Version :
Foundry, Inc. Router, IronWare Version 07.6.01b1T53 Compiled on Aug 29
2002 at 10:35:21 labeled as B2R07601b1

```

Syntax: `show fdp entry { * | device-id }`

The `*` parameter specifies the display of detailed updates for all neighbor devices.

The `device-id` variable specifies the display of update information for the specified device.

For information about the display, refer to [Displaying neighbor information](#) on page 165.

Displaying FDP information for an interface

To display FDP information for an interface, enter a command such as the following.

```

deviceA# show fdp interface ethernet 1/2/3
FastEthernet1/2/3 is up, line protocol is up
Encapsulation ethernet
Sending FDP packets every 5 seconds
Holdtime is 180 seconds

```

This example shows information for a specific ethernet port indicating how often the port sends FDP updates and how long neighbors that receive the updates, can hold them before discarding them.

Syntax: `show fdp interface [ethernet port]`

The optional *port* parameter lists only the information for the specified interface.

Displaying FDP and CDP statistics

To display FDP and CDP packet statistics, enter the following command.

```
deviceA# show fdp traffic
CDP/FDP counters:
  Total packets output: 6, Input: 5
  Hdr syntax: 0, Chksum error: 0, Encaps failed: 0
  No memory: 0, Invalid packet: 0, Fragmented: 0
  Internal errors: 0
```

Syntax: `show fdp traffic`

Clearing FDP and CDP information

You can clear the following FDP and CDP information:

- Information received in FDP and CDP updates
- FDP and CDP statistics

The same commands clear information for both FDP and CDP.

Clearing FDP and CDP neighbor information

To clear the information received in FDP and CDP updates from neighboring devices, enter the following command.

```
device# clear fdp table
```

Syntax: `clear fdp table`

NOTE

This command clears all the updates for FDP and CDP.

Clearing FDP and CDP statistics

To clear FDP and CDP statistics, enter the following command.

```
device# clear fdp counters
```

Syntax: `clear fdp counters`

CDP packets

Cisco Discovery Protocol (CDP) packets are used by Cisco devices to advertise themselves to other Cisco devices. By default, Brocade devices forward these packets without examining their contents. You can configure a Brocade device to intercept and display the contents of CDP packets. This feature is useful for learning device and interface information for Cisco devices in the network.

Brocade devices support intercepting and interpreting CDP version 1 and CDP version 2 packets.

NOTE

The Brocade device can interpret only the information fields that are common to both CDP version 1 and CDP version 2.

NOTE

When you enable interception of CDP packets, the Brocade device drops the packets. As a result, Cisco devices will no longer receive the packets.

Enabling interception of CDP packets globally

To enable the device to intercept and display CDP packets, enter the following command at the global CONFIG level of the CLI.

```
device(config)# cdp run
```

Syntax: [no] cdprun

The feature is disabled by default.

Enabling interception of CDP packets on an interface

You can disable and enable CDP at the interface level. By default, CDP is enabled on an interface once CDP is enabled on the device.

To enable CDP on an interface enter a command such as the following.

```
device(config)# int e 1/2/1
device(config-if-1/2/1)# cdp enable
```

Syntax: [no] cdp enable

Displaying CDP information

You can display the following CDP information:

- Cisco neighbors
- CDP entries for all Cisco neighbors or a specific neighbor
- CDP packet statistics

Displaying neighbors

To display the Cisco neighbors the Brocade device has learned from CDP packets, enter the **show fdp neighbors** command.

```
device# show fdp neighbors

Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater
(*) indicates a Cisco device
-----
Device ID      Local Int   Holdtm Capability Platform   Port ID
-----
(*)Router      Eth 1/1    124    R          cisco RSP4
FastEthernet5/0/0
```

To display detailed information for the neighbors, enter the **show fdp neighbors detail** command.

```
device# show fdp neighbors detail

Device ID: Router
Entry address(es):
  IP address: 10.95.6.143
Platform: cisco RSP4, Capabilities: Router
```



```

Interface: Eth 1/1, Port ID (outgoing port): FastEthernet5/0/0
Holdtime : 150 seconds
Version :
Cisco Internetwork Operating System Software
IOS (tm) RSP Software (RSP-JSV-M), Version 12.0(5)T1, RELEASE SOFTWARE
(fc1)
Copyright (c) 1986-1999 by cisco Systems, Inc.
Compiled Thu 19-Aug-99 04:12 by cmong

```

To display information about a neighbor attached to a specific port, enter a command such as the following.

```

device# show fdp neighbors ethernet 1/1
Device ID: Router
Entry address(es):
  IP address: 10.95.6.143
Platform: cisco RSP4, Capabilities: Router
Interface: Eth 1/1, Port ID (outgoing port): FastEthernet5/0/0
Holdtime : 127 seconds
Version :
Cisco Internetwork Operating System Software
IOS (tm) RSP Software (RSP-JSV-M), Version 12.0(5)T1, RELEASE SOFTWARE
(fc1)
Copyright (c) 1986-1999 by cisco Systems, Inc.
Compiled Thu 19-Aug-99 04:12 by cmong

```

Syntax: `show fdp neighbors [detail | ethernet port]`

Displaying CDP entries

To display CDP entries for all neighbors, enter the `show fdp entry` command.

```

device# show fdp entry *
Device ID: Router
Entry address(es):
  IP address: 10.95.6.143
Platform: cisco RSP4, Capabilities: Router
Interface: Eth 1/1, Port ID (outgoing port): FastEthernet5/0/0
Holdtime : 124 seconds
Version :
Cisco Internetwork Operating System Software
IOS (tm) RSP Software (RSP-JSV-M), Version 12.0(5)T1, RELEASE SOFTWARE
(fc1)
Copyright (c) 1986-1999 by cisco Systems, Inc.
Compiled Thu 19-Aug-99 04:12 by cmong

```

To display CDP entries for a specific device, specify the device ID, as shown in the following example.

```

device# show fdp entry Router1
Device ID: Router1
Entry address(es):
  IP address: 10.95.6.143
Platform: cisco RSP4, Capabilities: Router
Interface: Eth 1/1, Port ID (outgoing port): FastEthernet5/0/0
Holdtime : 156 seconds
Version :
Cisco Internetwork Operating System Software
IOS (tm) RSP Software (RSP-JSV-M), Version 12.0(5)T1, RELEASE SOFTWARE
(fc1)
Copyright (c) 1986-1999 by cisco Systems, Inc.
Compiled Thu 19-Aug-99 04:12 by cmong

```

Syntax: `show fdp entry * | device-id`

Displaying CDP statistics

To display CDP packet statistics, enter the **show fdp traffic** command.

```
device# show fdp traffic
CDP counters:
  Total packets output: 0, Input: 3
  Hdr syntax: 0, Chksum error: 0, Encaps failed: 0
  No memory: 0, Invalid packet: 0, Fragmented: 0
```

Syntax: show fdp traffic

Clearing CDP information

You can clear the following CDP information:

- Cisco neighbor information
- CDP statistics

To clear the Cisco neighbor information, enter the **clear fdp table** command.

```
device# clear fdp table
```

Syntax: clear fdptable

To clear CDP statistics, enter the following command.

```
device# clear fdp counters
```

Syntax: clear fdp counters

LLDP and LLDP-MED

• LLDP terms used in this chapter.....	171
• LLDP overview.....	172
• LLDP-MED overview.....	173
• General LLDP operating principles.....	175
• MIB support.....	179
• Syslog messages.....	180
• LLDP configuration.....	180
• LLDP-MED configuration.....	192
• LLDP-MED attributes advertised by the Brocade device.....	201
• Resetting LLDP statistics.....	209
• Clearing cached LLDP neighbor information.....	209

LLDP terms used in this chapter

Endpoint device - An LLDP-MED device located at the network edge, that provides some aspect of IP communications service based on IEEE 802 LAN technology. An Endpoint device is classified in one of three class types (I, II, or III) and can be an IP telephone, softphone, VoIP gateway, or conference bridge, among others.

Link Layer discovery protocol (LLDP) - The Layer 2 network discovery protocol described in the IEEE 802.1AB standard, *Station and Media Access Control Connectivity Discovery*. This protocol enables a station to advertise its capabilities to, and to discover, other LLDP-enabled stations in the same 802 LAN segments.

LLDP agent - The protocol entity that implements LLDP for a particular IEEE 802 device. Depending on the configured LLDP operating mode, an LLDP agent can send and receive LLDP advertisements (frames), or send LLDP advertisements only, or receive LLDP advertisements only.

LLDP media endpoint devices (LLDP-MED) - The Layer 2 network discovery protocol extension described in the ANSI/TIA-1057 standard, *LLDP for Media Endpoint Devices*. This protocol enables a switch to configure and manage connected Media Endpoint devices that need to send media streams across the network (for example, IP telephones and security cameras).

LLDPDU (LLDP Data Unit) - A unit of information in an LLDP packet that consists of a sequence of short variable length information elements, known as **TLVs**. LLDP pass-through is not supported in conformance to IEEE standard.

MIB (Management Information Base) - A virtual database that identifies each manageable object by its name, syntax, accessibility, and status, along with a text description and unique object identifier (OID). The database is accessible by a Network Management Station (NMS) using a management protocol such as the Simple Network Management Protocol (SNMP).

Network connectivity device - A forwarding 802 LAN device, such as a router, switch, or wireless access point.

Station - A node in a network.

TLV (Type-Length-Value) - An information element in an LLDPDU that describes the type of information being sent, the length of the information string, and the value (actual information) that will be transmitted.

TTL (Time-to-Live) - Specifies the length of time that the receiving device should maintain the information acquired through LLDP in its MIB.


LLDP overview

LLDP enables a station attached to an IEEE 802 LAN/MAN to advertise its capabilities to, and to discover, other stations in the same 802 LAN segments.

The information distributed by LLDP (the advertisement) is stored by the receiving device in a standard Management Information Base (MIB), accessible by a Network Management System (NMS) using a management protocol such as the Simple Network Management Protocol (SNMP). The information also can be viewed from the CLI, using **show LLDP** commands.

The following diagram illustrates LLDP connectivity

FIGURE 5 LLDP connectivity



Benefits of LLDP

LLDP provides the following benefits:


- Network Management:
 - Simplifies the use of and enhances the ability of network management tools in multi-vendor environments
 - Enables discovery of accurate physical network topologies such as which devices are neighbors and through which ports they connect
 - Enables discovery of stations in multi-vendor environments
- Network Inventory Data:
 - Supports optional system name, system description, system capabilities and management address
 - System description can contain the device product name or model number, version of hardware type, and operating system
 - Provides device capability, such as switch, router, or WLAN access point
- Network troubleshooting:
 - Information generated by LLDP can be used to detect speed and duplex mismatches
 - Accurate topologies simplify troubleshooting within enterprise networks
 - Can discover devices with misconfigured or unreachable IP addresses

LLDP-MED overview

LLDP-MED is an extension to LLDP. This protocol enables advanced LLDP features in a Voice over IP (VoIP) network. Whereas LLDP enables network discovery between Network Connectivity devices, LLDP-MED enables network discovery between Network Connectivity devices and media Endpoints such as, IP telephones, softphones, VoIP gateways and conference bridges.

The following diagram illustrates LLDP-MED connectivity.

FIGURE 6 LLDP-MED connectivity



Benefits of LLDP-MED

LLDP-MED provides the following benefits:

- Vendor-independent management capabilities, enabling different IP telephony systems to interoperate in one network.
- Automatically deploys network policies, such as Layer 2 and Layer 3 QoS policies and Voice VLANs.
- Supports E-911 Emergency Call Services (ECS) for IP telephony
- Collects Endpoint inventory information
- Network troubleshooting
 - Helps to detect improper network policy configuration

LLDP-MED class

An LLDP-MED class specifies an Endpoint type and its capabilities. An Endpoint can belong to one of three LLDP-MED class types:

- **Class 1 (Generic endpoint)** - A Class 1 Endpoint requires basic LLDP discovery services, but does not support IP media nor does it act as an end-user communication appliance. A Class 1 Endpoint can be an IP communications controller, other communication-related server, or other device requiring basic LLDP discovery services.
- **Class 2 (Media endpoint)** - A Class 2 Endpoint supports media streams and may or may not be associated with a particular end user. Device capabilities include media streaming, as well as all of the capabilities defined for Class 1 Endpoints. A Class 2 Endpoint can be a voice/media gateway, conference, bridge, media server, etc.
- **Class 3 (Communication endpoint)** - A Class 3 Endpoint supports end user IP communication. Capabilities include aspects related to end user devices, as well as all of the capabilities defined for Class 1 and Class 2 Endpoints. A Class 3 Endpoint can be an IP telephone, softphone (PC-based phone), or other communication device that directly supports the end user.

Discovery services defined in Class 3 include location identifier (ECS/E911) information and inventory management.

The LLDP-MED device class is advertised when LLDP-MED is enabled on a port.

General LLDP operating principles

LLDP and LLDP-MED use the services of the Data Link sublayers, Logical Link Control and Media Access Control, to transmit and receive information to and from other LLDP Agents (protocol entities that implement LLDP).

LLDP is a one-way protocol. An LLDP agent can transmit and receive information to and from another LLDP agent located on an adjacent device, but it cannot solicit information from another LLDP agent, nor can it acknowledge information received from another LLDP agent.

LLDP operating modes

When LLDP is enabled on a global basis, by default, each port on the Brocade device will be capable of transmitting and receiving LLDP packets. You can disable a port's ability to transmit and receive LLDP packets, or change the operating mode to one of the following:

- Transmit LLDP information only
- Receive LLDP information only

LLDP transmit mode

An LLDP agent sends LLDP packets to adjacent LLDP-enabled devices. The LLDP packets contain information about the transmitting device and port.

An LLDP agent initiates the transmission of LLDP packets whenever the transmit countdown timing counter expires, or whenever LLDP information has changed. When a transmit cycle is initiated, the LLDP manager extracts the MIB objects and formats this information into TLVs. The TLVs are inserted into an LLDPDU, addressing parameters are prepended to the LLDPDU, and the information is sent out LLDP-enabled ports to adjacent LLDP-enabled devices.

LLDP receive mode

An LLDP agent receives LLDP packets from adjacent LLDP-enabled devices. The LLDP packets contain information about the transmitting device and port.


When an LLDP agent receives LLDP packets, it checks to ensure that the LLDPDUs contain the correct sequence of mandatory TLVs, then validates optional TLVs. If the LLDP agent detects any errors in the LLDPDUs and TLVs, it drops them in software. TLVs that are not recognized but do not contain basic formatting errors, are assumed to be valid and are assigned a temporary identification index and stored for future possible alter retrieval by network management. All validated TLVs are stored in the neighbor database.

LLDP packets

LLDP agents transmit information about a sending device/port in packets called LLDP Data Units (LLDPDUs). All the LLDP information to be communicated by a device is contained within a single 1500 byte packet. A device receiving LLDP packets is not permitted to combine information from multiple packets.

As shown in the following figure, each LLDPDU has three mandatory TLVs, an End of LLDPDU TLV, plus optional TLVs as selected by network management.

FIGURE 7 LLDPDU packet format



M = mandatory TLV (required for all LLDPDUs)

Each LLDPDU consists of an untagged Ethernet header and a sequence of short, variable length information elements known as type, length, value (TLV).

TLVs have Type, Length, and Value fields, where:

- **Type** identifies the kind of information being sent
- **Length** indicates the length (in octets) of the information string
- **Value** is the actual information being sent (for example, a binary bit map or an alpha-numeric string containing one or more fields).

TLV support

This section lists the LLDP and LLDP-MED TLV support.

LLDP TLVs

There are two types of LLDP TLVs, as specified in the IEEE 802.3AB standard.

Basic management TLVs consist of both optional general system information TLVs as well as mandatory TLVs.

Mandatory TLVs cannot be manually configured. They are always the first three TLVs in the LLDPDU, and are part of the packet header.

General system information TLVs are optional in LLDP implementations and are defined by the Network Administrator.

Brocade devices support the following Basic Management TLVs:

- Chassis ID (mandatory)

- Port ID (mandatory)
- Time to Live (mandatory)
- Port description
- System name
- System description
- System capabilities
- Management address
- End of LLDPDU

Organizationally-specific TLVs are optional in LLDP implementations and are defined and encoded by individual organizations or vendors. These TLVs include support for, but are not limited to, the IEEE 802.1 and 802.3 standards and the TIA-1057 standard.

Brocade devices support the following Organizationally-specific TLVs:

- **802.1 organizationally-specific TLVs**
 - Port VLAN ID
 - VLAN name TLV
- **802.3 organizationally-specific TLVs**
 - MAC/PHY configuration/status
 - Power through MDI
 - Link aggregation
 - Maximum frame size

LLDP-MED TLVs

Brocade devices honor and send the following LLDP-MED TLVs, as defined in the TIA-1057 standard:

- LLDP-MED capabilities
- Network policy
- Location identification
- Extended power-via-MDI

Mandatory TLVs

When an LLDP agent transmits LLDP packets to other agents in the same 802 LAN segments, the following mandatory TLVs are always included:

- Chassis ID
- Port ID
- Time to Live (TTL)

This section describes the above TLVs in detail.

Chassis ID

The Chassis ID identifies the device that sent the LLDP packets.

There are several ways in which a device may be identified. A chassis ID subtype, included in the TLV and shown in the following table, indicates how the device is being referenced in the Chassis ID field.

TABLE 26 Chassis ID subtypes

ID subtype	Description
0	Reserved
1	Chassis component
2	Interface alias
3	Port component
4	MAC address
5	Network address
6	Interface name
7	Locally assigned
8 - 255	Reserved

Brocade devices use chassis ID subtype 4, the base MAC address of the device. Other third party devices may use a chassis ID subtype other than 4. The chassis ID will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
Chassis ID (MAC address): 0000.0033.e2c0
```

The chassis ID TLV is always the first TLV in the LLDPDU.

Port ID

The Port ID identifies the port from which LLDP packets were sent.

There are several ways in which a port may be identified, as shown in the following table. A port ID subtype, included in the TLV, indicates how the port is being referenced in the Port ID field.

TABLE 27 Port ID subtypes

ID subtype	Description
0	Reserved
1	Interface alias
2	Port component
3	MAC address
4	Network address
5	Interface name
6	Agent circuit ID
7	Locally assigned
8 - 255	Reserved


Brocade devices use port ID subtype 3, the permanent MAC address associated with the port. Other third party devices may use a port ID subtype other than 3. The port ID appears similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
Port ID (MAC address): 0000.0033.e2d3
```

The LLDPDU format is shown in [LLDP packets](#) on page 176.

The Port ID TLV format is shown below.

FIGURE 8 Port ID TLV packet format



TTL value

The Time to Live (TTL) Value is the length of time the receiving device should maintain the information acquired by LLDP in its MIB.

The TTL value is automatically computed based on the LLDP configuration settings. The TTL value will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (show lldp local-info).

```
Time to live: 40 seconds
```


If the TTL field has a value other than zero, the receiving LLDP agent is notified to completely replace all information associated with the LLDP agent/port with the information in the received LLDPDU.

If the TTL field value is zero, the receiving LLDP agent is notified that all system information associated with the LLDP agent/port is to be deleted. This TLV may be used, for example, to signal that the sending port has initiated a port shutdown procedure.

The LLDPDU format is shown in [LLDP packets](#) on page 176.

The TTL TLV format is shown below.

FIGURE 9 TTL TLV packet format



MIB support

Brocade devices support the following standard management information base (MIB) modules:

- LLDP-MIB
- LLDP-EXT-DOT1-MIB
- LLDP-EXT-DOT3-MIB
- LLDP-EXT-MED-MIB

Syslog messages

Syslog messages for LLDP provide management applications with information related to MIB data consistency and general status. These Syslog messages correspond to the `lldpRemTablesChange` SNMP notifications. Refer to [Enabling LLDP SNMP notifications and Syslog messages](#) on page 184.

Syslog messages for LLDP-MED provide management applications with information related to topology changes. These Syslog messages correspond to the `lldpXMedTopologyChangeDetected` SNMP notifications. Refer to [Enabling SNMP notifications and Syslog messages for LLDP-MED topology changes](#) on page 193.

LLDP configuration

This section describes how to enable and configure LLDP.

The following table lists the LLDP global-level tasks and the default behavior/value for each task.

TABLE 28 LLDP global configuration tasks and default behavior /value

Global task	Default behavior / value when LLDP is enabled
Enabling LLDP on a global basis	Disabled
Specifying the maximum number of LLDP neighbors per device	Automatically set to 392 neighbors per device
Specifying the maximum number of LLDP neighbors per port	Automatically set to 4 neighbors per port
Enabling SNMP notifications and Syslog messages	Disabled
Changing the minimum time between SNMP traps and Syslog messages	Automatically set to 2 seconds when SNMP notifications and Syslog messages for LLDP are enabled
Enabling and disabling TLV advertisements	When LLDP transmit is enabled, by default, the Brocade device will automatically advertise LLDP capabilities, except for the system description, VLAN name, and power-via-MDI information, which may be configured by the system administrator. Also, if desired, you can disable the advertisement of individual TLVs.
Changing the minimum time between LLDP transmissions	Automatically set to 2 seconds
Changing the interval between regular LLDP transmissions	Automatically set to 30 seconds
Changing the holdtime multiplier for transmit TTL	Automatically set to 4
Changing the minimum time between port reinitializations	Automatically set to 2 seconds

LLDP configuration notes and considerations

- LLDP is supported on Ethernet interfaces only.
- By default, if a port is 802.1X-enabled, the transmission and reception of LLDP packets will only take place while the port is authorized. The `lldp-pass-through` command overrides this behavior.
- Cisco Discovery Protocol (CDP) and Brocade Discovery Protocol (FDP) run independently of LLDP. Therefore, these discovery protocols can run simultaneously on the same device.
- By default, the Brocade device limits the number of neighbors per port to four, and staggers the transmission of LLDP packets on different ports, in order to minimize any high-usage spikes to the CPU.
- By default, the Brocade device forwards LLDP packets even though LLDP is not configured on the device. This ensures consistency with other protocols and allows transparent forwarding, though it amounts to noncompliance with IEEE Standards.
- Ports that are in blocking mode (spanning tree) can still receive LLDP packets from a forwarding port.

- Auto-negotiation status indicates what is being advertised by the port for 802.3 auto-negotiation.

Enabling and disabling LLDP

LLDP is enabled by default on individual ports. However, to run LLDP, you must first enable it on a global basis (on the entire device).

To enable LLDP globally, enter the following command at the global CONFIG level of the CLI.

```
device(config)#lldp run
```

Syntax: `[no] lldp run`

Enabling support for tagged LLDP packets

By default, Brocade devices do not accept tagged LLDP packets from other vendors' devices. To enable support, apply the command **lldp tagged-packets process** at the Global CONFIG level of the CLI. When enabled, the device will accept incoming LLDP tagged packets if the VLAN tag matches any of the following:

- a configured VLAN on the port
- the default VLAN for a tagged port
- the configured untagged VLAN for a dual-mode port

To enable support for tagged LLDP packets, enter the following command.

```
device(config)#lldp tagged-packets process
```

Syntax: `[no] lldptagged-packets process`

Changing a port LLDP operating mode

When LLDP is enabled on a global basis, by default, each port on the Brocade device will be capable of transmitting and receiving LLDP packets. You can disable a port's ability to transmit and receive LLDP packets, or change the operating mode to one of the following:

- Transmit LLDP information only
- Receive LLDP information only

You can configure a different operating mode for each port on the Brocade device. For example, you could disable the receipt and transmission of LLDP packets on port e 1/2/1, configure port e 1/2/3 to only receive LLDP packets, and configure port e 1/2/5 to only transmit LLDP packets.

The following sections show how to change the operating mode.

Enabling and disabling receive and transmit mode

To disable the receipt and transmission of LLDP packets on individual ports, enter a command such as the following at the Global CONFIG level of the CLI.

```
device(config)#no lldp enable ports e 1/2/4 e 1/2/5
```

The above command disables LLDP on ports 1/2/4 and 1/2/5. These ports will not transmit nor receive LLDP packets.

To enable LLDP on a port after it has been disabled, enter the following command.

```
device(config)#lldp enable ports e 1/2/4
```

Syntax: `[no] lldp enable ports ethernet port-list | all`

Use the [no] form of the command to disable the receipt and transmission of LLDP packets on a port.

NOTE

When a port is configured to both receive and transmit LLDP packets and the MED capabilities TLV is enabled, LLDP-MED is enabled as well. LLDP-MED is not enabled if the operating mode is set to receive only or transmit only.

Enabling and disabling receive only mode

When LLDP is enabled on a global basis, by default, each port on the Brocade device will be capable of transmitting and receiving LLDP packets. To change the LLDP operating mode from receive and transmit mode to receive only mode, simply disable the transmit mode. Enter a command such as the following at the Global CONFIG level of the CLI.

```
device(config)#no lldp enable transmit ports e 1/2/4 e 1/2/5 e 1/2/6
```

The above command changes the LLDP operating mode on ports 1/2/4, 1/2/5, and 1/2/6 from transmit and receive mode to receive only mode.

To change a port LLDP operating mode from transmit only to receive only, first disable the transmit only mode, then enable the receive only mode. Enter commands such as the following.

```
device(config)#no lldp enable transmit ports e 1/2/7 e 1/2/8 e 1/2/9
device(config)#lldp enable receive ports e 1/2/7 e 1/2/8 e 1/2/9
```

The above commands change the LLDP operating mode on ports 1/2/7, 1/2/8, and 1/2/9, from transmit only to receive only. Note that if you do not disable the transmit only mode, you will configure the port to both transmit and receive LLDP packets.

NOTE

LLDP-MED is not enabled when you enable the receive only operating mode. To enable LLDP-MED, you must configure the port to both receive and transmit LLDP packets. Refer to [Changing a port LLDP operating mode](#).

Syntax:[no] lldp enable receive ports ethernet *port-list* | all

Use the [no] form of the command to disable the receive only mode.

Enabling and disabling transmit only mode

When LLDP is enabled on a global basis, by default, each port on the Brocade device will be capable of transmitting and receiving LLDP packets. To change the LLDP operating mode to transmit only mode, simply disable the receive mode. Enter a command such as the following at the Global CONFIG level of the CLI.

```
device(config)#no lldp enable receive ports e 1/2/4 e 1/2/5 e 1/2/6
```

The above command changes the LLDP operating mode on ports 1/2/4, 1/2/5, and 1/2/6 from transmit and receive mode to transmit only mode. Any incoming LLDP packets will be dropped in software.

To change a port LLDP operating mode from receive only to transmit only, first disable the receive only mode, then enable the transmit only mode. For example, enter commands such as the following at the Global CONFIG level of the CLI.

```
device(config)#no lldp enable receive ports e 1/2/7 e 1/2/8
device(config)#lldp enable transmit ports e 1/2/7 e 1/2/8
```

The above commands change the LLDP operating mode on ports 1/2/7 and 1/2/8 from receive only mode to transmit only mode. Any incoming LLDP packets will be dropped in software. Note that if you do not disable receive only mode, you will configure the port to both receive and transmit LLDP packets.

NOTE

LLDP-MED is not enabled when you enable the transmit only operating mode. To enable LLDP-MED, you must configure the port to both receive and transmit LLDP packets. Refer to [Changing a port LLDP operating mode](#).

Syntax: `[no] lldp enabletransmit ports ethernet port-list | all`

Use the `[no]` form of the command to disable the *transmit only* mode.

Configuring LLDP processing on 802.1x blocked port

This feature adds support for reception and transmission of Link Layer Discovery Protocol (LLDP) packets over an 802.1x blocked port. The default behavior is to drop received LLDP packets and not to transmit LLDP packets over an 802.1x disabled port. To receive or transmit LLDP packets over 802.1x blocked port or in other words to enable the LLDP processing on 802.1x blocked ports, use the **lldp-pass-through** configuration command.

To enable the LLDP processing on all 802.1x blocked ports, enter the following command at the 802.1X configuration mode:

```
Brocade(config-dot1x)# lldp-pass-through all
```

Syntax: `[no] lldp-pass-through all`

To enable LLDP processing on a specific 802.1x blocked port, enter the following command at the 802.1X configuration mode:

```
Brocade(config-dot1x)# lldp-pass-through ethernet 1/1/1
```

Syntax: `[no] lldp-pass-through ethernet port`

Specify the *port* variable in the format *unit/slot/port*.

The **no** form of these commands disables LLDP processing on 802.1x blocked ports.

For more information on LLDP and 801.1x, refer IEEE 802.1AB and IEEE 802.1x.

NOTE

If **lldp-pass-through** is disabled, the neighboring information is lost only after LLDP timeout period (default is 120).

Maximum number of LLDP neighbors

You can change the limit of the number of LLDP neighbors for which LLDP data will be retained, per device as well as per port.

Specifying the maximum number of LLDP neighbors per device

You can change the maximum number of neighbors for which LLDP data will be retained for the entire system.

For example, to change the maximum number of LLDP neighbors for the entire device to 26, enter the following command.

```
device(config)#lldp max-total-neighbors 26
```

Syntax: `[no] lldp max-total-neighbors value`

Use the `[no]` form of the command to remove the static configuration and revert to the default value of 392.

where *value* is a number between 16 and 8192. The default number of LLDP neighbors per device is 392.

Use the **show lldp** command to view the configuration.

Specifying the maximum number of LLDP neighbors per port

You can change the maximum number of LLDP neighbors for which LLDP data will be retained for each port. By default, the maximum number is four and you can change this to a value between one and 64.

For example, to change the maximum number of LLDP neighbors to six, enter the following command.

```
device(config)#lldp max-neighbors-per-port 6
```

Syntax: `[no] lldp max-neighbors-per-port value`

Use the `[no]` form of the command to remove the static configuration and revert to the default value of four.

where *value* is a number from 1 to 64. The default is number of LLDP neighbors per port is four.

Use the **show lldp** command to view the configuration.

Enabling LLDP SNMP notifications and Syslog messages

SNMP notifications and Syslog messages for LLDP provide management applications with information related to MIB data updates and general status.

When you enable LLDP SNMP notifications, corresponding Syslog messages are enabled as well. When you enable LLDP SNMP notifications, the device will send traps and corresponding Syslog messages whenever there are changes to the LLDP data received from neighboring devices.

LLDP SNMP notifications and corresponding Syslog messages are disabled by default. To enable them, enter a command such as the following at the Global CONFIG level of the CLI.

```
device(config)#lldp enable snmp notifications ports e 1/4/2 to 1/4/6
```

The above command enables SNMP notifications and corresponding Syslog messages on ports 1/4/2 through 1/4/6. By default, the device will send no more than one SNMP notification and Syslog message within a five second period. If desired, you can change this interval. Refer to [Specifying the minimum time between SNMP traps and Syslog messages](#) on page 184.

Syntax: `[no] lldp enablesnmp notifications ports ethernet port-list | all`

Specifying the minimum time between SNMP traps and Syslog messages

When SNMP notifications and Syslog messages for LLDP are enabled, the device will send no more than one SNMP notification and corresponding Syslog message within a five second period. If desired, you can throttle the amount of time between transmission of SNMP traps (`lldpRemTablesChange`) and Syslog messages from five seconds up to a value equal to one hour (3600 seconds).

NOTE

Because LLDP Syslog messages are rate limited, some LLDP information given by the system will not match the current LLDP statistics (as shown in the **show lldp statistics** command output).

To change the minimum time interval between traps and Syslog messages, enter a command such as the following.

```
device(config)#lldp snmp-notification-interval 60
```

When the above command is applied, the LLDP agent will send no more than one SNMP notification and Syslog message every 60 seconds.

Syntax: `[no] lldp snmp-notification-interval seconds`

where *seconds* is a value between 5 and 3600. The default is 5 seconds.

Changing the minimum time between LLDP transmissions

The LLDP transmit delay timer limits the number of LLDP frames an LLDP agent can send within a specified time frame. When you enable LLDP, the system automatically sets the LLDP transmit delay timer to two seconds. If desired, you can change the default behavior from two seconds to a value between 1 and 8192 seconds.

NOTE

The LLDP transmit delay timer must not be greater than one quarter of the LLDP transmission interval (CLI command **lldp transmit-interval**).

The LLDP transmit delay timer prevents an LLDP agent from transmitting a series of successive LLDP frames during a short time period, when rapid changes occur in LLDP. It also increases the probability that multiple changes, rather than single changes, will be reported in each LLDP frame.

To change the LLDP transmit delay timer, enter a command such as the following at the Global CONFIG level of the CLI.

```
device(config)#lldp transmit-delay 7
```

The above command causes the LLDP agent to wait a minimum of seven seconds after transmitting an LLDP frame and before sending another LLDP frame.

Syntax: **[no] lldp transmit-delay** *seconds*

where *seconds* is a value between 1 and 8192. The default is two seconds. Note that this value must not be greater than one quarter of the LLDP transmission interval (CLI command **lldp transmit-interval**).

Changing the interval between regular LLDP transmissions

The LLDP transmit interval specifies the number of seconds between regular LLDP packet transmissions. When you enable LLDP, by default, the device will wait 30 seconds between regular LLDP packet transmissions. If desired, you can change the default behavior from 30 seconds to a value between 5 and 32768 seconds.

To change the LLDP transmission interval, enter a command such as the following at the Global CONFIG level of the CLI.

```
device(config)#lldp transmit-interval 40
```

The above command causes the LLDP agent to transmit LLDP frames every 40 seconds.

Syntax: **[no] lldp transmit-interval** *seconds*

where *seconds* is a value from 5 to 32768. The default is 30 seconds.

NOTE

Setting the transmit interval or transmit holdtime multiplier, or both, to inappropriate values can cause the LLDP agent to transmit LLDPDUs with TTL values that are excessively high. This in turn can affect how long a receiving device will retain the information if it is not refreshed.

Changing the holdtime multiplier for transmit TTL

The holdtime multiplier for transmit TTL is used to compute the actual time-to-live (TTL) value used in an LLDP frame. The TTL value is the length of time the receiving device should maintain the information in its MIB. When you enable LLDP, the device automatically sets the holdtime multiplier for TTL to four. If desired, you can change the default behavior from four to a value between two and ten.

To compute the TTL value, the system multiplies the LLDP transmit interval by the holdtime multiplier. For example, if the LLDP transmit interval is 30 and the holdtime multiplier for TTL is 4, then the value 120 is encoded in the TTL field in the LLDP header.

To change the holdtime multiplier, enter a command such as the following at the Global CONFIG level of the CLI.

```
device(config)#lldp transmit-hold 6
```

Syntax: `[no] lldp transmit-hold value`

where *value* is a number from 2 to 10. The default value is 4.

NOTE

Setting the transmit interval or transmit holdtime multiplier, or both, to inappropriate values can cause the LLDP agent to transmit LLDPDUs with TTL values that are excessively high. This in turn can affect how long a receiving device will retain the information if it is not refreshed.

Changing the minimum time between port reinitializations

The LLDP re-initialization delay timer specifies the minimum number of seconds the device will wait from when LLDP is disabled on a port, until it will honor a request to re-enable LLDP on that port. When you enable LLDP, the system sets the re-initialization delay timer to two seconds. If desired, you can change the default behavior from two seconds to a value between one and ten seconds.

To set the re-initialization delay timer, enter a command such as the following at the Global CONFIG level of the CLI.

```
device(config)#lldp reinit-delay 5
```

The above command causes the device to wait five seconds after LLDP is disabled, before attempting to honor a request to re-enable it.

Syntax: `[no] lldp reinit-delay seconds`

where *seconds* is a value from 1 - 10. The default is two seconds.

LLDP TLVs advertised by the Brocade device

When LLDP is enabled on a global basis, the Brocade device will automatically advertise the following information, except for the features noted:

General system information:

- Management address
- Port description
- System capabilities
- System description (not automatically advertised)
- System name

802.1 capabilities:

- VLAN name (not automatically advertised)
- Untagged VLAN ID

802.3 capabilities:

- Link aggregation information
- MAC/PHY configuration and status
- Maximum frame size
- Power-via-MDI information (not automatically advertised)

The above TLVs are described in detail in the following sections.

NOTE

The system description, VLAN name, and power-via-MDI information TLVs are not automatically enabled. The following sections show how to enable these advertisements.

General system information for LLDP

Except for the system description, the Brocade device will advertise the following system information when LLDP is enabled on a global basis:

- Management address
- Port description
- System capabilities
- System description (not automatically advertised)
- System name

Management address

A management address is normally an IPv4 or IPv6 address that can be used to manage the device. Management address advertising has two modes: default, or explicitly configured. The default mode is used when no addresses are configured to be advertised for a given port. If any addresses are configured to be advertised for a given port, then only those addresses are advertised. This applies across address types, so for example, if just one IPv4 address is explicitly configured to be advertised for a port, then no IPv6 addresses will be advertised for that port (since none were configured to be advertised), even if IPv6 addresses are configured within the system.

If no management address is explicitly configured to be advertised, the Brocade device will use the first available IPv4 address and the first available IPv6 address (so it may advertise IPv4, IPv6 or both). A Layer 3 switch will select the first available address of each type from those configured on the following types of interfaces, in the following order of preference:

- Physical port on which LLDP will be transmitting the packet
- Virtual router interface (VE) on a VLAN that the port is a member of
- Dedicated management port
- Loop back interface
- Virtual router interface (VE) on any other VLAN
- Other physical port
- Other interface

For IPv6 addresses, link-local and anycast addresses will be excluded from these searches.

If no IP address is configured on any of the above, the port's current MAC address will be advertised.

To advertise a IPv4 management address, enter a command such as the following:

```
device(config)#lldp advertise management-address ipv4 10.157.2.1 ports e 1/1/4
```

The management address will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**):

```
Management address (IPv4): 10.157.2.1
```

Syntax:[no] lldp advertise management-address ipv4 *ipv4 address* ports ethernet *port list* | all

To support an IPv6 management address, there is a similar command that has equivalent behavior as the IPv4 command.

To advertise an IPv6 management address, enter a command such as the following:

```
device(config)#lldp advertise management-address ipv6 2001:DB8::90 ports e 1/2/7
```

Syntax:[no] **lldp advertise management-address ipv6** *ipv6 address* **ports ethernet** *port list* | **all**

ipv4 address or *ipv6 address* or both are the addresses that may be used to reach higher layer entities to assist discovery by network management. In addition to management addresses, the advertisement will include the system interface number associated with the management address.

For *port list*, specify the ports in the format *unit/slot/port*. You can list all of the ports individually; use the keyword **to** specify a range of ports, or a combination of both. To apply the configuration to all ports on the device, use the keyword **all** instead of listing the ports individually.

Port description

The port description TLV identifies the port from which the LLDP agent transmitted the advertisement. The port description is taken from the ifDescr MIB object from MIB-II.

By default, the port description is automatically advertised when LLDP is enabled on a global basis. To disable advertisement of the port description, enter a command such as the following.

```
device(config)#no lldp advertise port-description ports e 1/2/4 to 1/2/12
```

The port description will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
Port description: "GigabitEthernet20"
```

Syntax:[no] **lldp advertise port-description ports ethernet** *port-list* | **all**

System capabilities

The system capabilities TLV identifies the primary functions of the device and indicates whether these primary functions are enabled. The primary functions can be one or more of the following (more than one for example, if the device is both a bridge and a router):

- Repeater
- Bridge
- WLAN access point
- Router
- Telephone
- DOCSIS cable device
- Station only (devices that implement end station capability)
- Other

System capabilities for Brocade devices are based on the type of software image in use (for example, Layer 2 switch or Layer 3 router). The enabled capabilities will be the same as the available capabilities, except that when using a router image (base or full Layer 3), if the global route-only feature is turned on, the bridge capability will not be included, since no bridging takes place.

By default, the system capabilities are automatically advertised when LLDP is enabled on a global basis. To disable this advertisement, enter a command such as the following.

```
device(config)#no lldp advertise system-capabilities ports e 1/2/4 to 1/2/12
```

The system capabilities will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
System capabilities :   bridge
Enabled capabilities:   bridge
```

Syntax: `[no] lldp advertisesystem-capabilities ports ethernet port-list | all`

System description

The system description is the network entity, which can include information such as the product name or model number, the version of the system hardware type, the software operating system level, and the networking software version. The information corresponds to the sysDescr MIB object in MIB-II.

To advertise the system description, enter a command such as the following.

```
device(config)#lldp advertise system-description ports e 1/2/4 to 1/2/12
```

The system description will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
+ System description : "Brocade Communications, Inc.,FCX_ADV_ROUTER_SOFT_PACKAGE,
IronWare Version 07.3.00T7f3 compiled on Sep 26 2011 at
21:15:14 labeled as FCXR07300
```

NOTE

The contents of the show command output will vary depending on which TLVs are configured to be advertised.

Syntax:`[no] lldp advertise system-description ports ethernet port-list | all`

System name

The system name is the system administratively assigned name, taken from the sysName MIB object in MIB-II. The sysName MIB object corresponds to the name defined with the CLI command **hostname**.

By default, the system name is automatically advertised when LLDP is enabled on a global basis. To disable this advertisement, enter a command such as the following.

```
device(config)#no lldp advertise system-name ports e 1/2/4 to 1/2/12
```

The system name will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
System name: "FCX624SHPOE-ADV Router"
```

Syntax:`[no] lldp advertise system-name ports ethernet port-list | all`

802.1 capabilities

Except for the VLAN name, the Brocade device will advertise the following 802.1 attributes when LLDP is enabled on a global basis:

- VLAN name (not automatically advertised)
- Untagged VLAN ID

VLAN name

The VLAN name TLV contains the name and VLAN ID of a VLAN configured on a port. An LLDPDU may include multiple instances of this TLV, each for a different VLAN.

To advertise the VLAN name, enter a command such as the following.

```
device(config)#lldp advertise vlan-name vlan 99 ports e 1/2/4 to 1/2/12
```

The VLAN name will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
VLAN name (VLAN 99): "Voice-VLAN-99"
```

Syntax: `[no] lldp advertise vlan-name vlan vlan ID ports ethernet port-list | all`

For *vlan ID*, enter the VLAN ID to advertise.

Untagged VLAN ID

The port VLAN ID TLV advertises the Port VLAN Identifier (PVID) that will be associated with untagged or priority-tagged frames. If the port is not an untagged member of any VLAN (i.e., the port is strictly a tagged port), the value zero will indicate that.

By default, the port VLAN ID is automatically advertised when LLDP is enabled on a global basis. To disable this advertisement, enter a command such as the following.

```
device(config)#no lldp advertise port-vlan-id ports e 1/2/4 to 1/2/12
```

The untagged VLAN ID will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
Port VLAN ID: 99
```

Syntax: `[no] lldp advertise port-vlan-id ports ethernet port-list | all`

802.3 capabilities

Except for Power-via-MDI information, the Brocade device will advertise the following 802.3 attributes when LLDP is enabled on a global basis:

- Link aggregation information
- MAC/PHY configuration and status
- Maximum frame size
- Power-via-MDI information (not automatically advertised)

Link aggregation TLV

The **link-aggregation** time, length, value (TLV) indicates the following:

- Whether the link is capable of being aggregated
- Whether the link is currently aggregated
- The primary trunk port

Brocade devices advertise link aggregation information about standard link aggregation (LACP) as well as static trunk configuration.

By default, link-aggregation information is automatically advertised when LLDP is enabled on a global basis. To disable this advertisement, enter a command such as the following.

```
device(config)#no lldp advertise link-aggregation ports e 1/2/12
```

Syntax: `[no] lldp advertise link-aggregation ports ethernet port-list | all`

The link aggregation advertisement will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
Link aggregation: not capable
```

MAC and PHY configuration status

The MAC and PHY configuration and status TLV includes the following information:

- Auto-negotiation capability and status
- Speed and duplex mode
- Flow control capabilities for auto-negotiation
- maximum port speed advertisement
- If applicable, indicates if the above settings are the result of auto-negotiation during link initiation or of a manual set override action

The advertisement reflects the effects of the following CLI commands:

- speed-duplex
- flow-control
- gig-default
- link-config

By default, the MAC/PHY configuration and status information are automatically advertised when LLDP is enabled on a global basis. To disable this advertisement, enter a command such as the following.

```
device(config)#no lldp advertise mac-phy-config-status ports e 1/2/4 to 1/2/12
```

The MAC/PHY configuration advertisement will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
+ 802.3 MAC/PHY      : auto-negotiation enabled
  Advertised capabilities: 10baseT-HD, 10baseT-FD, 100baseTX-HD, 100baseTX-FD,
  fdxSPause, fdxBPause, 1000baseT-HD, 1000baseT-FD
  Operational MAU type: 100BaseTX-FD
```

Syntax: `[no] lldp advertise mac-phy-config-status ports ethernet port-list | all`

Maximum frame size

The maximum frame size TLV provides the maximum 802.3 frame size capability of the port. This value is expressed in octets and includes the four-octet Frame Check Sequence (FCS). The default maximum frame size is 1522. The advertised value may change depending on whether the **aggregated-vlan** or **jumbo** CLI commands are in effect.

NOTE

On 48GC modules in non-jumbo mode, the maximum size of ping packets is 1486 bytes and the maximum frame size of tagged traffic is no larger than 1581 bytes.

By default, the maximum frame size is automatically advertised when LLDP is enabled on a global basis. To disable this advertisement, enter a command such as the following.

```
device(config)#no lldp advertise max-frame-size ports e 1/2/4 to 1/2/12
```

The maximum frame size advertisement will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
Maximum frame size: 1522 octets
```

Syntax:[no] lldp advertise max-frame-size ports ethernet *port-list* | all

Power-via-MDI

The power-via-MDI TLV provides general information about Power over Ethernet (POE) capabilities and status of the port. It indicates the following:

- POE capability (supported or not supported)
- POE status (enabled or disabled)
- Power Sourcing Equipment (PSE) power pair - indicates which pair of wires is in use and whether the pair selection can be controlled. The Brocade implementation always uses pair A, and cannot be controlled.
- Power class - Indicates the range of power that the connected powered device has negotiated or requested.

NOTE

The power-via-MDI TLV described in this section applies to LLDP. There is also a power-via-MDI TLV for LLDP-MED devices, which provides extensive POE information. Refer to [Extended power-via-MDI information](#) on page 202.

To advertise the power-via-MDI information, enter a command such as the following.

```
device(config)#lldp advertise power-via-mdi ports e 1/2/4 to 1/2/12
```

The power-via-MDI advertisement will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
+ 802.3 Power via MDI: PSE port, power enabled, class 0
  Power Pair          : A (not controllable)
```

Syntax:[no] lldp advertise power-via-mdi ports ethernet *port-list* | all

LLDP-MED configuration

This section provides the details for configuring LLDP-MED.

The following table lists the global and interface-level tasks and the default behavior/value for each task.

TABLE 29 LLDP-MED configuration tasks and default behavior / value

Task	Default behavior / value
Global CONFIG-level tasks	
Enabling LLDP-MED on a global basis	Disabled
Enabling SNMP notifications and Syslog messages for LLDP-MED topology change	Disabled
Changing the Fast Start Repeat Count	The system automatically sets the fast start repeat count to 3 when a Network Connectivity Device receives an LLDP packet from an Endpoint that is newly connected to the network.

TABLE 29 LLDP-MED configuration tasks and default behavior / value (continued)

Task	Default behavior / value
	<p>NOTE</p> <p>The LLDP-MED fast start mechanism is only intended to run on links between Network Connectivity devices and Endpoint devices. It does not apply to links between LAN infrastructure elements, including between Network Connectivity devices, or to other types of links.</p>
Interface-level tasks	
Defining a location ID	Not configured
Defining a network policy	Not configured

Enabling LLDP-MED

When LLDP is enabled globally, LLDP-MED is enabled if the LLDP-MED capabilities TLV is also enabled. By default, the LLDP-MED capabilities TLV is automatically enabled. To enable LLDP, refer to [Enabling and disabling LLDP](#) on page 181.

NOTE

LLDP-MED is not enabled on ports where the LLDP operating mode is receive only or transmit only. LLDP-MED is enabled on ports that are configured to both receive and transmit LLDP packets and have the LLDP-MED capabilities TLV enabled.

Enabling SNMP notifications and Syslog messages for LLDP-MED topology changes

SNMP notifications and Syslog messages for LLDP-MED provide management applications with information related to topology changes. For example, SNMP notifications can alert the system whenever a remote Endpoint device is connected to or removed from a local port.

SNMP notifications identify the local port where the topology change occurred, as well as the device capability of the remote Endpoint device that was connected to or removed from the port.

When you enable LLDP-MED SNMP notifications, corresponding Syslog messages are enabled as well. When you enable LLDP-MED SNMP notifications, the device will send traps and Syslog messages when an LLDP-MED Endpoint neighbor entry is added or removed.

SNMP notifications and corresponding Syslog messages are disabled by default. To enable them, enter a command such as the following at the Global CONFIG level of the CLI.

```
device(config)#lldp enable snmp med-topo-change-notifications ports e 1/4/4 to 1/4/6
```

Syntax:[no] lldp enable snmp med-topo-change-notifications ports ethernet *port-list* | all

Changing the fast start repeat count

The fast start feature enables a Network Connectivity Device to initially advertise itself at a faster rate for a limited time when an LLDP-MED Endpoint has been newly detected or connected to the network. This feature is important within a VoIP network, for example, where rapid availability is crucial for applications such as emergency call service location (E911).

The fast start timer starts when a Network Connectivity Device receives the first LLDP frame from a newly detected Endpoint.

The LLDP-MED fast start repeat count specifies the number of LLDP packets that will be sent during the LLDP-MED fast start period. By default, the device will send three packets at one-second intervals. If desired, you can change the number of packets the device will send per second, up to a maximum of 10.

NOTE

The LLDP-MED fast start mechanism is only intended to run on links between Network Connectivity devices and Endpoint devices. It does not apply to links between LAN infrastructure elements, including between Network Connectivity devices, or to other types of links.

To change the LLDP-MED fast start repeat count, enter commands such as the following.

```
device(config)#lldp med fast-start-repeat-count 5
```

The above command causes the device to send five LLDP packets during the LLDP-MED fast start period.

Syntax: `[no] lldp medfast-start-repeat-count value`

where value is a number from 1 to 10, which specifies the number of packets that will be sent during the LLDP-MED fast start period. The default is 3.

Defining a location id

The LLDP-MED Location Identification extension enables the Brocade device to set the physical location that an attached Class III Endpoint will use for location-based applications. This feature is important for applications such as IP telephony, for example, where emergency responders need to quickly determine the physical location of a user in North America that has just dialed 911.

For each port, you can define one or more of the following location ID formats:

- Geographic location (coordinate-based)
- Civic address
- Emergency Call Services (ECS) Emergency Location Identification Number (ELIN)

The above location ID formats are defined in the following sections.

Coordinate-based location

Coordinate-based location is based on the IETF RFC 3825 [6] standard, which specifies a Dynamic Host Configuration Protocol (DHCP) option for the coordinate-based geographic location of a client.

When you configure an Endpoint location information using the coordinate-based location, you specify the latitude, longitude, and altitude, along with resolution indicators (a measure of the accuracy of the coordinates), and the reference datum (the map used for the given coordinates).

To configure a coordinate-based location for an Endpoint device, enter a command such as the following at the Global CONFIG level of the CLI.

```
device(config)#lldp med location-id coordinate-based latitude
-78.303 resolution 20 longitude 34.27 resolution 18 altitude meters 50 resolution 16 wgs84
```

Syntax: `[no] lldp med location-id coordinate-based latitude degrees resolution bits longitude degrees resolution bits altitude floors number resolution bits | meters number resolution bits datum`

latitude degrees is the angular distance north or south from the earth equator measured through 90 degrees. Positive numbers indicate a location north of the equator and negative numbers indicate a location south of the equator.

resolution bits specifies the precision of the value given for latitude. A smaller value increases the area within which the device is located. For latitude, enter a number between 1 and 34.

longitude degrees is the angular distance from the intersection of the zero meridian. Positive values indicate a location east of the prime meridian and negative numbers indicate a location west of the prime meridian.

resolution bits specifies the precision of the value given for longitude. A smaller value increases the area within which the device is located. For longitude resolution, enter a number between 1 and 34.

altitude floors number is the vertical elevation of a building above the ground, where 0 represents the floor level associated with the ground level at the main entrance and larger values represent floors that are above (higher in altitude) floors with lower values. For example, 2 for the 2nd floor. Sub-floors can be represented by non-integer values. For example, a mezzanine between floor 1 and floor 2 could be represented as 1.1. Similarly, the mezzanines between floor 4 and floor 5 could be represented as 4.1 and 4.2 respectively. Floors located below ground level could be represented by negative values.

resolution bits specifies the precision of the value given for altitude. A smaller value increases the area within which the device is located. For floors resolution, enter the value 0 if the floor is unknown, or 30 if a valid floor is being specified.

altitude meters number is the vertical elevation in number of meters, as opposed to floors.

resolution bits specifies the precision of the value given for altitude. A smaller value increases the area within which the device is located. For meters resolution, enter a value from 0 to 30.

Datum is the map used as the basis for calculating the location. Specify one of the following:

- **wgs84** - (geographical 3D) - World Geodesic System 1984, CRS Code 4327, Prime Meridian Name: Greenwich
- **nad83-navd88** - North American Datum 1983, CRS Code 4269, Prime Meridian Name: Greenwich; The associated vertical datum is the North American Vertical Datum of 1988 (NAVD88). Use this datum when referencing locations on land. If land is near tidal water, use nad83-mllw (below).
- **nad83-mllw** - North American Datum 1983, CRS Code 4269, Prime Meridian Name: Greenwich; The associated vertical datum is mean lower low water (MLLW). Use this datum when referencing locations on water, sea, or ocean.

Example coordinate-based location configuration

The following shows an example coordinate-based location configuration for the Sears Tower, at the following location.

103rd Floor 233 South Wacker Drive Chicago, IL 60606

```
device(config)#lldp med location-id coordinate-based latitude 41.87884 resolution 18 longitude 87.63602
resolution 18 altitude floors 103 resolution 30 wgs84
```

The above configuration shows the following:

- Latitude is 41.87884 degrees north (or 41.87884 degrees).
- Longitude is 87.63602 degrees west (or 87.63602 degrees).
- The latitude and longitude resolution of 18 describes a geo-location area that is latitude 41.8769531 to latitude 41.8789062 and extends from -87.6367188 to -87.6347657 degrees longitude. This is an area of approximately 373412 square feet (713.3 ft. x 523.5 ft.).
- The location is inside a structure, on the 103rd floor.
- The WGS 84 map was used as the basis for calculating the location.

Example coordinate-based location advertisement

The coordinate-based location advertisement will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
+ MED Location ID
  Data Format: Coordinate-based
  Latitude Resolution : 20 bits
```

```

Latitude Value      : -78.303 degrees
Longitude Resolution : 18 bits
Longitude Value     : 34.27 degrees
Altitude Resolution : 16 bits
Altitude Value      : 50. meters
Datum               : WGS 84

```

Configuring civic address location

When you configure a media Endpoint location using the address-based location, you specify the location the entry refers to, the country code, and the elements that describe the civic or postal address.

To configure a civic address-based location for LLDP-MED, enter commands such as the following at the Global CONFIG level of the CLI.

```

device(config)#lldp med location-id civic-address refers-to client country US elem 1 CA elem 3 "Santa
Clara" elem 6 "4980 Great America Pkwy" elem 24 95054 elem 27 5 elem 28 551 elem 29 office elem 23 "John
Doe"

```

Syntax: `[no] lldp med location-id civic-address refers-to elem country country code elem CA type value [elem CA type value] [elem CA type value]`

refers-to elem describes the location that the entry refers to. Specify one of the following:

- client
- dhcp-server
- network-element

where **dhcp-server** or **network-element** should only be used if it is known that the Endpoint is in close physical proximity to the DHCP server or network element.

country code is the two-letter ISO 3166 country code in capital ASCII letters.

- CA - Canada
- DE - Germany
- JP - Japan
- KR - Korea
- US - United States

CA type is a value from 0 - 255, that describes the civic address element. For example, a CA type of 24 specifies a postal or zip code. Valid elements and their types are listed in the following table.

value is the actual value of the elem CA type, above. For example, 95123 for the postal or zip code. Acceptable values are also listed in the following table.

NOTE

If the value of an element contains one or more spaces, use double quotation marks (") at the beginning and end of the string. For example, elem 3 "Santa Clara".

TABLE 30 Elements used with civic address

Civic Address (CA) type	Description	Acceptable values / examples
0	Language	The ISO 639 language code used for presenting the address information.
1	National subdivisions (state, canton, region, province, or prefecture)	Examples: Canada - Province Germany - State

TABLE 30 Elements used with civic address (continued)

Civic Address (CA) type	Description	Acceptable values / examples
		Japan - Metropolis Korea - Province United States - State
2	County, parish, gun (JP), or district (IN)	Examples: Canada - County Germany - County Japan - City or rural area Korea - County United States - County
3	City, township, or shi (JP)	Examples: Canada - City or town Germany - City Japan - Ward or village Korea - City or village United States - City or town
4	City division, borough, city district, ward, or chou (JP)	Examples: Canada - N/A Germany - District Japan - Town Korea - Urban district United States - N/A
5	Neighborhood or block	Examples: Canada - N/A Germany - N/A Japan - City district Korea - Neighborhood United States - N/A
6	Street	Examples: Canada - Street Germany - Street Japan - Block Korea - Street United States - Street
16	Leading street direction	N (north), E (east), S (south), W (west), NE, NW, SE, SW
17	Trailing street suffix	N (north), E (east), S (south), W (west), NE, NW, SE, SW

TABLE 30 Elements used with civic address (continued)

Civic Address (CA) type	Description	Acceptable values / examples
18	Street suffix	Acceptable values for the United States are listed in the United States Postal Service Publication 28 [18], Appendix C. Example: Ave, Place
19	House number	The house number (street address) Example: 1234
20	House number suffix	A modifier to the house number. It does not include parts of the house number. Example: A, 1/2
21	Landmark or vanity address	A string name for a location. It conveys a common local designation of a structure, a group of buildings, or a place that helps to locate the place. Example: UC Berkeley
22	Additional location information	An unstructured string name that conveys additional information about the location. Example: west wing
23	Name (residence and office occupant)	Identifies the person or organization associated with the address. Example: Textures Beauty Salon
24	Postal / zip code	The valid postal / zip code for the address. Example: 95054-1234
25	Building (structure)	The name of a single building if the street address includes more than one building or if the building name is helpful in identifying the location. Example: Law Library
26	Unit (apartment, suite)	The name or number of a part of a structure where there are separate administrative units, owners, or tenants, such as separate companies or families who occupy that structure. Common examples include suite or apartment designations. Example: Apt 27
27	Floor	Example: 4
28	Room number	The smallest identifiable subdivision of a structure. Example: 7A
29	Placetype	The type of place described by the civic coordinates. For example, a home, office, street, or other public space. Example: Office
30	Postal community name	When the postal community name is defined, the civic community name (typically CA type 3) is replaced by this value. Example: Alviso

TABLE 30 Elements used with civic address (continued)

Civic Address (CA) type	Description	Acceptable values / examples
31	Post office box (P.O. box)	When a P.O. box is defined, the street address components (CA types 6, 16, 17, 18, 19, and 20) are replaced with this value. Example: P.O. Box 1234
32	Additional code	An additional country-specific code that identifies the location. For example, for Japan, this is the Japan Industry Standard (JIS) address code. The JIS address code provides a unique address inside of Japan, down to the level of indicating the floor of the building.
128	Script	The script (from ISO 15924 [14]) used to present the address information. Example: Latn NOTE If not manually configured, the system assigns the default value Latn
255	Reserved	

Example civic address location advertisement

The Civic address location advertisement will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
+ MED Location ID
  Data Format: Civic Address
  Location of: Client
  Country    : "US"
  CA Type    : 1
  CA Value   : "CA"
  CA Type    : 3
  CA Value   : "Santa Clara"
  CA Type    : 6
  CA Value   : "4980 Great America Pkwy."
  CA Type    : 24
  CA Value   : "95054"
  CA Type    : 27
  CA Value   : "5"
  CA Type    : 28
  CA Value   : "551"
  CA Type    : 29
  CA Value   : "office"
  CA Type    : 23
  CA Value   : "John Doe"
```

Configuring emergency call service

The Emergency Call Service (ECS) location is used specifically for Emergency Call Services applications.

When you configure a media Endpoint location using the emergency call services location, you specify the Emergency Location Identification Number (ELIN) from the North America Numbering Plan format, supplied to the Public Safety Answering Point (PSAP) for ECS purposes.

To configure an ECS-based location for LLDP-MED, enter a command such as the following at the Global CONFIG level of the CLI.

```
device(config)#lldp med location-id ecs-elin 4082071700
```

Syntax: `[no] lldp med location-id ecs-elin number ports ethernet port-list | all`

number is a number from 10 to 25 digits in length.

Example ECS ELIN location advertisements

The ECS ELIN location advertisement will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
+ MED Location ID
  Data Format: ECS ELIN
  Value      : 4082071700
```

Defining an LLDP-MED network policy

An LLDP-MED network policy defines an Endpoint VLAN configuration (VLAN type and VLAN ID) and associated Layer 2 and Layer 3 priorities that apply to a specific set of applications on a port.

NOTE

This feature applies to applications that have specific real-time network policy requirements, such as interactive voice or video services. It is not intended to run on links other than between Network Connectivity devices and Endpoints, and therefore does not advertise the multitude of network policies that frequently run on an aggregated link.

To define an LLDP-MED network policy for an Endpoint, enter a command such as the following.

```
device(config)#lldp med network-policy application voice tagged vlan 99 priority 3 dscp 22 port e 1/2/6
```

The network policy advertisement will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
+ MED Network Policy
  Application Type : Voice
  Policy Flags    : Known Policy, Tagged
  VLAN ID        : 99
  L2 Priority     : 3
  DSCP Value     : 22
```

NOTE

Endpoints will advertise a policy as "unknown" in the **show lldp neighbor detail** command output, if it is a policy that is required by the Endpoint and the Endpoint has not yet received it.

LLDP-MED network policy configuration syntax

The CLI syntax for defining an LLDP-MED network policy differs for tagged, untagged, and priority tagged traffic. Refer to the appropriate syntax, below.

For tagged traffic

Syntax: `[no] lldp med network-policy application application type taggedvlan vlan ID priority 0-7 dscp 0-63 ports ethernet port-list | all`

For untagged traffic

Syntax:`[no] lldp med network-policy application application type untagged dscp 0-63 ports ethernet port-list | all`

For priority-tagged traffic

Syntax: `[no] lldp med network-policy application application type priority-tagged priority 0-7 dscp 0-63 ports ethernet port-list | all`

application type indicates the primary function of the applications defined by this network policy. Application type can be one of the following:

- **guest-voice** - Limited voice service for guest users and visitors with their own IP telephony handsets or similar devices that support interactive voice services.
- **guest-voice-signaling** - Limited voice service for use in network topologies that require a different policy for guest voice signaling than for guest voice media.
- **softphone-voice** - Softphone voice service for use with multi-media applications that work in association with VoIP technology, enabling phone calls direct from a PC or laptop. Softphones do not usually support multiple VLANs, and are typically configured to use an untagged VLAN or a single tagged data-specific VLAN. Note that when a network policy is defined for use with an untagged VLAN, the Layer 2 priority field is ignored and only the DSCP value is relevant.
- **streaming-video** - Applies to broadcast- or multicast-based video content distribution and similar applications that support streaming video services requiring specific network policy treatment. Video applications that rely on TCP without buffering would not be an intended use of this application type.
- **video-conferencing** - Applies to dedicated video conferencing equipment and similar devices that support real-time interactive video/audio services.
- **video-signaling** - For use in network topologies that require a separate policy for video signaling than for video media. Note that this application type should not be advertised if all the same network policies apply as those advertised in the video conferencing policy TLV.
- **voice** - For use by dedicated IP telephony handsets and similar devices that support interactive voice services.
- **voice-signaling** - For use in network topologies that require a different policy for voice signaling than for voice media. Note that this application type should not be advertised if all the same network policies apply as those advertised in the voice policy TLV.
- *tagged vlan *vlan id** specifies the tagged VLAN that the specified application type will use.
- *untagged* indicates that the device is using an untagged frame format.
- *priority-tagged* indicates that the device uses priority-tagged frames. In this case, the device uses the default VLAN (PVID) of the ingress port.
- *priority 0 -7* indicates the Layer 2 priority value to be used for the specified application type. Enter 0 to use the default priority.
- *dscp 0 - 63* specifies the Layer 3 Differentiated Service codepoint priority value to be used for the specified application type. Enter 0 to use the default priority.

LLDP-MED attributes advertised by the Brocade device

LLDP-MED attributes are only advertised on a port if LLDP-MED is enabled (which is done by enabling the LLDP-MED capabilities TLV), the port operating mode is *receive* and *transmit* (the default), and the port has received an LLDP-MED advertisement from an Endpoint. By default, the Brocade device will automatically advertise the following LLDP-MED attributes when the above criteria are met:

- LLDP-MED capabilities
- Location ID
- Network policy
- Power-via-MDI information

NOTE

Although the Location ID and Network policy attributes are automatically advertised, they will have no effect until they are actually defined.

LLDP-MED capabilities

When enabled, LLDP-MED is enabled, and the LLDP-MED capabilities TLV is sent whenever any other LLDP-MED TLV is sent. When disabled, LLDP-MED is disabled and no LLDP-MED TLVs are sent.

The LLDP-MED capabilities advertisement includes the following information:

- The supported LLDP-MED TLVs
- The device type (Network Connectivity device or Endpoint (Class 1, 2, or 3))

By default, LLDP-MED information is automatically advertised when LLDP-MED is enabled. To disable this advertisement, enter a command such as the following.

```
device(config)#no lldp advertise med-capabilities ports e 1/2/4 to 1/2/12
```

NOTE

Disabling the LLDP-MED capabilities TLV disables LLDP-MED.

To re-enable the LLDP-MED Capabilities TLV (and LLDP-MED) after it has been disabled, enter a command such as the following.

```
device(config)#lldp advertise med-capabilities ports e 1/2/4 to 1/2/12
```

The LLDP-MED capabilities advertisement will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
+ MED capabilities: capabilities, networkPolicy, location, extendedPSE      MED device type : Network
Connectivity
```

Syntax: [no] lldp advertised-med-capabilities ports ethernet *port-list* | all

Extended power-via-MDI information

The extended Power-via-MDI TLV enables advanced power management between LLDP-MED Endpoints and Network Connectivity Devices.

This TLV provides significantly more information than the 802.1AB Power-via-MDI TLV referenced in [802.3 capabilities](#) on page 190. For example, this TLV enables an Endpoint to communicate a more precise required power level, thereby enabling the device to allocate less power to the Endpoint, while making more power available to other ports.

The LLDP-MED Power-via-MDI TLV advertises an Endpoint IEEE 802.3af power-related information, including the following:

- **Power type** - indicates whether the LLDP-MED device transmitting the LLDPDU is a power sourcing device or a powered device:
 - **Power sourcing device/equipment (PSE)** - This is the source of the power, or the device that integrates the power onto the network. Power sourcing devices/equipment have embedded POE technology. In this case, the power sourcing device is the Brocade POE device.
 - **Powered device (PD)** - This is the Ethernet device that requires power and is situated on the other end of the cable opposite the power sourcing device.
- **Power source** - The power source being utilized by a PSE or PD, for example, primary power source, backup power source, or unknown.

For Endpoint devices, the power source information indicates the power capability of the Network Connectivity Device it is attached to. When the Network Connectivity device advertises that it is using its primary power source, the Endpoint should expect to have uninterrupted access to its available power. Likewise, if the Network Connectivity device advertises that it is using backup power, the Endpoint should not expect continuous power. The Endpoint may additionally choose to power down non-essential subsystems or to conserve power as long as the PSE is advertising that it is operating on backup power.

NOTE

Brocade devices always advertise the power source as "unknown".

- **Power priority** - The in-line power priority level for the PSE or PD:
 - 3 - low
 - 2 - high
 - 1 - critical
 - unknown
- **Power level** - The total power, in tenths of watts, required by a PD from a PSE, or the total power a PSE is capable of sourcing over a maximum length cable based on its current configuration.

If the exact power is not known for a PSE or PD, it will advertise the power level associated with its 802.3af power class listed in the following table.

TABLE 31 802.3af power classes

Power class	Minimum power level output at the PSE	Maximum power levels at the PD
0	15.4 watts	0.44 - 12.95 watts
1	4.0 watts	0.44 - 3.84 watts
2	7.0 watts	3.84 - 6.49 watts
3	15.4 watts	6.49 - 12.95 watts

For a PD (Endpoint device), the power level represents the maximum power it can consume during normal operations in its current configuration, even if its actual power draw at that instance is less than the advertised power draw.

For a PSE (Network Connectivity device), the power level represents the amount of power that is available on the port at the time. If the PSE is operating in reduced power (i.e., it is using backup power), the reduced power capacity is advertised as long as the condition persists.

By default, LLDP-MED power-via-MDI information is automatically advertised when LLDP-MED is enabled, the port is a POE port, and POE is enabled on the port. To disable this advertisement, enter a command such as the following.

```
device(config)#no lldp advertise med-power-via-mdi ports e 1/2/4 to 1/2/12
```

The LLDP-MED power-via-MDI advertisement will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
+ MED Extended Power via MDI
  Power Type      : PSE device
  Power Source    : Unknown Power Source
  Power Priority   : Low (3)
  Power Value     : 6.5 watts (PSE equivalent: 7005 mWatts)
```

Syntax:[no] lldp advertise med-power-via-mdi ports ethernet *port-list* | all

Displaying LLDP statistics and configuration settings

You can use the following CLI **show** commands to display information about LLDP settings and statistics:

- **show lldp** - Displays a summary of the LLDP configuration settings.
- **show lldp statistics** - Displays LLDP global and per-port statistics.
- **show lldp neighbors** - Displays a list of the current LLDP neighbors.
- **show lldp neighbors detail** - Displays the details of the latest advertisements received from LLDP neighbors.
- **show lldp local-info** - Displays the details of the LLDP advertisements that will be transmitted on each port.

This above **show** commands are described in this section.

LLDP configuration summary

To display a summary of the LLDP configuration settings on the device, enter the **show lldp** command at any level of the CLI.

The following shows an example report.

```
device#show lldp
LLDP transmit interval      : 10 seconds
LLDP transmit hold multiplier : 4 (transmit TTL: 40 seconds)
LLDP transmit delay        : 1 seconds
LLDP SNMP notification interval : 5 seconds
LLDP reinitialize delay    : 1 seconds
LLDP-MED fast start repeat count : 3
LLDP maximum neighbors     : 392
LLDP maximum neighbors per port : 4
```

Syntax: show lldp

The following table describes the information displayed by the **show lldp statistics** command.

Field	Description
LLDP transmit interval	The number of seconds between regular LLDP packet transmissions.
LLDP transmit hold multiplier	The multiplier used to compute the actual time-to-live (TTL) value of an LLDP advertisement. The TTL value is the transmit interval multiplied by the transmit hold multiplier.
LLDP transmit delay	The number of seconds the LLDP agent will wait after transmitting an LLDP frame and before transmitting another LLDP frame.
LLDP SNMP notification interval	The number of seconds between transmission of SNMP LLDP traps (lldpRemTablesChange) and SNMP LLDP-MED traps (lldpXMedTopologyChangeDetected).
LLDP reinitialize delay	The minimum number of seconds the device will wait from when LLDP is disabled on a port, until a request to re-enable LLDP on that port will be honored.
LLDP-MED fast start repeat count	The number of seconds between LLDP frame transmissions when an LLDP-MED Endpoint is newly detected.
LLDP maximum neighbors	The maximum number of LLDP neighbors for which LLDP data will be retained, per device.
LLDP maximum neighbors per port	The maximum number of LLDP neighbors for which LLDP data will be retained, per port.

Displaying LLDP statistics

The **show lldp statistics** command displays an overview of LLDP neighbor detection on the device, as well as packet counters and protocol statistics. The statistics are displayed on a global basis.

The following shows an example report.

```
device#show lldp statistics
Last neighbor change time: 23 hours 50 minutes 40 seconds ago
Neighbor entries added      : 14
Neighbor entries deleted    : 5
Neighbor entries aged out   : 4
Neighbor advertisements dropped : 0
Port      Tx Pkts  Rx Pkts  Rx Pkts  Rx Pkts  Rx TLVs  Rx TLVs  Neighbors
          Total   Total   w/Errors Discarded Unrecognz Discarded Aged Out
1         60963  75179   0         0         0         0         0         4
2         0      0         0         0         0         0         0         0
3         60963  60963   0         0         0         0         0         0
4         60963  121925  0         0         0         0         0         0
5         0      0         0         0         0         0         0         0
6         0      0         0         0         0         0         0         0
7         0      0         0         0         0         0         0         0
8         0      0         0         0         0         0         0         0
9         0      0         0         0         0         0         0         0
10        60974  0         0         0         0         0         0         0
11        0      0         0         0         0         0         0         0
12        0      0         0         0         0         0         0         0
13        0      0         0         0         0         0         0         0
14        0      0         0         0         0         0         0         0
```

Syntax: show lldp statistics

NOTE

You can reset LLDP statistics using the CLI command **clear LLDP statistics**. Refer to [Resetting LLDP statistics](#) on page 209.

The following table describes the information displayed by the **show lldp statistics** command.

Field	Description
Last neighbor change time	The elapsed time (in hours, minutes, and seconds) since a neighbor last advertised information. For example, the elapsed time since a neighbor was last added, deleted, or its advertised information changed.
Neighbor entries added	The number of new LLDP neighbors detected since the last reboot or since the last time the clear lldp statistics all command was issued.
Neighbor entries deleted	The number of LLDP neighbors deleted since the last reboot or since the last time the clear lldp statistics all command was issued.
Neighbor entries aged out	The number of LLDP neighbors dropped on all ports after the time-to-live expired. Note that LLDP entries age out naturally when a port cable or module is disconnected or when a port becomes disabled. However, if a disabled port is re-enabled, the system will delete the old LLDP entries.
Neighbor advertisements dropped	The number of valid LLDP neighbors the device detected, but could not add. This can occur, for example, when a new neighbor is detected and the device is already supporting the maximum number of neighbors possible. This can also occur when an LLDPDU is missing a mandatory TLV or is not formatted correctly.
Port	The local port number.
Tx Pkts Total	The number of LLDP packets the port transmitted.
Rx Pkts Total	The number of LLDP packets the port received.

Field	Description
Rx Pkts w/Errors	The number of LLDP packets the port received that have one or more detectable errors.
Rx Pkts Discarded	The number of LLDP packets the port received then discarded.
Rx TLVs Unrecognz	The number of TLVs the port received that were not recognized by the LLDP local agent. Unrecognized TLVs are retained by the system and can be viewed in the output of the show LLDP neighbors detail command or retrieved through SNMP.
Rx TLVs Discarded	The number of TLVs the port received then discarded.
Neighbors Aged Out	The number of times a neighbor information was deleted because its TTL timer expired.

Displaying LLDP neighbors

The **show lldp neighbors** command displays a list of the current LLDP neighbors per port.

The following shows an example report.

```
device#show lldp neighbors
Lcl Port Chassis ID      Port ID      Port Description      System Name
1         0000.0034.0fc0 0000.0034.0fc0 GigabitEthernet9/1    FastIron Supe~
1         0000.0001.4000 0000.0001.4000 GigabitEthernet0/1/1  FastIron SX Swi~
3         0000.0011.0200 0000.0011.0203 GigabitEthernet4      FastIron SX 8~
4         0000.0011.0200 0000.0011.0202 GigabitEthernet3      FastIron SX 8~
4         0000.0011.0200 0000.0011.0210 GigabitEthernet17     FastIron SX 8~
15        0000.0011.0200 0000.0011.020f GigabitEthernet16     FastIron SX 8~
16        0000.0011.0200 0000.0011.020e GigabitEthernet15     FastIron SX 8~
17        0000.0011.0200 0000.0011.0211 GigabitEthernet18     FastIron SX 8~
18        0000.0011.0200 0000.0011.0210 GigabitEthernet17     FastIron SX 8~
```

Syntax:show lldp neighbors

The following table describes the information displayed by the **show lldp neighbors** command.

Field	Description
Lcl Port	The local LLDP port number.
Chassis ID	The identifier for the chassis. Brocade devices use the base MAC address of the device as the Chassis ID.
Port ID	The identifier for the port. Brocade devices use the permanent MAC address associated with the port as the port ID.
Port Description	The description for the port. Brocade devices use the ifDescr MIB object from MIB-II as the port description.
System Name	The administratively-assigned name for the system. Brocade devices use the sysName MIB object from MIB-II, which corresponds to the CLI hostname command setting. NOTE A tilde (~) at the end of a line indicates that the value in the field is too long to display in full and is truncated.

Displaying LLDP neighbors detail

The **show lldp neighbors detail** command displays the LLDP advertisements received from LLDP neighbors.

The following shows an example **show lldp neighbors detail** report.

NOTE

The **show lldp neighbors detail** output will vary depending on the data received. Also, values that are not recognized or do not have a recognizable format, may be displayed in hexadecimal binary form.

```
device#show lldp neighbors detail ports e 1/1/9
Local port: 1/1/9
Neighbor: 0000.0018.cc03, TTL 101 seconds
+ Chassis ID (network address): 10.43.39.151
+ Port ID (MAC address): 0000.0018.cc03
+ Time to live: 120 seconds
+ Port description      : "LAN port"
+ System name          : "regDN 1015,MITEL 5235 DM"
+ System description   : "regDN 1015,MITEL 5235 DM,h/w rev 2,ASIC rev 1,f/w\
                          Boot 02.01.00.11,f/w Main 02.01.00.11"
+ System capabilities  : bridge, telephone
  Enabled capabilities: bridge, telephone
+ Management address (IPv4): 10.43.39.151
+ 802.3 MAC/PHY       : auto-negotiation enabled
  Advertised capabilities: 10BaseT-HD, 10BaseT-FD, 100BaseTX-HD,
                          100BaseTX-FD
  Operational MAU type   : 100BaseTX-FD
+ MED capabilities: capabilities, networkPolicy, extendedPD
  MED device type       : Endpoint Class III
+ MED Network Policy
  Application Type      : Voice
  Policy Flags          : Known Policy, Tagged
  VLAN ID               : 300
  L2 Priority            : 7
  DSCP Value            : 7
+ MED Extended Power via MDI
  Power Type            : PD device
  Power Source          : Unknown Power Source
  Power Priority         : High (2)
  Power Value           : 6.2 watts (PSE equivalent: 6656 mWatts)
+ MED Hardware revision : "PCB Version: 2"
+ MED Firmware revision : "Boot 02.01.00.11"
+ MED Software revision : "Main 02.01.00.11"
+ MED Serial number     : ""
+ MED Manufacturer      : "Mitel Corporation"
+ MED Model name        : "MITEL 5235 DM"
+ MED Asset ID          : ""
```

A backslash (\) at the end of a line indicates that the text continues on the next line.

Except for the following field, the fields in the above output are described in the individual TLV advertisement sections in this chapter.

Field	Description
Neighbor	The source MAC address from which the packet was received, and the remaining TTL for the neighbor entry.

Syntax: `show lldp neighbors detail [ports ethernet port-list | all]`

If you do not specify any ports or use the keyword **all**, by default, the report will show the LLDP neighbor details for all ports.

Displaying LLDP configuration details

The **show lldp local-info** command displays the local information advertisements (TLVs) that will be transmitted by the LLDP agent.

NOTE

The `show lldp local-info` output will vary based on LLDP configuration settings.

The following shows an example report.

```

device#show lldp local-info ports e 20
Local port: 20
+ Chassis ID (MAC address): 0000.0033.e2c0
+ Port ID (MAC address): 0000.0033.e2d3
+ Time to live: 40 seconds
+ System name: "FCX624SHPOE-ADV Router"
+ Port description: "GigabitEthernet20"
+ System description : "Brocade Communications,
Inc.

FCX_ADV_ROUTER_SOFT_PACKAGE,
IronWare Version 07.3.00T7f3 compiled on Sep 26 2011 at 21:15:14 labeled
as FCXR07300"
+ System capabilities : bridge
Enabled capabilities: bridge
+ 802.3 MAC/PHY : auto-negotiation enabled
Advertised capabilities: 10BaseT-HD, 10BaseT-FD, 100BaseTX-HD,
100BaseTX-FD, fdxSPause, fdxBPause, 1000BaseT-HD,
1000BaseT-FD
Operational MAU type: 100BaseTX-FD
+ 802.3 Power via MDI: PSE port, power enabled, class 2
Power Pair : A (not controllable)
+ Link aggregation: not capable
+ Maximum frame size: 1522 octets
+ MED capabilities: capabilities, networkPolicy, location, extendedPSE
MED device type : Network Connectivity
+ MED Network Policy
Application Type : Voice
Policy Flags : Known Policy, Tagged
VLAN ID : 99
L2 Priority : 3
DSCP Value : 22
+ MED Network Policy
Application Type : Video Conferencing
Policy Flags : Known Policy, Tagged
VLAN ID : 100
L2 Priority : 5
DSCP Value : 10
+ MED Location ID
Data Format: Coordinate-based location
Latitude Resolution : 20 bits
Latitude Value : -78.303 degrees
Longitude Resolution : 18 bits
Longitude Value : 34.27 degrees
Altitude Resolution : 16 bits
Altitude Value : 50. meters
Datum : WGS 84
+ MED Location ID
Data Format: Civic Address
Location of: Client
Country : "US"
CA Type : 1
CA Value : "CA"
CA Type : 3
CA Value : "Santa Clara"
CA Type : 6
CA Value : "4980 Great America Pkwy."
CA Type : 24
CA Value : "95054"
CA Type : 27
CA Value : "5"
CA Type : 28
CA Value : "551"
CA Type : 29

```



```

CA Value   : "office"
CA Type    : 23
CA Value   : "John Doe"
+ MED Location ID
  Data Format: ECS ELIN
  Value      : "1234567890"
+ MED Extended Power via MDI
  Power Type   : PSE device
  Power Source : Unknown Power Source
  Power Priority : Low (3)
  Power Value  : 6.5 watts (PSE equivalent: 7005 mWatts) + Port VLAN ID: 99
+ Management address (IPv4): 10.1.1.121
+ VLAN name (VLAN 99): "Voice-VLAN-99"

```

NOTE

The contents of the **show** output will vary depending on which TLVs are configured to be advertised.

A backslash (\) at the end of a line indicates that the text continues on the next line.

The fields in the above output are described in the individual TLV advertisement sections in this chapter.

Syntax: `show lldp local-info [ports ethernet port-list | all]`

If you do not specify any ports or use the keyword **all**, by default, the report will show the local information advertisements for all ports.

Resetting LLDP statistics

To reset LLDP statistics, enter the **clear lldp statistics** command at the Global CONFIG level of the CLI. The Brocade device will clear the global and per-port LLDP neighbor statistics on the device (refer to [Displaying LLDP statistics](#) on page 205).

```
device#clear lldp statistics
```

Syntax: `clear lldp statistics [ports ethernet port-list | all]`

If you do not specify any ports or use the keyword **all**, by default, the system will clear lldp statistics on all ports.

Clearing cached LLDP neighbor information

The Brocade device clears cached LLDP neighbor information after a port becomes disabled and the LLDP neighbor information ages out. However, if a port is disabled then re-enabled before the neighbor information ages out, the device will clear the cached LLDP neighbor information when the port is re-enabled.

If desired, you can manually clear the cache. For example, to clear the cached LLDP neighbor information for port e 20, enter the following command at the Global CONFIG level of the CLI.

```
device#clear lldp neighbors ports e 20
```

Syntax: `clear lldp neighbors [ports ethernet port-list | all]`

If you do not specify any ports or use the keyword **all**, by default, the system will clear the cached LLDP neighbor information for all ports.

Hardware Component Monitoring

- Traffic Limitations in Mixed Environments..... 211
- Virtual cable testing..... 211
- Digital optical monitoring..... 214
- FastIron Fiber-optic Transceivers..... 219

Traffic Limitations in Mixed Environments

Packet loss can occur in environments that mix SX hardware modules of different generations.

The following table lists the modules by generation:

TABLE 32 SX Hardware Generations

First	Second	Third
SX-FI2XGMR4	SX-FI2XGMR6	SX-FI48GPP
SX-FI2XGMR4-PREM	SX-FI2XGMR6-PREM	SX-FI-2XG
SX-FI424100FX	SX-FI2XGMR6-PREM6	SX-FI-8XG
SX-FI42XG-BNDL-2CX4	SX-FI624100FX	SX-FI-24HF
SX-FI424C	SX-FI624C	SX-FI-24GPP
SX-FI424P	SX-FI624HF	
SX-FI424F	SX-FI624P	
SX-FI424HF	SX-FI62XG	
SX-FI42XG		

Throughput is 100 percent when only SX third-generation modules are used.

Virtual cable testing

Most FastIron devices support Virtual Cable Test (VCT) technology. VCT technology enables the diagnosis of a conductor (wire or cable) by sending a pulsed signal into the conductor, then examining the reflection of that pulse. This method of cable analysis is referred to as Time Domain Reflectometry (TDR). By examining the reflection, the Brocade device can detect and report cable statistics such as local and remote link pair, cable length, and link status.

Virtual Cable Testing configuration notes

- VCT is supported on copper ports only. It is not supported on fiber ports.
- VCT is only supported when Ethernet port speed is configured to Auto. VCT does not work on ports with fixed speeds.
- VCT is not supported on the following:
 - ICX 6610-24F
 - SX-FI24GPP
 - SX-FI48GPP
 - SX-FI2XG

- SX-FI8XG
- SX-FI24HF
- The port to which the cable is connected must be enabled when you issue the command to diagnose the cable. If the port is disabled, the command is rejected.
- If the port is operating at 100 Mbps half-duplex, the TDR test on one pair will fail.
- If the remote pair is set to forced 100 Mbps, any change in MDI/MDIX may cause the device to interpret the Multilevel Threshold-3 (MLT-3) as a reflected pulse, in which case, the device will report a faulty condition. In this scenario, it is recommended that you run the TDR test a few times, clearing the registers before each test, for accurate results.

Virtual Cable Test command syntax

To diagnose a cable using TDR, enter commands such as the following at the Privileged EXEC level of the CLI.

```
device# phy cable-diagnostics tdr 1/1/1
```

The **clear-diag tdr** command clears results of any previous TDR test from test registers for port 1/1/1 (port 1 on slot 1 on device 1).

NOTE

It is recommended that you clear the TDR test registers before each test.

```
device# clear cable-diagnostics tdr 1/1/1
```

The command in the previous example diagnoses the cable attached to port 1/1/1.

When you issue the **phy cable-diagnostics** command, the command brings the port down for a second or two, and then immediately brings the port back up.

Syntax: clear cable-diagnostics tdr *stackid/slot/port*

Syntax: phy cable-diagnostics tdr *stackid/slot/port*

Viewing the results of the cable analysis

To display the results of the cable analysis, enter a command such as the one shown in the following examples at the Privileged EXEC level of the CLI.

In the first example, the command displays TDR test results for port 1, slot 1 on device 1 in the stack. The results indicate that the port is down or the cable is not connected.

```
device>show cable-diagnostics tdr 1/1/1
Port      Speed Local pair Pair Length Remote pair Pair status
-----
01        UNKWN                Pair A    <=3 M                Open
                Pair B    <=3 M                Open
                Pair C    <=3 M                Open
                Pair D    <=3 M                Open
```

In the second test example, the TDR test results for the same port show details for an active port.

```
device>show cable-diagnostics tdr 1/1/1
Port      Speed Local pair Pair Length Remote pair Pair status
-----
01        1000M                Pair A    <50M                Pair B    Terminated
                Pair B    <50M                Pair A    Terminated
                Pair C    <50M                Pair D    Terminated
                Pair D    <50M                Pair C    Terminated
```

Syntax: `show cable-diagnostics tdr stackid/slot/port`


In the output shown, "Local pair" indicates the assignment of wire pairs from left to right, where Pair A is the left-most pair. The following table shows the "Local pair" mapping to the T568A pin/pair and color assignment from the TIA/EIA-568-B standard.

TABLE 33 Local pair definition

Local pair	T568A pair and color assignment
Pair A	Pair 3 (green)
Pair B	Pair 2 (orange)
Pair C	Pair 1 (blue)
Pair D	Pair 4 (brown)

The following figure illustrates the T568A pin/pair assignment.

FIGURE 10 T568A pin/pair assignment



The following table describes the fields shown in the `show cable-diagnostics port` command output.

TABLE 34 Cable statistics

Field	Meaning
Port	The port that was tested.
Speed	The port current line speed.
Local pair	The local link name. Refer to the previous local pair definition table.
Pair Length	The cable length when terminated, or the distance to the point of fault when the line is not up.

TABLE 34 Cable statistics (continued)

Field	Meaning
Remote pair	The remote link name.
Pair status	The status of the link. This field displays one of the following: <ul style="list-style-type: none"> Terminated: The link is up. Shorted: A short is detected in the cable. Open: An opening is detected in the cable. ImpedMis: The impedance is mismatched. Failed: The TDR test failed.

Digital optical monitoring

You can configure your Brocade device to monitor optical transceivers in the system, either globally or by specified ports. When this feature is enabled, the system will monitor the temperature and signal power levels for the optical transceivers in the specified ports. Console messages and Syslog messages are sent when optical operating conditions fall below or rise above the XFP, SFP, and SFP+ manufacturer recommended thresholds.

Digital optical monitoring configuration limitations

- A Brocade chassis device can monitor a maximum of 24 SFPs and 12 XFPs.
- Brocade ICX 6650, ICX 7450, and ICX 7750 devices allow all ports to support Digital Optical Monitoring (DOM).
- Brocade FCX, FSX (SX), ICX 6610, and ICX 64x0 all support up to 24 ports running DOM.
- Only the FCX 6xx-S and FSX (SX) products support XFP media.

Enabling digital optical monitoring

To enable optical monitoring on all Brocade-qualified optics installed in the device, use the following command.

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

To enable optical monitoring on a specific port, use the following command.

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

To enable optical monitoring on a range of ports, use the following command.

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

Syntax: [no] optical-monitor

Use the **no** form of the command to disable digital optical monitoring.

Setting the alarm interval

You can optionally change the interval between which alarms and warning messages are sent.

To change the interval, use the following command.

```
device(config)#interface ethernet 1/1/1 to 1/1/2
device(config-mif-e10000-1/1/1-1/1/2)#optical-monitor 10
```

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

For *alarm-interval*, the default is 3 minutes with a 1 minute minimum on all devices except the ICX 6650, ICX 7450, and ICX 7750. These devices have a default and minimum value of 8 minutes. The maximum value for all devices is 65535.

NOTE

The commands `no optical-monitor` and `optical-monitor 0` perform the same function. That is, they both disable digital optical monitoring.

Displaying information about installed media

Use the `show media`, `show media slot`, and `show media ethernet` commands to obtain information about the media devices installed per device, per slot, and per port. The results displayed from these commands provide the Type, Vendor, Part number, Version and Serial number of the SFP, SFP+, or XFP optical device installed in the port. If there is no SFP, SFP+, or XFP optical device installed in a port, the "Type" field will display "EMPTY".

On ICX 6430 and ICX 6450 devices, 1G copper ports will always be shown with the type as 1G M-C (Gig-Copper), even if the ports are not connected.

Use the `show media` command to obtain information about the media devices installed in a device.

```
device# show media
Port 1/1/1: Type : 1G M-C (Gig-Copper)
Port 1/1/2: Type : 1G M-C (Gig-Copper)
Port 1/1/3: Type : 1G M-C (Gig-Copper)
Port 1/1/4: Type : 1G M-C (Gig-Copper)
Port 1/1/5: Type : 1G M-C (Gig-Copper)
Port 1/1/6: Type : 1G M-C (Gig-Copper)
Port 1/1/7: Type : 1G M-C (Gig-Copper)
Port 1/1/8: Type : 1G M-C (Gig-Copper)
Port 1/1/9: Type : 1G M-C (Gig-Copper)
Port 1/1/10: Type : 1G M-C (Gig-Copper)
Port 1/1/11: Type : 1G M-C (Gig-Copper)
Port 1/1/12: Type : 1G M-C (Gig-Copper)
Port 1/1/13: Type : 1G M-C (Gig-Copper)
Port 1/1/14: Type : 1G M-C (Gig-Copper)
Port 1/1/15: Type : 1G M-C (Gig-Copper)
Port 1/1/16: Type : 1G M-C (Gig-Copper)
Port 1/1/17: Type : 1G M-C (Gig-Copper)
Port 1/1/18: Type : 1G M-C (Gig-Copper)
Port 1/1/19: Type : 1G M-C (Gig-Copper)
Port 1/1/20: Type : 1G M-C (Gig-Copper)
Port 1/1/21: Type : 1G M-C (Gig-Copper)
Port 1/1/22: Type : 1G M-C (Gig-Copper)
Port 1/1/23: Type : 1G M-C (Gig-Copper)
Port 1/1/24: Type : 1G M-C (Gig-Copper)
Port 1/2/1: Type : 10GE SR 300m (SFP +)
Port 1/2/2: Type : EMPTY
Port 1/2/3: Type : 1G Twinax 1m (SFP)
Port 1/2/4: Type : 1G Twinax 1m (SFP)
```

Use the `show media slot` command to obtain information about the media device installed in a slot.

```
device# show media slot 1
Port 1/1/1: Type : 1G M-SX(SFP)
Vendor: Brocade Communications, Inc. Version:
Part# : PL-XPL-VC-S13-19 Serial#: 425HC109
Port 1/1/2: Type : 1G M-SX(SFP)
Vendor: Brocade Communications, Inc. Version:
Part# : PL-XPL-VC-S13-19 Serial#: 411HC0AH
Port 1/1/3: Type : EMPTY
Port 1/1/4: Type : 1G M-SX(SFP)
Vendor: Brocade Communications, Inc. Version: X1
Part# : FTRJ-8519-3 Serial#: H11654K
Port 1/1/5: Type : EMPTY
```

```

Port 1/1/6: Type : EMPTY
Port 1/1/7: Type : 100M M-FX-IR (SFP)
Vendor: Brocade Communications, Inc. Version: A
Part# : FTLF1323P1BTR-FD Serial#: UCT000T
Port 1/1/8: Type : EMPTY
Port 1/1/9: Type : 100M M-FX-LR (SFP)
Vendor: Brocade Communications, Inc. Version: A
Part# : FTLF1323P1BTL-FD Serial#: UD3085J
Port 1/1/10: Type : EMPTY
Port 1/1/11: Type : 100M M-FX-SR (SFP)
Vendor: Brocade Communications, Inc. Version: A
Part# : FTLF1217P2BTL-F1 Serial#: UCQ003J
Port 1/1/12: Type : EMPTY
Port 1/1/13: Type : 100M M-FX-IR (SFP)
Vendor: Brocade Communications, Inc. Version: A
Part# : FTLF1323P1BTR-F1 Serial#: PCA2XC5

```

Use the **show media ethernet** command to obtain information about the media device installed in a port.

```

device# show media e 1/1/17
Port 1/1/17: Type : 1G M-SX (SFP)
Vendor: Brocade Communications, Inc. Version:
Part# : PL-XPL-VC-S13-19 Serial#: 425HC109

```

Use the **show media validation** command to find out whether the connected optic modules are supported or not on Brocade devices.

```

device# show media validation
Port      Supported Vendor      Type
-----
1/2/1     Yes      FINISAR CORP.      1GE M-SX (SFP)
1/2/2     Yes      BROCADE             10GE Twinax 1m (SFP +)
2/2/1     Yes      BROCADE             10GE SR 300m (SFP +)
2/2/3     Yes      BROCADE             10GE SR 300m (SFP +)

```

Syntax: `show media [validation [ethernet [unit / slot / port] | slot slot-num | stack stack-id] | ethernet [unit / slot / port] | slot slot-num | stack stack-id]`

Viewing optical monitoring information

You can view temperature and power information for qualified XFPs, SFPs, and SFP+ installed in a FastIron device.

Use the **show optic** command to view information about an XFP, SFP, or SFP+ installed in a particular port. The following shows example output.

Optical monitoring feature will not work in the following scenarios:

- The port is DOWN.
- The port is configured as a stacking port.
- The the optic module does not support optical monitoring.
- For ICX 6430 devices only:
 - If an SFP+ optic is inserted in an SFP only port, the optic will not initialize.
 - If an SFP optic is inserted in an SFP+ only port, the optic will not initialize.
 - If an optic is inserted into a device that supports both SFP and SFP+ optics, use the **speed-duplex** command to set the port speed correctly.

```

device#show optic 13
Port Temperature Tx Power Rx Power Tx Bias Current
+-----+-----+-----+-----+
13 33.2968 C -005.4075 dBm -007.4328 dBm 6.306 mA
Normal Normal Normal Normal

```

Syntax: `showoptic port-number`

Use the **show optic slot** on a FastIron X Series chassis to view information about all qualified XFPs, SFPs, and SFP+ in a particular slot. The following shows example output.

```
device>show optic slot 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          Low-Alarm    -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
```

Syntax: `show optic slot slot-number`

NOTE

The **show optic slot** command is supported on the FSX 800 and FSX 1600 only.

NOTE

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

The following table describes the information displayed by the **show optic** command.

TABLE 35 Output from the show optic command

Field	Description
Port	The Brocade port number.
Temperature	<ul style="list-style-type: none"> The operating temperature, in degrees Celsius, of the optical transceiver. The alarm status, as described in the next table.
Tx Power	<ul style="list-style-type: none"> The transmit power signal, in decibels (dB), of the measured power referenced to one milliwatt (mW). The alarm status, as described in the next table.
Rx Power	<ul style="list-style-type: none"> The receive power signal, in decibels (dB), of the measured power referenced to one milliwatt (mW). The alarm status, as described in the next table.
Tx Bias Current	<ul style="list-style-type: none"> The transmit bias power signal, in milliamperes (mA). The alarm status, as described in the next table.

For Temperature, Tx Power, Rx Power, and Tx Bias Current in the **show optic** command output, values are displayed along with one of the following alarm 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 optical transceivers. The following table describes each of these status values.

TABLE 36 Alarm status value description

Status value	Description
Low-Alarm	Monitored level has dropped below the "low-alarm" threshold set by the manufacturer of the optical transceiver.
Low-Warn	Monitored level has dropped below the "low-warn" threshold set by the manufacturer of the optical transceiver.

TABLE 36 Alarm status value description (continued)

Status value	Description
Normal	Monitored level is within the "normal" range set by the manufacturer of the optical transceiver.
High-Warn	Monitored level has climbed above the "high-warn" threshold set by the manufacturer of the optical transceiver.
High-Alarm	Monitored level has climbed above the "high-alarm" threshold set by the manufacturer of the optical transceiver.

Viewing optical transceiver thresholds

The thresholds that determine the alarm status values for an optical transceiver are set by the manufacturer of the XFP, SFP, or SFP+. To view the thresholds for a qualified optical transceiver in a particular port, use the **show optic threshold** command as shown below.

```
device>show optic threshold 1/2/2
Port 1/2/2 sfp monitor thresholds:
Temperature High alarm      5a00      90.0000 C
Temperature Low alarm       d300     -45.0000 C
Temperature High warning    5500      85.0000 C
Temperature Low warning     d800     -40.0000 C
Supply Voltage High alarm   9088
Supply Voltage Low alarm    7148
Supply Voltage High warning 8ca0
Supply Voltage Low warning  7530
TX Bias High alarm          7530      60.000 mA
TX Bias Low alarm           01f4      1.000 mA
TX Bias High warning        61a8      50.000 mA
TX Bias Low warning         05dc      3.000 mA
TX Power High alarm         1f07     -001.0001 dBm
TX Power Low alarm          02c4     -011.4996 dBm
TX Power High warning       18a6     -001.9997 dBm
TX Power Low warning        037b     -010.5012 dBm
RX Power High alarm         2710      000.0000 dBm
RX Power Low alarm          0028     -023.9794 dBm
RX Power High warning       1f07     -001.0001 dBm
RX Power Low warning        0032     -023.0102 dBm
```

Syntax:show optic threshold port

For Temperature, Supply Voltage, TX Bias, TX Power, and RX Power, values are displayed for each of the following four alarm and warning settings: High alarm, Low alarm, High warning, and Low warning. The hexadecimal values are the manufacturer internal calibrations, as defined in the SFF-8472 standard. The other values indicate at what level (above the high setting or below the low setting) the system should send a warning message or an alarm. Note that these values are set by the manufacturer of the optical transceiver, and cannot be configured.

Syslog messages for optical transceivers

The system generates Syslog messages for optical transceivers in the following circumstances:

- The temperature, supply voltage, TX Bias, TX power, or TX power value goes above or below the high or low warning or alarm threshold set by the manufacturer.
- The optical transceiver does not support digital optical monitoring.
- The optical transceiver is not qualified, and therefore not supported by Brocade.

For details about the above Syslog messages, refer to [Syslog messages for optical transceivers](#).

FastIron Fiber-optic Transceivers

The fiber-optic transceivers listed in the following table are supported on all FastIron devices.

TABLE 37 FastIron fiber-optic transceivers

Label	Type	Brocade part number	Supports Digital Optical Monitoring?
E1MG-BXD	1000Base-BXD	33005-000	No
E1MG-BXU	1000Base-BXU	33006-000	No
E1MG-LHA-OM	1000Base-LHA	33212-100	Yes
E1MG-LX-OM	1000Base-LX	33211-100	Yes
E1MG-100FX-LR-OM	100Base-FX-LR, 40 km	33226-100	Yes
E1MG-100FX-OM	100Base-FX	33224-100	Yes
E1MG-100FX-IR-OM	100Base-FX-IR, 15 km	33225-100	Yes
E1MG-SX-OM	1000Base-SX	33210-100	Yes
E1MG-TX	1000Base-T Copper	33002-100	No
10G-XFP-ER	10GBase-ER XFP, 40 km	33013-000	Yes
10G-XFP-LR	10GBase-LR XFP, 10 km	33012-000	Yes
10G-XFP-SR	10GBase-SR XFP	33011-000	Yes
10G-XFP-ZR	10GBase-ZR XFP, 80 km	33014-000	Yes
10G-XFP-ZRD	10GBase-ZRD XFP, 80 km	33063-000 to 33107-000	Yes
10G-SFPP-SR	10GE SR SFP+	57-0000075-01	Yes
10G-SFPP-LR	10GE LR SFP+	57-0000076-01	Yes
10G-SFPP-TWX-0101	FCoE 1M Active Cable	58-1000026-01	No
10G-SFPP-TWX-0301	FCoE 3M Active Cable	58-1000027-01	No
10G-SFPP-TWX-0501	FCoE 5M Active Cable	58-1000023-01	No
10G-SFPP-ER	10GBase-ER SFP+, 40 km	57-0000085-01	Yes
10G-SFPP-LRM	10GBase-LRM SFP+	57-0000084-01	Yes
E1MG-LHB	1000Base-LHB	33004-000	No
10G-SFPP-USR	10GE Ultra Short Reach (USR) SFP+ 100m on OM3 MMF	57-1000130-01	Yes
10GE ZR SFPP	10GE ZR 80km ((SFP+)) 1550.0 nm Used on Port 1/3/8	57-1000180-01	Yes
40G-QSFP-C-0101	40GE QSFP Direct Attached Copper Cable, 1m (stacking) Used for stacking only.	58-0000033-01	No
40G-QSFP-C-0501	40GE QSFP Direct Attached Copper Cable, 5m (stacking) Used for stacking only.	58-0000035-01	No
40Ge LR4	40GE-LR4 10km (QSFP + LC)	57-1000263-01	Yes
40GE-SR4	100m (QSFP+)	57-1000128-1	Yes

Network Monitoring

- Basic system management.....221
- RMON support.....231
- sFlow.....235
- Utilization list for an uplink port.....251

Basic system management

The following sections contain procedures for basic system management tasks.

Viewing system information

You can access software and hardware specifics for a Brocade Layer 2 switch or Layer 3 switch. For software specifics, refer to the section [Software versions installed and running on a device](#) on page 86 in the *FastIron Ethernet Switch Administration Guide*.

To view the software and hardware details for the system, enter the **show version** command. The following shows example output.

NOTE

The output may vary depending on the device on which you run the **show version**.

```
device# show version
=====
Active Management CPU [Slot-9]:
  SW: Version 04.3.00b17T3e3 Copyright (c) 1996-2008 Brocade Communications, Inc., Inc.
      Compiled on Sep 25 2008 at 04:09:20 labeled as SXR04300b17
      (4031365 bytes) from Secondary sxr04300b17.bin
      BootROM: Version 04.0.00T3e5 (FEv2)
  HW: ANR-Chassis FastIron SX 1600-PREM (PROM-TYPE SX-FIL3U)
      Serial #: TExxxxxxxx
=====
SL 3: SX-FI424C 24-port Gig Copper
      Serial #: CYxxxxxxxx
      P-ASIC 4: type 00D1, rev D2 subrev 00
      P-ASIC 5: type 00D1, rev D2 subrev 00
=====
SL 9: SX-FI8GMR4 8-port Management
      Serial #: CHxxxxxxxx
      P-ASIC 16: type 00D1, rev D2 subrev 00
=====
SL 14: SX-FI42XGW 2-port 10G LAN/WAN
      Serial #: Invalid
      P-ASIC 26: type 01D1, rev 00 subrev 00
      P-ASIC 27: type 01D1, rev 00 subrev 00
=====
Active Management Module:
  660 MHz Power PC processor 8541 (version 32/0020) 66 MHz bus
  512 KB boot flash memory
  16384 KB code flash memory
  512 MB DRAM
The system uptime is 2 minutes 13 seconds
The system : started=warm start   reloaded=by "reload"
*** NOT FOR PRODUCTION ***
*** AUTO SHUTDOWN IS OFF. PLEASE ACTIVATE WITH auto-shutdown ***
```

The following hardware details are listed in the output of the **show version** command:

- Chassis type
- PROM type (if applicable)

- Chassis serial number
- Management and interface module serial numbers and ASIC types

For a description of the software details in the output of the **show version** command, refer to the section [Software versions installed and running on a device](#) on page 86 in the *FastIron Ethernet Switch Administration Guide*.

Starting with FastIron 08.0.30, you can view the serial number pluggable modules. If there are no pluggable modules on the device, the serial number of the fixed modules on the device is displayed. The following is an example of the **show version** output on an ICX 7750.

```
device# show version
Copyright (c) 1996-2014 Brocade Communications Systems, Inc. All rights reserved.
UNIT 1: compiled on Dec 22 2014 at 12:35:56 labeled as SWR08030b1
(20833985 bytes) from Secondary SWR08030b1.bin
SW: Version 08.0.30b1T203
UNIT 2: compiled on Dec 22 2014 at 12:35:56 labeled as SWR08030b1
(20833985 bytes) from Secondary SWR08030b1.bin
SW: Version 08.0.30b1T203
Compressed Boot-Monitor Image size = 1835008, Version:10.1.03T205 (swz10103b003)
HW: Stackable ICX7750-26Q
Internal USB: Serial #: 40D41E003CF90029
Vendor: UNIGEN, Total size = 1910 MB
=====
UNIT 1: SL 1: ICX7750-20QXG 20-port Management Module
Serial #:CRK2234J00V
License: ICX7750_L3_SOFT_PACKAGE (LID: etmHHIJlFFx)
P-ASIC 0: type B850, rev 03 Chip BCM56850_A2
=====
UNIT 1: SL 2: ICX7750-QSFP 6-port QSFP 240G Module
=====
UNIT 1: SL 3: ICX7750-6Q 6-port QSFP 240G Module
Serial #:PR320400290
=====
UNIT 2: SL 1: ICX7750-48XGF 48-port Management Module
Serial #:CRH2234J00M
License: ICX7750_L3_SOFT_PACKAGE (LID: etjHHIJlFFo)
=====
UNIT 2: SL 2: ICX7750-QSFP 6-port QSFP 240G Module
=====
UNIT 2: SL 3: ICX7750-6Q 6-port QSFP 240G Module
Serial #:PR320400289
=====
1500 MHz Power PC processor (version 8023/0022) 88 MHz bus
8192 KB boot flash memory
2048 MB code flash memory
256 MB DRAM
STACKID 1 system uptime is 14 minute(s) 30 second(s)
STACKID 2 system uptime is 14 minute(s) 6 second(s)
The system: started=warm start reloaded=by "reload"
```

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 Layer 2 Switches and Layer 3 Switches and by configuration level.

To determine the available show commands for the system or a specific level of the CLI, enter the following command.

```
device#show ?
```

Syntax: show option

You also can enter "show" at the command prompt, then press the TAB key.

Enabling the display of the elapsed timestamp for port statistics reset

Whenever the port statistics of a device are cleared globally or on an interface, the counter values of the received and transmitted packets on the device are reset for all the ports or for an interface, respectively.

The elapsed time after the most recent reset of the port statistics counters can be displayed in the output of the **show statistics** command by configuring the **port-statistics-reset-timestamp enable** command. By default, the display of the elapsed timestamp information is disabled.

The elapsed time is calculated as the time between the most recent reset of the port statistics counters and the time when the **show statistics** command is executed.

The following list provides details of the conditions under which the port statistics counters are reset and also explains the elapsed time calculation considerations.

- When the port statistics are cleared individually using the **clear statistics ethernet** command. The elapsed time is calculated and displayed only for that particular interface.
- When the port statistics are cleared globally using the **clear statistics** command. The port statistics counters for all the ports, including management ports, are cleared and the elapsed time is calculated and displayed for each of the interfaces.
- When the management interface is cleared using the **clear statistics management** command. The port statistics counters specific to management ports are cleared. The elapsed time is calculated and displayed for the management interface.
- If the system is reloaded (hard reboot or soft reboot), the port statistics on the device are cleared automatically. In this case, the time when the ports are cleared during the reload is considered as the most recent reset time.
- In a stacking device, the Elapsed Timestamp information is applicable for other unit's ports. In case of a switchover, all the port statistics are cleared and the elapsed time is calculated and displayed for all ports.
- If hitless failover is enabled and if any unit is reloaded, the statistics of the reloading device's interfaces are cleared. In this case, the time when the ports are cleared during the reload is considered as the most recent reset time.
- The elapsed time is not impacted when the Network Time Protocol (NTP) syncs up with a different time other than the recorded time.

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**
- **show statistics**

The Elapsed Timestamp information is displayed in the output of the following **show** commands:

- **show statistics**
- **show statistics brief**
- **show statistics ethernet**
- **show statistics management**

NOTE

The **port-statistics-reset-timestamp enable** command must be configured to have the Elapsed Timestamp information displayed in the output.

To display the statistics, enter a command such as the following.

```
device# show statistics ethernet 1/1/13
Port      Link      State Dupl Speed Trunk Tag Pvid Pri   MAC           Name
1/1/13   Up        Forward Full 1G   None No  1     0   748e.f893.065c

Port 1/1/13 Counters:
*Last time counter reset (Elapsed Timestamp): 1 hour(s) 21 minute(s) 12 second(s)
InOctets      50218819740      OutOctets      50216689676
InPkts        63180119         OutPkts        63428168
InBroadcastPkts 5      OutBroadcastPkts 3
InMulticastPkts 63180114      OutMulticastPkts 63428165
InUnicastPkts           OutUnicastPkts
InBadPkts
InFragments
InDiscards           OutErrors
CRC                  Collisions
InErrors             LateCollisions
InGiantPkts         0
InShortPkts
InJabber
InFlowCtrlPkts      OutFlowCtrlPkts
InBitsPerSec      97441855      OutBitsPerSec      97432612
InPktsPerSec      153280        OutPktsPerSec      153972
InUtilization      100.00%      OutUtilization      100.00%
```

Syntax: `show statistics [ethernet | port]`

TABLE 38 Port statistics in the show statistics command

Parameter	Description
Port configuration	
Port	The port number.
Link	The link state.
State	The STP state.
Dupl	The mode (full-duplex or half-duplex).
Speed	The port speed.
Trunk	The trunk group number, if the port is a member of a trunk group.
Tag	Whether the port is a tagged member of a VLAN.
Pvid	The port default VLAN ID of the port.
Priori	The QoS forwarding priority of the port (level0 - level7).
MAC	The MAC address of the port.
Name	The name of the port, if you assigned a name.
Statistics	
*Last time counter reset (Elapsed Timestamp)	The elapsed time between the most recent reset of the port statistics counters and the time when the show statistics command is executed.
InOctets	The total number of good octets and bad octets received.
OutOctets	The total number of good octets and bad octets sent.
InPkts	The total number of packets received. The count includes rejected and local packets that are not sent to the switching core for transmission.
OutPkts	The total number of good packets sent. The count includes unicast, multicast, and broadcast packets.
InBroadcastPkts	The total number of good broadcast packets received.
OutBroadcastPkts	The total number of good broadcast packets sent.
InMulticastPkts	The total number of good multicast packets received.
OutMulticastPkts	The total number of good multicast packets sent.

TABLE 38 Port statistics in the show statistics command (continued)

Parameter	Description
InUnicastPkts	The total number of good unicast packets received.
OutUnicastPkts	The total number of good unicast packets sent.
InBadPkts	The total number of packets received for which one of the following is true: <ul style="list-style-type: none"> • The CRC is invalid. • The packet is oversized. • Jabbers: The packets are longer than 1518 octets and have a bad FCS. • Fragments: The packets was less than 64 octets long and have a bad FCS. • The packet is undersized (short).
InFragments	The total number of packets received for which both of the following is true: <ul style="list-style-type: none"> • The length is less than 64 bytes. • The CRC is invalid.
InDiscards	The total number of packets that were received and then dropped due to a lack of receive buffers.
OutErrors	The total number of packets with internal transmit errors such as TX underruns.
CRC	The total number of packets received for which all of the following is true: <ul style="list-style-type: none"> • The data length is between 64 bytes and the maximum allowable frame size. • No Collision or Late Collision is detected. • The CRC is invalid.
Collisions	The total number of packets received in which a Collision event was detected.
InErrors	The total number of packets received that had Alignment errors or PHY errors. <p>NOTE Excessive errors for some counters usually indicate a problem. When you operate at a half-duplex setting, some data link errors incrementing in Frame Check Sequence (FCS), alignment, runts, and collision counters are normal. Generally, a one percent ratio of errors to total traffic is acceptable for half-duplex connections. If the ratio of errors to input packets is greater than two or three percent, performance degradation could be noticed.</p> <p>In half-duplex environments, it is possible for both the switch and the connected device to sense the wire and transmit at exactly the same time and result in a collision. Collisions may cause runts, errors in FCS, and alignment errors due to the frame not being completely copied to the wire, resulting in fragmented frames. When you operate at full-duplex, errors in FCS, Cyclic Redundancy Check (CRC), alignment, and runt counters must be minimal.</p>
LateCollisions	The total number of packets received in which a Collision event was detected, but for which a receive error (Rx Error) event was not detected.
InGiantPkts	The total number of packets for which all of the following is true: <ul style="list-style-type: none"> • The data length is longer than the maximum allowable frame size. • No Rx Error is detected.

TABLE 38 Port statistics in the show statistics command (continued)

Parameter	Description
	<p>NOTE Packets are counted for this statistic regardless of whether the CRC is valid or invalid.</p>
InShortPkts	<p>The total number of packets received for which all of the following is true:</p> <ul style="list-style-type: none"> • The data length is less than 64 bytes. • No Rx Error is detected. • No Collision or Late Collision is detected. <p>NOTE Packets are counted for this statistic regardless of whether the CRC is valid or invalid.</p>
InJabber	<p>The total number of packets received for which all of the following is true:</p> <ul style="list-style-type: none"> • The data length is longer than the maximum allowable frame size. • No Rx Error is detected. • The CRC is invalid.
InFlowCtrlPkts	The total number of flow control packets received.
OutFlowCtrlPkts	The total number of flow control packets transmitted.
InBitsPerSec	The number of bits received per second.
OutBitsPerSec	The number of bits sent per second.
InPktsPerSec	The number of packets received per second.
OutPktsPerSec	The number of packets sent per second.
InUtilization	The percentage of the port bandwidth used by received traffic.
OutUtilization	The percentage of the port bandwidth used by sent traffic.

Viewing STP statistics

You can view a summary of STP statistics for Layer 2 Switches and Layer 3 Switches. 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 using the **clear** command.

To determine the available **clear** commands for the system, enter the **clear** command at the Privileged EXEC level of the CLI.

```
device#clear ?
```

Syntax: **clear** *option*

You also can enter "clear" at the command prompt, then press the TAB key.

Traffic counters for outbound traffic

You can configure traffic counters (also called transmit counters) that enable the Brocade device to count the following packet types on a port or port region:

- broadcast packets
- multicast packets
- unicast packets
- dropped packets due to congestion and egress filtering

Depending on the parameters specified with the traffic counter configuration, traffic counters record the number of outbound packets from any combination of the following sources:

- a specific port or all ports in a specific port region
- a specific VLAN or all VLANs
- a specific 802.1p priority queue or all priority queues

Traffic counters configuration notes

Consider the following rules when configuring traffic counters for outbound traffic.

- This feature is supported on FastIron X Series devices only.
- This feature is supported in the Layer 2 and Layer 3 codes.
- This feature applies to physical ports only, including 10 Gbps Ethernet ports and trunk ports. It does not apply to virtual interfaces.
- Once the enhanced traffic counters are read using the **show transmit-counter values** command, the counters are cleared (reset to zero).
- For each port region, you can enable a maximum of two traffic counters, regardless of whether traffic counters are enabled on individual ports or on all ports in the port region.
- Traffic counters increase for bridged filtered outbound traffic when any of the following conditions occur:
 - The port is disabled or the link is down.
 - The port or port region does not belong to the VLAN specified in the transmit counter configuration.
 - A Layer 2 protocol (e.g., spanning tree) has the port in a Blocked state.
 - The source port needs to be suppressed for multi-target packets.
 - The priority queue specified in the traffic counter is not allowed for some other reason.
 - Unknown unicast and unregistered multicast packets are filtered.

Traffic counters configuration syntax

This section provides the syntax and configuration examples for enhanced traffic counters.

To configure traffic counters for outbound traffic on a specific port, enter a command such as the following.

```
device(config)#transmit-counter 4 port 18 only vlan 1 prio 7 enable
```

The above command creates and enables traffic counter 4 on port 18. The device will count the number of packets sent out on port 18 that are in VLAN 1 and have a priority queue of 7.

To configure traffic counters for outbound traffic in a specific port region, enter a command such as the following.

```
device(config)#transmit-counter 1 port 1 region vlan all prio all enable
```

The above command creates and enables traffic counter 1 on all ports that are in the same port region as port 1. The device will count the number of packets transmitted in this port region that belong to any VLAN and have any assigned priority queue.

Syntax: `[no] transmit-counter counter-ID port [slotnum /] port-num { only | region } vlan {vlan-ID | all } priority {priority-queue | all} enable`

Enter the **no** form of the command to remove the outbound traffic counter.

The *counter-ID* parameter identifies the traffic counter. You can configure up to 64 traffic counters. Enter a number from 1 - 64.

The *slotnum* parameter is required on chassis devices.

The *port-num* parameter is the port number to which enhanced traffic counters will apply. Enter the port number followed by **only** to apply the enhanced traffic counter to a specific port, or enter the port number followed by **region** to apply the enhanced traffic counter to all of the ports in the port region.

The *vlan-ID* parameter identifies the VLAN ID for which outbound traffic will be counted. Enter a number from 0 - 4095 or enter **all** to indicate all VLANs.

The *priority-queue* parameter identifies the 802.1p priority queue for which traffic will be counted. Enter a number from 0 - 7 or enter **all** to indicate all priority queues.

Displaying enhanced traffic counter profiles

To display the details of the traffic counters configured on your device, enter the **show transmit-counter profiles** command. The following shows an example output.

```
device#show transmit-counter profiles
Tx Counter      Port(s)  Vlan Id  Priority  Device  Set
      1         1 - 12      All      All     Dev 0   Set0
      4         4 - 18      1        7     Dev 1   Set0
     10        10 - 24     100      All     Dev 1   Set1
```

Displaying enhanced traffic counter statistics

To display the traffic counters for outbound traffic, enter the **show transmit-counter profiles** command.

NOTE

Once the enhanced traffic counters are displayed, the counters are cleared (reset to zero).

The following shows an example output.

```
device#show transmit-counter values 1
Transmit Queue Counter Values for Counter 1:
Transmitted Frames:
  Known Unicast           : 17204
  Multicast & Unknown Unicast : 2797
  Broadcast                : 5
Dropped Frames:
  Bridge Egress Filtered   : 2
  Congestion Drops         : 0
device#show transmit-counter values 4
Transmit Queue Counter Values for Counter 4:
Transmitted Frames:
  Known Unicast           : 124
  Multicast & Unknown Unicast : 2752
  Broadcast                : 0
Dropped Frames:
  Bridge Egress Filtered   : 37
  Congestion Drops         : 0
```

Syntax: `show transmit-counter values number`

where *number* identifies a valid enhanced traffic counter and is a value from 1 - 64.

TABLE 39 Outbound traffic counter statistics

This line...	Displays...
Transmitted frames	
Known Unicast	The number of known unicast packets transmitted.
Multicast & Unknown Unicast	The number of multicast and unknown unicast packets transmitted.
Broadcast	The number of broadcast packets transmitted.
Dropped Frames	
Bridge Egress Filtered	The number of bridged outbound packets that were filtered and dropped. This number includes the number of packets that were dropped because of any one of the following conditions: <ul style="list-style-type: none"> • The port was disabled or the link was down. • The port or port region does not belong to the VLAN specified in the transmit counter configuration. • A Layer 2 protocol (e.g., spanning tree) had the port in a Blocked state. • The source port was suppressed for multi-target packets. • The priority queue specified in the traffic counter was not allowed for some other reason. • Unknown unicast and unregistered multicast packets were filtered.
Congestion Drops	The number of outbound packets that were dropped because of traffic congestion.

Viewing egress queue counters on ICX 6610 and FCX devices

The **show interface** command displays the number of packets on a port that were queued for each QoS priority (traffic class) and dropped because of congestion.

NOTE

These counters do not include traffic on management ports or for a stack member unit that is down.

The egress queue counters display at the end of the **show interface** command output as shown in the following example.

```
device#show interface e 1/1/1
GigabitEthernet1/1/1 is up, line protocol is up
Hardware is GigabitEthernet, address is 0000.0077.8080 (bia 0000.0077.8080)
Configured speed auto, actual 1Gbit, configured duplex fdx, actual fdx
Configured mdi mode AUTO, actual none
Member of L2 VLAN ID 52, port is untagged, port state is FORWARDING
BPDU guard is Disabled, ROOT protect is Disabled
Link Error Dampening is Disabled
STP configured to ON, priority is level0, mac-learning is enabled
Flow Control is config enabled, oper enabled, negotiation disabled
mirror disabled, monitor disabled
Not member of any active trunks
Not member of any configured trunks
No port name
Inter-Packet Gap (IPG) is 96 bit times
IP MTU 1500 bytes
300 second input rate: 0 bits/sec, 0 packets/sec, 0.00% utilization
300 second output rate: 256 bits/sec, 0 packets/sec, 0.00% utilization
0 packets input, 0 bytes, 0 no buffer
Received 0 broadcasts, 0 multicasts, 0 unicasts
0 input errors, 0 CRC, 0 frame, 0 ignored
0 runts, 0 giants
```

```

215704 packets output, 13805066 bytes, 0 underruns
Transmitted 0 broadcasts, 215704 multicasts, 0 unicasts
0 output errors, 0 collisions
Relay Agent Information option: Disabled
Egress queues:
Queue counters    Queued packets    Dropped Packets
0                 0                 0
1                 0                 0
2                 1                 0
3                 0                 0
4                 0                 0
5                 0                 0
6                 0                 0
7                 215703           0

```

Syntax: `show interface [ethernet port]`

Specify the *port* variable in the format *stack-unit/slotnum/portnum*.

TABLE 40 Egress queue statistics

Parameter	Description
Queue counters	The QoS traffic class.
Queued packets	The number of packets queued on the port for the given traffic class.
Dropped packets	The number of packets for the given traffic class that were dropped because of congestion.

Viewing egress queue counters on ICX 7750 devices

Viewing egress queue counters on ICX 7750 devices.

For a port, the **show interface** command displays the number of packets that were queued for each QoS priority (traffic class) and dropped because of congestion. The egress queue counters are displayed at the end of the **show interface** command output as shown in the following example.

NOTE

This command output displays the total of unicast and multicast counters for any particular QOS priority.

```

Brocade# show interface ethernet 1/1/1
10GigabitEthernet 1/1/1 is down, line protocol is down
Port down for 16 hours 16 minutes 48 seconds
Hardware is 10GigabitEthernet , address is 748e.f8f9.6280 (bia 748e.f8f9.6280)
Interface type is 40Gig Fiber
Configured speed 40Gbit, actual unknown, configured duplex fdx, actual unknown
Configured mdi mode AUTO, actual unknown
Member of L2 VLAN ID 1, port is untagged, port state is BLOCKING
BPDU guard is Disabled, ROOT protect is Disabled, Designated protect is Disabled
Link Error Dampening is Disabled
STP configured to ON, priority is level0, mac-learning is enabled
Flow Control is enabled
Mirror disabled, Monitor disabled
Mac-notification is disabled
Not member of any active trunks
Not member of any configured trunks
No port name
IPG MII 96 bits-time, IPG GMII 96 bits-time
MTU 1500 bytes
300 second input rate: 0 bits/sec, 0 packets/sec, 0.00% utilization
300 second output rate: 0 bits/sec, 0 packets/sec, 0.00% utilization
0 packets input, 0 bytes, 0 no buffer
Received 0 broadcasts, 0 multicasts, 0 unicasts
0 input errors, 0 CRC, 0 frame, 0 ignored
0 runts, 0 giants
0 packets output, 0 bytes, 0 underruns

```

```
Transmitted 0 broadcasts, 0 multicasts, 0 unicasts
0 output errors, 0 collisions
Relay Agent Information option: Disabled
```

```
Egress queues:
Queue counters   Queued packets   Dropped Packets
0                 0                 0
1                 0                 0
2                 0                 0
3                 0                 0
4                 0                 0
5                 0                 0
6                 0                 0
7                 0                 0
```

Clearing the egress queue counters

You can clear egress queue statistics (reset them to zero), using the **clear statistics** and **clear statistics ethernet** port command.

Syntax: **clear statistics** [**ethernet port**]

Specify the *port* variable in the format *stack-unit/slotnum/portnum*.

RMON support

The Brocade RMON agent supports the following groups. The group numbers come from the RMON specification (RFC 1757):

NOTE

RFC 1757 is obsolete and is replaced by RFC 2819 for the Brocade ICX devices.

- 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.

Maximum number of entries allowed in the RMON control table

You can specify the maximum number of entries allowed in the RMON control table, including alarms, history, and events. The default number of RMON entries allowed in the RMON control table is 2048 on the FSX 800 and FSX 1600. The maximum number of RMON entries supported is 32768.

To set the maximum number of allowable entries to 3000 in the RMON history table, enter commands such as the following.

```
device(config)#system-max rmon-entries 3000
device(config)#write mem
device(config)#exit
device#reload
```

NOTE

You must save the change to the startup-config file and reload or reboot. The change does not take effect until you reload or reboot.

Syntax: **system-max rmon-entries** *value*

where *value* can be:

- 1536 - 32768 for FSX 800 and FSX 1600 devices

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 Brocade Layer 2 Switch or Layer 3 Switch.

The statistics group collects statistics on promiscuous traffic across an interface. The interface group collects statistics on total traffic into and out of the agent interface.

No configuration is required to activate collection of statistics for the Layer 2 Switch or Layer 3 Switch. This activity is by default automatically activated at system start-up.

You can view a textual summary of the statistics for all ports by entering the following CLI command.

```
device#show rmon statistics
Ethernet statistics 1 is active, owned by monitor
Interface 1/1/1 (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 [ethernet port]`

NOTE

Though 48GC modules receive oversized packets and jabbers, they do not support count information for oversized packets and jabbers and the output of the **show rmon statistics** command reports 0 for both of these counters.

The *port* 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 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 1/2/1.

This command shows the following information.

TABLE 41 Export configuration and statistics

Parameter	Definition
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.

TABLE 41 Export configuration and statistics (continued)

Parameter	Definition
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. NOTE 48GC modules do not support count information on oversized packets and report 0.
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. NOTE 48GC modules do not support count information on jabbers and report 0.
Collisions	The best estimate of the total number of collisions on this Ethernet segment.
64 octets pkts	The total number of packets received that were 64 octets long. This number includes bad packets. This number does not include framing bits but does include FCS octets.
65 to 127 octets pkts	The total number of packets received that were 65 - 127 octets long.

TABLE 41 Export configuration and statistics (continued)

Parameter	Definition
	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.

History (RMON group 2)

All active ports by default will generate two history control data entries per active Brocade Layer 2 Switch port or Layer 3 Switch interface. An active port is defined as one with a link up. If the link goes down the two entries are automatically deleted.

Two history entries are generated for each device:

- A sampling of statistics every 30 seconds
- A sampling of statistics every 30 minutes

The history data can be accessed and displayed using any of the popular RMON applications

A sample RMON history command and its syntax is shown below.

```
device(config)#rmon history 1 interface 1 buckets 10 interval 10 owner nyc02
```

Syntax: `rmon history`*entry-number interface port buckets number interval sampling-interval owner text-string*

You can modify the sampling interval and the bucket (number of entries saved before overwrite) using the CLI. In the above example, owner refers to the RMON station that will request the information.

NOTE

To review the control data entry for each port or interface, enter the **show rmon history** command.

Alarm (RMON group 3)

Alarm is designed to monitor configured thresholds for any SNMP integer, time tick, gauge or counter MIB object. Using the CLI, you can define what MIB objects are monitored, the type of thresholds that are monitored (falling, rising or both), the value of those thresholds, and the sample type (absolute or delta).

An alarm event is reported each time that a threshold is exceeded. The alarm entry also indicates the action (event) to be taken if the threshold be exceeded.

A sample CLI alarm entry and its syntax is shown below.

```
device(config)#rmon alarm 1 ifInOctets.6 10 delta rising-threshold 100 1 falling threshold 50 1 owner nyc02
```

Syntax: `rmon alarm entry-number MIB-object . interface numsampling timesample type-threshold type-threshold value event number -threshold type-threshold valueevent-number owner text-string`

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.

```
device(config)# rmon event 1 description 'testing a longer string' trap public owner nyc02
```

Syntax: `rmon eventevent-entry description text-string {log | trap | log-and-trap} owner rmon-station`

NOTE

FastIron devices currently support only the **trap** option.

sFlow

NOTE

FastIron devices support sFlow version 5 by default.

sFlow is a standards-based protocol that allows network traffic to be sampled at a user-defined rate for the purpose of monitoring traffic flow patterns and identifying packet transfer rates on user-specified interfaces.

When sFlow is enabled on a Layer 2 or Layer 3 switch, the system performs the following sFlow-related tasks:

- Samples traffic flows by copying packet header information
- Identifies ingress and egress interfaces for the sampled flows
- Combines sFlow samples into UDP packets and forwards them to the sFlow collectors for analysis
- Forwards byte and packet count data, or counter samples, to sFlow collectors

sFlow is described in RFC 3176, "InMon Corporation's sFlow: A Method for Monitoring Traffic in Switched and Routed Networks".

On ICX and FCX Series devices, you can use QoS queue 1 for priority traffic, even when sFlow is enabled on the port. This differs from FastIron X Series devices, which support seven priorities instead of eight when sFlow is enabled. In this case, QoS queue 1 is reserved for sFlow and is not used by other packets. Any non-sFlow packets assigned to QoS queue 1 will be directed to QoS queue 0.

sFlow version 5

sFlow version 5 enhances and modifies the format of the data sent to the sFlow collector. sFlow version 5 introduces several new sFlow features and also defines a new datagram syntax used by the sFlow agent to report flow samples and interface counters to the sFlow collector.

sFlow version 5 adds support for the following:

- sFlow version 5 datagrams
- Sub-agent support

- Configurable sFlow export packet size
- Support for the new data field and sample type length in flow samples
- Configurable interval for exporting Brocade-specific data structure

sFlow version 5 is backward-compatible with sFlow version 2. By default, the sFlow agent exports sFlow version 5 flow samples by default, but you can configure the device to export the data in sFlow version 2 format. You can switch between sFlow version 2 and sFlow version 5 formats. The sFlow collector automatically parses each incoming sample and decodes it based on the version number.

The configuration procedures for sFlow version 5 are the same as for sFlow version 2, except where explicitly noted. Configuration procedures for sFlow are in the section [Configuring and enabling sFlow](#) on page 239. The features and CLI commands that are specific to sFlow version 5 are described in the section [sFlow version 5 feature configuration](#) on page 245.

sFlow support for IPv6 packets

The Brocade implementation of sFlow features support IPv6 packets. This support includes extended router information and extended gateway information in the sampled packet. Note that sFlow support for IPv6 packets exists only on devices running software that supports IPv6.

The configuration procedures for this feature are the same as for IPv4, except where the collector is a link-local address on a Layer 3 switch. For details refer to [Specifying the collector](#) on page 240.

Extended router information

IPv6 sFlow sampled packets include the following extended router information:

- IP address of the next hop router
- Outgoing VLAN ID
- Source IP address prefix length
- Destination IP address prefix length

Note that in IPv6 devices, the prefix lengths of the source and destination IP addresses are collected if BGP is configured and the route lookup is completed. In IPv4 devices, this information is collected only if BGP is configured on the devices.

Extended gateway information

If BGP is enabled, extended gateway information is included in IPv6 sFlow sampled packets, including the following BGP information about a packet destination route:

- The Autonomous System number for the router
- The source IP Autonomous System of the route
- The source peer Autonomous System for the route
- The Autonomous System patch to the destination

NOTE

Autonomous System communities and local preferences are not included in the sampled packets.

To obtain extended gateway information, use "struct extended_gateway" as described in RFC 3176.

IPv6 packet sampling

IPv6 sampling is performed by the packet processor. The system uses the sampling rate setting to selectively mark the monitoring bit in the header of an incoming packet. Marked packets tell the CPU that the packets are subject to sFlow sampling.

sFlow configuration considerations

This section lists the sFlow configuration considerations on Brocade devices.

On ICX and FCX Series devices, you can use QoS queue 1 for priority traffic, even when sFlow is enabled on the port. This differs from FastIron X Series devices, which support seven priorities instead of eight when sFlow is enabled. In this case, QoS queue 1 is reserved for sFlow and is not used by other packets. Any non-sFlow packets assigned to QoS queue 1 will be directed to QoS queue 0.

If ICX and FCX stacks are rebooted, sFlow is disabled on standby and member units until the configuration is synchronized between the Active and Standby Controllers.

sFlow and hardware support

- Brocade devices support sFlow packet sampling of inbound traffic only. These devices do not sample outbound packets. However, Brocade devices support byte and packet count statistics for both traffic directions.
- sFlow is supported on all Ethernet ports (10/100, Gbps, and 10 Gbps)

sFlow and CPU utilization

Enabling sFlow may cause a slight and noticeable increase of up to 20% in CPU utilization. In typical scenarios, this is normal behavior for sFlow, and does not affect the functionality of other features on the switch.

sFlow and agent address

The sampled sFlow data sent to the collectors includes an `agent_address` field. This field identifies the IP address of the device that sent the data:

- On a Layer 2 switch, `agent_address` is the Layer 2 switch management IP address. You must configure the management IP address in order to export sFlow data from the device. If the switch has both an IPv4 and IPv6 address, the `agent_address` is the IPv4 address. If the switch has an IPv6 address only, the `agent_address` is the global IPv6 address.
- On a Layer 3 switch with IPv6 interfaces only, sFlow looks for an IPv6 address in the following order, and uses the first address found:
 - The first IPv6 address on the lowest-numbered loopback interface
 - The first IPv6 address on the lowest-numbered VE interface
 - The first IPv6 address on any interface
- On a Layer 3 switch with both IPv4 and IPv6 interfaces, or with IPv4 interfaces only, sFlow looks for an IP address in the following order, and uses the first address found:
 - The IPv4 router ID configured by the `ip router-id` command
 - The first IPv4 address on the lowest-numbered loopback interface
 - The first IPv4 address on the lowest-numbered virtual interface
 - The first IPv4 address on any interface

NOTE

The device uses the router ID only if the device also has an IP interface with the same address. Router ID is not supported on IPv6 devices.

NOTE

If an IP address is not already configured when you enable sFlow, the feature uses the source address 0.0.0.0. To display the agent_address, enable sFlow, then enter the **show sflow** command. Refer to [Enabling sFlow forwarding](#) on page 243 and [Displaying sFlow information](#) on page 249.

NOTE

In sFlow version 5, you can set an arbitrary IPv4 or IPv6 address as the sFlow agent IP address. Refer to [Specifying the sFlow agent IP address](#) on page 245.

sFlow and source IP address

When the sFlow packet is sent to the sFlow collector, by default, the IP address of the outgoing interface is used in the sFlow datagram.

However, you can specify the source interface, from which the IP address is selected for the sFlow datagram, using the **sflow source** command. The Ethernet, VE, or loopback interface can be configured as the source interface for both IPv4 and IPv6 addresses.

sFlow source IP address configuration notes

- The first IP address in the interface IP address list is considered the source IP address.
- If the sFlow destination is IPv6, and the sFlow source is configured for an IPv6 address, then an IPv6 address will be selected from the configured interface.
- If the sFlow destination is IPv4, and the sFlow source is configured for IPv4 address, then an IPv4 address will be selected from the configured interface.
- At any point of time, only one source of the Ethernet, VE, or loopback interface can be specified as the source interface.
- Upon configuring another source for an IPv4 or IPv6 address, any previously configured source for the IPv4 or IPv6 address will be deleted.
- If the source IP address is not configured, by default, the IP address of the outgoing interface will be used in the sFlow datagram.
- You can configure IPv4 and IPv6 source interfaces independently.
- In case of LAG, the sFlow source configuration is valid only for the primary port.
- The sFlow source IP configuration is supported on sFlow version 2 and sFlow version 5 and is valid only for the router build.
- Addition and deletion of IPv4 and IPv6 addresses on an sFlow source interface will trigger the following events:
 - If the added IP address is the first IP address in the table, then it will be considered as the source IP address.
 - If the added IP address is positioned on top of the IP table (due to IP address sequence order), then it will be reassigned as the source IP address.
 - If the IP address that is used as the source IP is deleted, the next IP address on the same interface will be considered as the source IP address.
 - If all the IP addresses are deleted from the source interface, the IP address of the outgoing interface is used in the sFlow datagram.

sFlow and source port

By default, sFlow sends data to the collector out of UDP source port 8888, but you can specify a different source port. For more information, refer to [Changing the sFlow source port](#) on page 243.

sFlow and sampling rate

The *sampling rate* is the average ratio of the number of packets incoming on an sFlow enabled port, to the number of flow samples taken from those packets. sFlow sampling can affect performance in some configurations.

Note that on the FastIron devices, the configured sampling rate and the actual rate are the same. The software does not adjust the configured sampling rate as on other Brocade devices.

NOTE

The value range for sampling rate is from 256 through 1073741823 on Brocade ICX 7750, ICX 7450, and ICX 7250 devices. The default value is 4096 for all devices.

sFlow and port monitoring

- ICX and FCX Series devices support sFlow and port monitoring together on the same port.
- FastIron X Series devices support port monitoring and sFlow together on the same device. The caveat is that these features cannot be configured together within the same port region on non-third-generation modules. The following third-generation SX modules support sFlow and mirroring on the same port:
 - SX-FI48GPP
 - SX-FI-24GPP
 - SX-FI-24HF
 - SX-FI-2XG
 - SX-FI-8XG

Configuring and enabling sFlow

NOTE

The commands in this section apply to sFlow version 2 and sFlow version 5. CLI commands that are specific to sFlow version 5 are documented in [sFlow version 5 feature configuration](#) on page 245.

To configure sFlow, perform the following tasks:

- Optional - If your device supports sFlow version 5, change the version used for exporting sFlow data
- Specify collector information. The collector is the external device to which you are exporting the sFlow data. You can specify up to four collectors.
- Optional - Change the polling interval
- Optional - Change the sampling mode to include dropped packets
- Optional - Change the sampling rate
- Optional - Change the sFlow source IP address
- Optional - Change the sFlow source port
- Enable sFlow globally
- Enable sFlow forwarding on individual interfaces
- Enable sFlow forwarding on individual trunk ports
- If your device supports sFlow version 5, configure sFlow version 5 features

Specifying the collector

sFlow exports traffic statistics to an external collector. You can specify up to four collectors. You can specify more than one collector with the same IP address if the UDP port numbers are unique. You can have up to four unique combinations of IP addresses and UDP port numbers.

Specifying an sFlow collector on IPv4 devices

To specify an sFlow collector on an IPv4 device, enter a command such as the following.

```
device(config)#sflow destination 10.10.10.1
```

This command specifies a collector with IPv4 address 10.10.10.1, listening for sFlow data on UDP port 6343.

Syntax: `[no] sflow destination ip-addr [dest-udp-port | vrf]`

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

The *dest-udp-port* parameter specifies the UDP port on which the sFlow collector will be listening for exported sFlow data. The default port number is 6343. For information on VRF parameter, see the *FastIron Layer 3 Routing Configuration Guide*.

The sampled sFlow data sent to the collectors includes an *agent_address* field. This field identifies the device that sent the data. Refer to [sFlow and agent address](#) on page 237.

Specifying an sFlow collector on IPv6 devices

To specify an sFlow collector on an IPv6 device, enter a command such as the following.

```
device(config)#sflow destination ipv6 2001:DB8:0::0b:02a
```

This command specifies a collector with IPv6 address 2001:DB8:0b:02a, listening for sFlow data on UDP port 6343.

Syntax: `[no] sflow destination ipv6 ip-addr [dest-udp-port]`

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

The *dest-udp-port* parameter specifies the UDP port on which the sFlow collector will be listening for exported sFlow data. The default port number is 6343.

If the IPv6 address you specify is a link-local address on a Layer 3 switch, you must also specify the *outgoing-interface ethernet port-num* or the *ve port-num*. This identifies the outgoing interface through which the sampled packets will be sent.

The sampled sFlow data sent to the collectors includes an *agent_address* field. This field identifies the device that sent the data. Refer to [sFlow and agent address](#) on page 237.

Changing the polling interval

The polling interval defines how often sFlow byte and packet counter data for a port are sent to the sFlow collectors. If multiple ports are enabled for sFlow, the Brocade device staggers transmission of the counter data to smooth performance. For example, if sFlow is enabled on two ports and the polling interval is 20 seconds, the Brocade device sends counter data every ten seconds. The counter data for one of the ports are sent after ten seconds, and counter data for the other port are sent after an additional ten seconds. Ten seconds later, new counter data for the first port are sent. Similarly, if sFlow is enabled on five ports and the polling interval is 20 seconds, the Brocade device sends counter data every four seconds.

The default polling interval is 20 seconds. You can change the interval to a value from 0 to 4294967295 seconds. The interval value applies to all interfaces on which sFlow is enabled. If you set the polling interval to 0, counter data sampling is disabled.

To change the polling interval, enter a command such as the following at the global CONFIG level of the CLI.

```
device(config)#sflow polling-interval 30
```

Syntax: [no] sflow polling-interval secs

The secs parameter specifies the interval and can be from 0 through 4294967295 seconds. The default is 20 seconds. If you specify 0, counter data sampling is disabled.

Changing the sampling mode

On all devices other than the Brocade ICX 7750, Brocade ICX 7450, and Brocade ICX 7250, by default, only the non-dropped packets are included for sFlow sampling.

The dropped packets are not included in the sFlow samples that are sent to the sFlow collector. The sampling mode can be changed to include the dropped packets using the **sflow sample-mode** command.

NOTE

The **sflow sample-mode** command is not supported on Brocade ICX 7750, Brocade ICX 7450, Brocade ICX 7250, Brocade ICX 6430, Brocade ICX 6650, and FSX 800/1600 devices.

Changing the sampling rate

The sampling rate is the average ratio of the number of packets incoming on an sFlow-enabled port, to the number of flow samples taken from those packets.

You can change the default (global) sampling rate. You also can change the rate on an individual port, overriding the default sampling rate of 4096. With a sampling rate of 4096, on average, one in every 4096 packets forwarded on an interface is sampled.

Configuration considerations

The sampling rate is a fraction in the form $1/N$, meaning that, on average, one out of every N packets is sampled. The **sflow sample** command from the global configuration mode or port mode specifies N , the denominator of the fraction. Thus a higher number for the denominator means a lower sampling rate since fewer packets are sampled. Likewise, a lower number for the denominator means a higher sampling rate because more packets are sampled. For example, if you change the denominator from 512 to 128, the sampling rate increases because four times as many packets are sampled.

NOTE

Brocade recommends that you do not change the denominator to a value lower than the default. Sampling requires CPU resources. Using a low denominator for the sampling rate can cause high CPU utilization.

On Brocade ICX 7750, ICX 7450, and ICX 7250, the CPU-bound sFlow sample packets are rate-limited to 50 samples per second to avoid high CPU utilization.

If the input traffic rate is more on the interface, the sampling rate must be configured to a higher value to keep the number of sample packets within the CPU rate limit. Else, the excess sample packets are dropped by the CPU.

The following examples show the ideal sample rate configurations for various input rates that keep the sample packets within the CPU rate limit.

- If the input traffic rate is 200,000 packets/sec, the interface sample rate must be set to 4096. ($200000/4096 =$ less than 50 samples)
- If the input traffic rate is 400,000 packets/sec, the interface sample rate must be set to 8192. ($400000/8192 =$ less than 50 samples)

As the sample packets are generated within the CPU rate limit (50 samples/sec) in the above example, the packets are forwarded to the sFlow collector.

Configured rate and actual rate — When you enter a sampling rate value, this value is the configured rate as well as the actual sampling rate.

Change to global rate — If you change the global sampling rate, the change is applied to all sFlow-enabled ports except those ports on which you have already explicitly set the sampling rate. For example, suppose that sFlow is enabled on ports 1/1/1, 1/1/2, and 1/5/1. If you configure the sampling rate on port 1/1/1 but leave the other two ports using the default rate, then a change to the global sampling rate applies to ports 1/1/2 and 1/5/1 but not port 1/1/1. sFlow assumes that you want to continue using the sampling rate you explicitly configured on an individual port even if you globally change the sampling rate for the other ports.

Module rate — While different ports on a module may be configured to have different sampling rates, the hardware for the module is programmed to take samples at a single rate (the module sampling rate). The module sampling rate is the highest sampling rate (that is, the lowest number) configured for any of the ports on the module.

When ports on a given module are configured with different sampling rates, the CPU discards some of the samples supplied by the hardware for ports with configured sampling rates that are lower than the module sampling rate. This is referred to as subsampling, and the ratio between the port sampling rate and the module sampling rate is known as the subsampling factor. For example, if the module in slot 4 has sFlow enabled on ports 1/4/2 and 1/4/8, and port 1/4/2 is using the default sampling rate of 512, and port 1/4/8 is configured explicitly for a rate of 2048, then the module sampling rate is 512 because this is this highest port sampling rate (lowest number). The subsampling factor for port 1/4/2 is 1, meaning that every sample taken by the hardware is exported, while the subsampling factor for port 1/4/8 is 4, meaning that one out of every four samples taken by the hardware is exported. Whether the port sampling rate is configured explicitly, or whether it uses the global default setting, has no effect on the calculations.

You do not need to perform any of these calculations to change a sampling rate. You can display the rates you entered for the default sampling rate, module rates, and all sFlow-enabled ports by entering the **show sflow** command. Refer to [Displaying sFlow information](#) on page 249.

Sampling rate for new ports — When you enable sFlow on a port, the port's sampling rate is set to the global default sampling rate. This also applies to ports on which you disable and then re-enable sFlow. The port does not retain the sampling rate it had when you disabled sFlow on the port, even if you had explicitly set the sampling rate on the port.

Changing the default sampling rate

To change the default (global) sampling rate, enter a command such as the following at the global CONFIG level of the CLI.

```
device(config)# sflow sample 2048
```

Changing the sampling rate of a module

You cannot change a module sampling rate directly. You can change a module sampling rate only by changing the sampling rate of a port on that module.

Changing the sampling rate on a port

You can configure an individual port to use a different sampling rate than the global default sampling rate. This is useful in cases where ports have different bandwidths. For example, if you are using sFlow on 10/100 ports and Gbps Ethernet ports, you might want to configure the Gbps ports to use a higher sampling rate (and thus gather fewer samples per number of packets) than the 10/100 ports.

To change the sampling rate on an individual port, enter a command such as the following from the configuration mode for the port.

```
device(config-if-1/1/1)# sflow sample 8192
```

NOTE

Configuring a sampling rate on a port that is the primary port of a trunk applies that same sampling rate to all ports in the trunk.

Changing the sampling rate for a trunk port

You can configure an individual static trunk port to use a different sampling rate than the global default sampling rate. This feature is also supported on LACP trunk ports. This feature is useful in cases where ports have different bandwidths. For example, if you are using sFlow on 10/100 ports and Gbps Ethernet ports, you might want to configure the Gbps ports to use a higher sampling rate (and thus gather fewer samples per number of packets) than the 10/100 ports.

To configure a static trunk port to use a different sampling rate than the global default sampling rate, enter commands such as the following:

```
device(config)# trunk ethernet 1/4/1 to 1/4/8
device(config-trunk-1/4/1-1/4/8)# sflow sample 8192
```

NOTE

Configuring a sampling rate on only the port that is the primary port of a trunk automatically applies that same sampling rate to all ports in the trunk.

Changing the sFlow source port

By default, sFlow sends data to the collector using UDP source port 8888, but you can change the source UDP port to any port number in the range 1025-65535.

To change the source UDP port, enter a command such as the following:

```
device(config)#sflow source-port 8000
```

Syntax: `[no] sflow source-port num`

The *num* parameter specifies the sFlow source port.

Enabling sFlow forwarding

sFlow exports data only for the interfaces on which you enable sFlow forwarding. You can enable sFlow forwarding on Ethernet interfaces.

NOTE

When management port is used, sFlow can be received only from active units in a stack (not from all units). However, if you use management VLAN with data port, sFlow is received normally. To receive sFlow from all units in a stack, you must use a data port.

To enable sFlow forwarding, perform the following:

- Globally enable the sFlow feature
- Enable sFlow forwarding on individual interfaces
- Enable sFlow forwarding on individual trunk ports

NOTE

Before you enable sFlow, make sure the device has an IP address that sFlow can use as its source address. Refer to [sFlow and agent address](#) on page 237 for the source address requirements.

NOTE

When you enable sFlow forwarding on an 802.1X-enabled interface, the samples taken from the interface include the username used to obtain access to either or both the inbound and outbound ports, if that information is available. For information about 802.1X, refer to "Flexible authentication" chapter in the *FastIron Ethernet Switch Security Configuration Guide*

Command syntax for enabling sFlow forwarding

This section shows how to enable sFlow forwarding.

Globally enabling sFlow forwarding

To enable sFlow forwarding, you must first enable it on a global basis, then on individual interfaces or trunk ports, or both.

To globally enable sFlow forwarding, enter the following command.

```
device(config)#sflow enable
```

You can now enable sFlow forwarding on individual ports as described in the next two sections.

Syntax: [no] sflow enable

Enabling sFlow forwarding on individual interfaces

To enable sFlow forwarding enter commands such as the following.

```
device(config)#sflow enable
device(config)#interface ethernet 1/1/1 to 1/1/8
device(config-mif-1/1/1-1/1/8)#sflow forwarding
```

These commands globally enable sFlow, then enable sFlow forwarding on Ethernet ports 1/1/1 - 1/1/8. You must use both the **sflow enable** and **sflow forwarding** commands to enable the feature.

Syntax: [no] sflow enable

Syntax: [no] sflow forwarding

Enabling sFlow forwarding on individual trunk ports

This feature is supported on individual ports of a static trunk group. It is also supported on LACP trunk ports.

NOTE

When you enable sFlow forwarding on a trunk port, only the primary port of the trunk group forwards sFlow samples.

To enable sFlow forwarding on a trunk port, enter commands such as the following.

```
device(config)#sflow enable
device(config)#trunk e 1/4/1 to 1/4/8
device(config-trunk-1/4/1-1/4/8)#config-trunk-ind
device(config-trunk-1/4/1-1/4/8)#sflow forwarding e 1/4/2
```

These commands globally enable sFlow, then enable sFlow forwarding on trunk port e 1/4/2. You must use both the **sflow enable** and **sflow forwarding** commands to enable the feature.

Syntax: [no] sflow enable

Syntax: [no] sflow forwarding

sFlow version 5 feature configuration

NOTE

The commands in this section are supported when sFlow version 5 is enabled on the device. These commands are not supported with sFlow version 2. sFlow version 5 also supports all of the sFlow configuration commands in [Configuring and enabling sFlow](#) on page 239.

When sFlow version 5 is enabled on the device, you can do the following:

- Specify the sFlow version (version 2 or version 5)
- Specify the sFlow agent IP address
- Specify the maximum flow sample size
- Export CPU and memory usage information to the sFlow collector
- Specify the polling interval for exporting CPU and memory usage information to the sFlow collector
- Export CPU-directed data (management traffic) to the sFlow collector

Egress interface ID for sampled broadcast and multicast packets

For broadcast and multicast traffic, the egress interface ID for sampled traffic is always 0x80000000. When broadcast and multicast packets are sampled, they are usually forwarded to more than one port. However, the output port field in an sFlow datagram supports the display of one egress interface ID only. Therefore, the sFlow version 5 agent always sets the output port ID to 0x80000000 for broadcast and multicast packets that are sampled.

Specifying the sFlow version format

If your device supports sFlow version 5, you can optionally specify the version used for exporting sFlow data. Refer [Specifying the sFlow agent IP address](#) on page 245.

Specifying the sFlow agent IP address

The sampled sFlow data sent to the collectors includes an `agent_address` field. This field identifies the device (the sFlow agent) that sent the data. By default, the device automatically selects the sFlow agent IP address based on the configuration, as described in the section [sFlow and agent address](#) on page 237. Alternatively, you can configure the device to instead use an arbitrary IPv4 or IPv6 address as the sFlow agent IP address.

To specify an IPv4 address as the sFlow agent IP address, enter a command such as the following

```
device(config)#sflow agent-ip 10.10.10.1
```

Syntax: `[no] sflow agent-ipipv4-addr`

The `ipv4-addr` specifies the address of the device that sent the data.

To specify an IPv6 address as the sFlow agent IP address, enter a command such as the following.

```
device(config)#sflow agent-ip FE80::240:D0FF:FE48:4672
```

Syntax: `[no] sflow agent-ipipv6-addr`

The `ipv6-addr` the address of the device that sent the data.

Specifying the version used for exporting sFlow data

By default, when sFlow is enabled globally on the Brocade device, the sFlow agent exports sFlow data in version 5 format. You can change this setting so that the sFlow agent exports data in version 2 format. You can switch between versions without rebooting the device or disabling sFlow.

NOTE

When the sFlow version number is changed, the system will reset sFlow counters and flow sample sequence numbers.

To specify the sFlow version used for exporting sFlow data, enter the following command.

```
device(config)#sflow version 2
```

Syntax: [no] sflow version[2 | 5]

The default is 5.

Specifying the maximum flow sample size

With sFlow version 5, you can specify the maximum size of the flow sample sent to the sFlow collector. If a packet is larger than the specified maximum size, only the data of the packet up to the specified maximum number of bytes is exported. If the size of the packet is smaller than the specified maximum, then the entire packet is exported.

For example, to specify 1024 bytes as the maximum flow sample size, enter the following command.

```
device(config)# sflow max-packet-size 1024
```

Syntax: [no] sflow max-packet-size size

The range of the maximum flow sample size can be from 0 through 1300 bytes. The default value is 128 bytes.

The **sflow max-packet-size** command is supported on all ICX platforms.

The following sample list provides information about the sFlow sample size sent to the sFlow collector, when the max-packet-size is configured with different values.

TABLE 42 sFlow sample size sent to the sFlow collector with varying max-packet-size values

Maximum packet size	Size of the sFlow sample sent to the sFlow collector
0 bytes	Only the information about the packet is captured and no data from the packet is sent to the sFlow collector.
1 byte	1 byte from the packet is sent to the sFlow collector. However, it is padded with zero to make it 4 bytes.
2 bytes	2 bytes from the packet is sent to the sFlow collector. However, it is padded with zero to make it 4 bytes.
100 bytes	100 bytes from packet is sent to the sFlow collector.
200 bytes	200 bytes from packet is sent to the sFlow collector.
1200 bytes	1200 bytes from the packet is sent to the sFlow collector.

Exporting CPU and memory usage information to the sFlow collector

With sFlow version 5, you can optionally configure the sFlow agent on the Brocade device to export information about CPU and memory usage to the sFlow collector.

To export CPU usage and memory usage information, enter the following command.

```
device(config)# sflow export system-info
```

Syntax: [no] sflow export system-info

By default, CPU usage information and memory usage information are not exported.

Specifying the polling interval for exporting CPU and memory usage information to the sFlow collector

The polling interval defines how often sFlow data for a port is sent to the sFlow collector. With sFlow version 5, you can optionally set the polling interval used for exporting CPU and memory usage information.

For example, to set the polling interval for exporting CPU and memory usage information to 30 seconds, enter the following command.

```
device(config)# sflow export system-info 30
```

Syntax: [no] sflow export system-infoseconds

You can specify a polling interval from 5 seconds to 1,800 seconds (30 minutes). The default polling interval for exporting CPU and memory usage information is 300 seconds (5 minutes).

Exporting CPU-directed data (management traffic) to the sFlow collector

You can select which and how often data destined to the CPU (for example, Telnet sessions) is sent to the sFlow collector.

CLI commands allow you to do the following:

- Enable the sFlow agent to export CPU-directed data
- Specify the sampling rate for exported CPU-directed data

Enabling the sFlow agent to export CPU-directed data

To enable the sFlow agent on a Brocade device to export data destined to the CPU to the sFlow collector, enter the following command.

```
device(config)# sflow export cpu-traffic
```

Syntax: [no] sflow export cpu-traffic

By default, this feature is disabled. The sFlow agent does not send data destined to the CPU to the sFlow collector.

Specifying the sampling rate for exported CPU-directed data

The sampling rate is the average ratio of the number of packets incoming on an sFlow-enabled port, to the number of flow samples taken from those packets. You can optionally set the sampling rate for CPU-directed data exported to the sFlow collector. For example, to set this sampling rate to 2048, enter the following command.

```
device(config)# sflow export cpu-traffic 2048
```

Syntax: [no] sflow export cpu-traffic

The default sampling rate depends on the Brocade device being configured. Refer to [Changing the sampling rate](#) on page 241 for the default sampling rate for each kind of Brocade device.

Configuring sFlow with Multi-VRFs

sFlow is a traffic monitoring protocol that supports VRFs. sFlow provides traffic sampling on configured ports based on sample rate and port information to a collector. By default, sFlow uses the management VRF to send the samples to the collector. See the section [Management VRFs](#) on page 19 for information on management VRFs. If no management VRF is configured, sFlow uses the default VRF, and this default VRF ID will be assigned to any configured collector that does not have a user-included VRF.

Collectors can be added and per VRF so that collectors can be spread out across different VRFs. The sFlow forwarding port can belong to a non-default VRF, and captured sFlow packets will have correct sample routing next hop information.

sFlow forwarding ports can come from ports belonging to any VRF. The port does not have to be in the same VRF as the collector. sFlow collects packets from all sFlow forwarding ports, even if they do not belong to a VRF, compiles the packets into the sFlow samples, and sends the samples to the particular collector with no filtering for VRF membership. For counter samples, sample statistics from each port are sent to each collector specified, even if the port and collector do not belong to a VRF instance.

To distinguish collected packets from different VRFs, refer to the **in vlan** and **out vlan** data fields for each captured ingress packet. For example, when two collected packets are from different VRFs but have the same source/destination IP and the same incoming/outgoing port, the VLAN field differs in the two samples. A VLAN/VE can only belong to one VRF. The collector does not have any VRF knowledge, but, based on the VLAN fields, the collector can distinguish which packet came from which VLAN/VRF.

To configure an sFlow collector and specify a VRF, enter the following command.

```
device(config)# sflow destination 10.10.10.vrf customer1
device(config)#
```

Syntax: `[no] sflow destination [ipaddress | ipv6 ipv6-address] [udp-port-number] [vrf vrf-name]`

To disable the management VRF in sFlow, enter the following command.

```
device(config)# sflow management-vrf disable
device(config)#
```

Syntax: `[no] sflow management-vrf-disable`

To display sFlow configuration and statistics, enter the following command.

```
device(config)# show sflow
sFlow version: 5
sFlow services are enabled.
sFlow management VRF is disabled.
sFlow agent IP address: 10.37.230.21
Collector IP 10.37.224.233, UDP 6343, Configured VRF: green
UDP source port: 8888 (Default)
Polling interval is 20 seconds.
Configured default sampling rate: 1 per 500 packets.
Actual default sampling rate: 1 per 500 packets.
The maximum sFlow sample size: 128.
sFlow exporting cpu-traffic is disabled.
100 UDP packets exported
80 sFlow flow samples collected.
sFlow ports: ethe 4/1/5
Module Sampling Rates
-----
Port Sampling Rates
-----
Port=4/1/5, configured rate=500, actual rate=500
```

Syntax: `show sflow`

Displaying sFlow information

To display sFlow configuration information and statistics, enter the following command at any level of the CLI.

```
device#show sflow
sFlow version:5
sFlow services are enabled.
sFlow agent IP address: 10.123.123.1
sFlow source IP address: 5.5.5.5
sFlow source IPv6 address: 4545::2
4 collector destinations configured:
Collector IP 192.168.4.204, UDP 6343
Collector IP 192.168.4.200, UDP 6333
Collector IP 192.168.4.202, UDP 6355
Collector IP 192.168.4.203, UDP 6565
Configured UDP source port: 33333
Polling interval is 0 seconds.
Configured default sampling rate: 1 per 512 packets
Actual default sampling rate: 1 per 512 packets
Sample mode: Non-dropped packets
The maximum sFlow sample size:512
exporting cpu-traffic is enabled
exporting cpu-traffic sample rate:16
exporting system-info is enabled
exporting system-info polling interval:20 seconds
10552 UDP packets exported
24127 sFlow samples collected.
sFlow ports: ethe 1/1/2 to 1/1/12 ethe 1/1/15 ethe 1/1/25 to 1/1/26 ethe 1/4/1 ethe 1/5/10 to
1/5/20 ethe 1/8/1 ethe 1/8/4
Module Sampling Rates
-----
Slot 1 configured rate=512, actual rate=512
Slot 3 configured rate=0, actual rate=0
Slot 4 configured rate=10000, actual rate=32768
Slot 5 configured rate=512, actual rate=512
Slot 7 configured rate=0, actual rate=0
Slot 8 configured rate=512, actual rate=512
Port Sampling Rates
-----
Port 1/8/4, configured rate=512, actual rate=512, Subsampling factor=1
Port 1/8/1, configured rate=512, actual rate=512, Subsampling factor=1
Port 1/5/20, configured rate=3000, actual rate=8192, Subsampling factor=16
Port 1/5/19, configured rate=512, actual rate=512, Subsampling factor=1
Port 1/5/18, configured rate=512, actual rate=512, Subsampling factor=1
Port 1/5/17, configured rate=1500, actual rate=2048, Subsampling factor=4
...Output truncated...
```

Syntax: show sflow

The show sflow command displays the following information.

TABLE 43 sFlow information

Parameter	Definition
sFlow version	The version of sFlow enabled on the device, which can be one of the following: <ul style="list-style-type: none"> • 2 • 5
sFlow services	The feature state, which can be one of the following: <ul style="list-style-type: none"> • disabled • enabled
sFlow agent IP address	The IP address that sFlow is using in the agent_address field of packets sent to the collectors. Refer to sFlow and agent address on page 237.
sFlow source IP address	The IPv4 address that sFlow is using as the source IP address in the sFlow datagram.

TABLE 43 sFlow information (continued)

Parameter	Definition
sFlow source IPv6 address	The IPv6 address that sFlow is using as the source IP address in the sFlow datagram.
Collector	The collector information. The following information is displayed for each collector: <ul style="list-style-type: none"> • IP address • UDP port If more than one collector is configured, the line above the collectors indicates how many have been configured.
Configured UDP source port	The UDP source port used to send data to the collector.
Polling interval	The port counter polling interval.
Configured default sampling rate	The configured global sampling rate. If you changed the global sampling rate, the value you entered is shown here. The actual rate calculated by the software based on the value you entered is listed on the next line, "Actual default sampling rate".
Actual default sampling rate	The actual default sampling rate.
Sample mode	Indicates whether only the non-dropped packets or all the packets, including the dropped packets, are included for sFlow sampling.
The maximum sFlow sample size	The maximum size of a flow sample sent to the sFlow collector.
exporting cpu-traffic	Indicates whether the sFlow agent is configured to export data destined to the CPU (for example, Telnet sessions) to the sFlow collector: <ul style="list-style-type: none"> • enabled • disabled
exporting cpu-traffic sample rate	The sampling rate for CPU-directed data, which is the average ratio of the number of incoming packets on an sFlow-enabled port, to the number of flow samples taken from those packets.
exporting system-info	Indicates whether or not the sFlow agent is configured to export information about CPU and memory usage to the sFlow collector: <ul style="list-style-type: none"> • enabled • disabled
exporting system-info polling interval	Specifies the interval, in seconds, that sFlow data is sent to the sFlow collector.
UDP packets exported	The number of sFlow export packets the Brocade device has sent. <p style="text-align: center;">NOTE Each UDP packet can contain multiple samples.</p>
sFlow samples collected	The number of sampled packets that have been sent to the collectors.
sFlow ports	The ports on which you enabled sFlow.
Module Sampling Rates	The configured and actual sampling rates for each module. If a module does not have any sFlow-enabled ports, the rates are listed as 0.
Port Sampling Rates	The configured and actual sampling rates for each sFlow-enabled port. <p>The Subsampling factor indicates how many times the sampling rate of the port's module is multiplied to achieve the port's sampling rate. Because of the way the actual sampling rates are computed, the Subsampling factors are always whole numbers.</p>

Clearing sFlow statistics

To clear the UDP packet and sFlow sample counters in the **show sflow** display, enter the following command.

```
device#clear statistics
```

Syntax: clear statistics

This command clears the values in the following fields of the **show sflow** display:

- UDP packets exported
- sFlow samples collected

NOTE

This command also clears the statistics counters used by other features.

Utilization list for an uplink port

You can configure uplink utilization lists that display the percentage of a given uplink port bandwidth that is used by a specific list of downlink ports. The percentages are based on 30-second intervals of RMON packet statistics for the ports. Both transmit and receive traffic is counted in each percentage.

NOTE

This feature is intended for ISP or collocation environments in which downlink ports are dedicated to various customers' traffic and are isolated from one another. If traffic regularly passes between the downlink ports, the information displayed by the utilization lists does not provide a clear depiction of traffic exchanged by the downlink ports and the uplink port.

Each uplink utilization list consists of the following:

- Utilization list number (1, 2, 3, or 4)
- One or more uplink ports
- One or more downlink ports

Each list displays the uplink port and the percentage of that port bandwidth that was utilized by the downlink ports over the most recent 30-second interval.

You can configure up to four bandwidth utilization lists.

Utilization list for an uplink port command syntax

To configure an uplink utilization list, enter commands such as the following. The commands in this example configure a link utilization list with port 1/1/1 as the uplink port and ports 1/1/2 and 1/1/3 as the downlink ports.

```
device(config)#relative-utilization 1 uplink eth 1/1/1 downlink eth 1/1/2 to 1/1/3
device(config)#write memory
```

Syntax: **[no] relative-utilization** *num* **uplink ethernet** [**to port** | **port...**] **downlink ethernet** port [**to port** | **[port...**]

The *num* parameter specifies the list number. You can configure up to four lists. Specify a number from 1 - 4.

The **uplink ethernet** parameters and the port numbers you specify after the parameters indicate the uplink ports.

The **downlink ethernet** parameters and the port numbers you specify after the parameters indicate the downlink ports.

Displaying utilization percentages for an uplink

After you configure an uplink utilization list, you can display the list to observe the percentage of the uplink bandwidth that each of the downlink ports used during the most recent 30-second port statistics interval. The number of packets sent and received between the two ports is listed, as well as the ratio of each individual downlink port packets relative to the total number of packets on the uplink.

To display an uplink utilization list, enter a command such as the following at any level of the CLI.

```
device#show relative-utilization 1
uplink: ethe 1
30-sec total uplink packet count = 3011
packet count ratio (%)
  1/ 2:60   1/ 3:40
```

In this example, ports 1/1/2 and 1/1/3 are sending traffic to port 1/1/1. Port 1/1/2 and port 1/1/3 are isolated (not shared by multiple clients) and typically do not exchange traffic with other ports except for the uplink port, 1/1/1.

Syntax: `show relative-utilization num`

The *num* parameter specifies the list number.

NOTE

The example above represents a pure configuration in which traffic is exchanged only by ports 1/1/2 and 1/1/1, and by ports 1/1/3 and 1/1/1. For this reason, the percentages for the two downlink ports equal 100%. In some cases, the percentages do not always equal 100%. This is true in cases where the ports exchange some traffic with other ports in the system or when the downlink ports are configured together in a port-based VLAN.

In the following example, ports 1/1/2 and 1/1/3 are in the same port-based VLAN.

```
device#show relative-utilization 1
uplink: ethe 1
30-sec total uplink packet count = 3011
packet count ratio (%)
  1/ 2:100  1/ 3:100
```

Here is another example showing different data for the same link utilization list. In this example, port 1/1/2 is connected to a hub and is sending traffic to port 1/1/1. Port 1/1/3 is unconnected.

```
device#show relative-utilization 1
uplink: ethe 1
30-sec total uplink packet count = 2996
packet count ratio (%)
  1 /2:100  1/ 3:---
```

System Monitoring

• Overview of system monitoring.....	253
• Configure system monitoring.....	254
• System monitoring on FCX and ICX devices.....	256
• System monitoring for Fabric Adapters.....	257
• System monitoring for Cross Bar.....	259
• System monitoring for Packet Processors.....	260

Overview of system monitoring

System monitoring (sysmon) is a utility that runs as a background process and monitors connections and components of the device for specific errors and logs them. It has a default policy that controls the parameters that are monitored and actions to be taken if a fault is detected. These policies include the type of errors, the threshold for errors to be logged, and the frequency of checking for errors. You can use the CLI commands to configure these policies.

The sysmon utility monitors the hardware error registers to identify errors and failures. You can configure the sysmon timer to define how frequently the sysmon utility queries the hardware error registers. The data generated by the sysmon utility is written to either the sysmon internal log or to the syslog.

Sysmon starts the timer based on the specified timer setting, with the default value as three minutes. After the interval specified by the timer, the utility checks the hardware error registers. If the sysmon utility detects an error in a hardware error register, it increments the relevant error count by 1. Otherwise, it restarts the timer and waits for the given interval. Hardware error registers are cleared when read, so after Sysmon reads the value, they are reset to zero.

Sysmon checks the value of the error counters it maintains and the values specified in the sysmon threshold. If the value of the error counters exceeds the matching threshold, it takes the action specified (logs internally or to the syslog). Otherwise, it restarts the timer and waits for the specified interval before checking for errors again.

To ensure that logging repeating errors does not cause the logs to overflow, you can specify a back-off value that allows the utility to skip the specified number of error instances before logging again. If the error count is smaller than the specified log back-off value, the utility logs the error to the internal log or syslog, restarts the timer and waits for the specified interval before checking for errors again.

Configuration notes and feature limitations

- While system monitoring is supported on all FastIron devices, the types of errors monitored vary according to devices. On FSX devices, the sysmon utility monitors the following for errors:
 - Fabric Adapter (FA) for processing and link errors.
 - Cross Bar (XBAR) or Switch Fabric Module (SFM) for processing and link errors.
 - Packet processor (PP) for link errors.

On FCX and ICX devices, the sysmon utility monitors the following errors:

- Link errors.
- ECC errors.
- By default, system monitoring starts on system boot up and runs in the background every three minutes. You can configure, disable, or enable, the time interval through the CLI; however, if you define the system monitoring interval at the global level, this value overrides the individual settings. Valid range for the sysmon timer is 1 to 60 minutes.

- You can define a system monitoring threshold that is defined as N/W , where N is the number of error events in a specified window (W) of consecutive polling periods. When the threshold is reached, the action that is defined is performed. The threshold enables the `sysmon` utility to ignore random errors that occur because of corrupted data coming in to the device, and perform the action only for errors generated because of device failure. A threshold of $1/W$ means no threshold.
- You can choose the log action as either to the internal `sysmon` buffer or to the `syslog`. If you choose the internal `sysmon` buffer, logs that are written beyond the limit of the `sysmon` buffer rolls over. If you choose logging to `syslog`, messages are sent to the configured `syslog` servers.

Configure system monitoring

You can use the following commands at the privileged EXEC level to globally configure the `sysmon` utility:

- [disable system-monitoring all](#) on page 254
- [enable system-monitoring all](#) on page 254
- [sysmon timer](#) on page 254

In addition, you can enable or disable system monitoring for each event type from the CLI, with each event type having separate threshold and log back off values.

disable system-monitoring all

Disables system monitoring at the global level for all types.

disable system-monitoring all

Privileged EXEC mode.

Disabling `sysmon` at the global level disables any individually configured and enabled `sysmon` tasks as well. However, any `sysmon` configuration that is made, including global and event-specific configuration are retained.

The following example disables system monitoring:

```
Brocade# disable system-monitoring all
```

enable system-monitoring all

Enables system monitoring at the global level for all event types.

enable system-monitoring all

Privileged EXEC mode.

This command enables system monitoring globally, and covers all event-specific system monitoring configuration as well. If specific configuration is not made for different types, default values defined at the global level are used.

The following example enables all system monitoring tasks at the global level:

```
Brocade# enable system-monitoring all
```

sysmon timer

Configures the global system monitoring timer.

sysmon timer minutes

minutes Specifies the system monitoring timer in minutes. The range of values is 1 through 60. The default value is 3.

Global configuration mode.

The following example sets the system monitoring timer to five minutes:

```
Brocade(config)# sysmon timer 5
```

sysmon log-backoff

Defines the number of times to skip logging an event before logging again at the global level. The **no** form of this command resets the parameter to default value.

sysmon log-backoff number

no sysmon log-backoff

number Specifies the number of times to skip an event logging before logging again.

Global configuration mode.

Logging every error may not provide any new information, but adds significantly to the number of error entries that need to be analyzed. You can configure the system monitoring utility to ignore a certain number of errors (within a stream of consecutive errors) before writing the entry to the log again.

This option helps you further isolate issues that randomly occur from issues because of device failure. The sysmon utility keeps a counter of the number of times the threshold value is exceed. If the number exceeds the back-off value, the error is logged as specified by the action option.

The following example sets the number of times to skip logging to 20.

```
Brocade(config)# sysmon log-backoff 20
```

sysmon threshold

Defines the threshold for errors at the global level. The **no** form of this command resets the threshold configuration to default values.

sysmon threshold events polling-interval

no sysmon threshold

events Specifies the threshold in terms of the number of events. Valid values are 1 through 10. When expressed in the command, the default value is 2.

polling-interval Specifies the number of polling windows. The device polls the internal registers at the interval specified by the sysmon timer value. Valid values 1-32. However, the polling window number must be equal or greater than the number of events. When expressed in the command, the default value is 10.

Global configuration mode.

The type-specific threshold values that you define overrides the global threshold value for each event. However, if you define the global value later, the latest value prevails. The threshold is defined as N/W, where N is the number of events, and W is the number of consecutive polling periods. When the threshold is reached, actions configured for this event type will take place. Note that a threshold of 1/W implies that there is no threshold, and the action will always be triggered.

The following example sets the threshold to 3 events over 7 consecutive polling periods:

```
Brocade(config)# sysmon threshold 3 7
```

System monitoring on FCX and ICX devices

On FCX and ICX devices, system monitoring monitors the following errors:

- ECC errors.
- Link errors.

These errors are monitored on a stack unit basis.

Use the following commands configure and display the status of system monitoring on fabric adaptors:

- [sysmon ecc-error](#) on page 256
- [sysmon link-error](#) on page 256

sysmon ecc-error

Configures how sysmon handles ECC errors. The **no** version of this command disables system monitoring on internal ECC errors.

sysmon ecc-error -count { threshold *events polling-interval* | log-backoff *value* | action { none | syslog } }

no sysmon fa error-count

threshold	Defines the threshold for errors. The threshold is defined as N/W, where N is the number of events, and W is the number of consecutive polling periods. When the threshold is reached, actions configured for this event type will take place. Note that a threshold of 1/W implies that there is no threshold, and the action will always be triggered.
events	Specifies the threshold in terms of the number of events. Valid values are 1 through 10.
polling-interval	Specifies the number of polling windows. The device polls the internal registers at the interval specified by the sysmon timer value. Valid values 1-32. However, the polling window number must be equal or greater than the number of events.
log-backoff	If an error condition persists, it will be continuously logged (internally and/or externally to syslog as defined by the action). The log back-off count skips configured number of logs before logging again.
action	Specifies the action to take when error count exceeds the specified threshold and log back-off values.
none	The error is logged in the internal sysmon logs. This is the default value.
syslog	The error is logged to syslog.

Global configuration mode.

This command is supported only on FCX and ICX devices.

The following example configures system monitoring for fabric adaptor errors:

```
Brocade(config)# sysmon ecc-error threshold 3 7
Brocade(config)# sysmon ecc-error action syslog
Brocade(config)# sysmon ecc-error log-backoff 15
```

sysmon link-error

Configures how sysmon handles link errors. The **no** version of this command disables system monitoring on link errors.

sysmon link-error { threshold *events polling-interval* | log-backoff *value* | action { none | syslog } }

no sysmon link-error

threshold	Defines the threshold for errors. The threshold is defined as N/W, where N is the number of events, and W is the number of consecutive polling periods. When the threshold is reached, actions configured for this event type will take place. Note that a threshold of 1/W implies that there is no threshold, and the action will always be triggered.
------------------	--

	events	Specifies the threshold in terms of the number of events. Valid values are 1 through 10.
	polling-interval	Specifies the number of polling windows. The device polls the internal registers at the interval specified by the sysmon timer value. Valid values 1-32. However, the polling window number must be equal or greater than the number of events.
log-backoff		If an error condition persists, it will be continuously logged (internally and/or externally to syslog as defined by the action). The log back-off count skips configured number of logs before logging again.
action		Specifies the action to take when the error count exceeds the specified threshold and log back-off values.
	none	The error is logged in the internal sysmon logs. This is the default value.
	syslog	The error is logged to syslog.

Global configuration mode.

This command is supported only on FCX and ICX devices.

The following example configures system monitoring for fabric adaptor errors:

```
Brocade(config)# sysmon link-error threshold 3 7
Brocade(config)# sysmon link-error action syslog
Brocade(config)# sysmon link-error log-backoff 15
```

System monitoring for Fabric Adapters

On FSX devices, system monitoring for fabric adaptors monitor errors such as the following:

- End of Packet (EoP) or Start of Packet (SoP) errors
- Cyclic Redundancy Check (CRC) errors
- Packets dropped due to congestion

In addition to the error count, sysmon also checks for connectivity of FA links. This happens at the interval defined by the sysmon-timer command generally or specifically for FA.

Use the following commands configure and display the status of system monitoring on fabric adaptors:

- [sysmon fa error-count](#) on page 257
- [sysmon fa link](#) on page 258
- [show sysmon counters](#) on page 263
- [show sysmon logs](#) on page 262
- [show sysmon config](#) on page 266

sysmon fa error-count

Configures how sysmon handles fabric adaptor-related errors. The **no** version of this command disables system monitoring on fabric adaptors.

sysmon fa error-count { **threshold** *events polling-interval* | **log-backoff** *value* | **action** { **none** | **syslog** } }

no sysmon fa error-count

threshold Defines the threshold for errors. The threshold is defined as N/W, where N is the number of events, and W is the number of consecutive polling periods. When the threshold is reached, actions configured for this event type will take place. Note that a threshold of 1/W implies that there is no threshold, and the action will always be triggered.

events Specifies the threshold in terms of the number of events. Valid values are 1 through 10.

	<i>polling-interval</i>	Specifies the number of polling windows. The device polls the internal registers at the interval specified by the sysmon timer value. Valid values 1-32. However, the polling window number must be equal or greater than the number of events.
log-backoff		If an error condition persists, it will be continuously logged (internally and/or externally to syslog as defined by the action). The log back-off count skips configured number of logs before logging again.
action		Specifies the action to take when a fabric adapter error count exceeds the specified threshold and log back-off values.
	none	The error is logged in the internal sysmon logs. This is the default value.
	syslog	The error is logged to syslog.

Global configuration mode.

This command is supported only on FSX devices.

The following example configures system monitoring for fabric adaptor errors:

```
Brocade(config)# sysmon fa error-count threshold 3 7
Brocade(config)# sysmon fa error-count action syslog
Brocade(config)# sysmon fa error-count log-backoff 15
```

sysmon fa link

Configures system monitoring for link errors on all or specified fabric adaptors. The **no** form of this command resets the parameters to default values.

sysmon fa link { **threshold** *events* *polling-interval* | **log-backoff** *value* | **action** { **none** | **syslog** } }

no sysmon fa link

threshold		Defines the failure threshold for the fabric adapter link error event. The threshold is defined as N/W, where N is the number of events, and W is the number of consecutive polling periods. When the threshold is reached, actions configured for this event type will take place. Note that a threshold of 1/W implies that there is no threshold, and no event will be triggered.
	events	Specifies the threshold in terms of the number of events. Valid values are 1 through 10.
	<i>polling-interval</i>	Specifies the number of polling windows. The device polls the internal registers at the interval specified by the sysmon timer value. Valid values 1-32. However, the polling window number must be equal or greater than the number of events.
log-backoff		If an error condition persists, it will be continuously logged (internally and/or externally). The log back-off count skips configured number of logs before logging again. This avoids overflow of the internal log or of the syslog.
action		Specifies the action to take when a fabric adapter link error exceeds the specified threshold and log back-off values.
	none	No action is taken. This is the default.
	syslog	The error is logged to syslog.

Global configuration mode.

This command is supported only on FSX devices.

The following example configures the sysmon options for fabric adaptor links:

```
Brocade(config)# sysmon fa link threshold 3 7
Brocade(config)# sysmon fa link action syslog
Brocade(config)# sysmon fa link log-backoff 15
```

System monitoring for Cross Bar

On FSX devices, errors typically detected in the cross bar include:

- Bad (IP) headers
- Bad length errors
- Reformat errors

Besides the error count, sysmon also checks for connectivity of SFM/XBAR links. This happens at the interval defined by the sysmon-timer command generally or specifically for cross bar.

Use the following commands to configure and display the statistics of cross bar or switch fabric module:

- [sysmon xbar error-count](#) on page 259
- [sysmon xbar link](#) on page 260
- [show sysmon logs](#) on page 262
- [show sysmon counters](#) on page 263
- [show sysmon config](#) on page 266
- [show sysmon system sfm](#) on page 267

sysmon xbar error-count

Configures system monitoring for cross bar errors. The **no** form of this command resets the parameters to default values.

sysmon xbar error-count { *threshold events polling-interval* | *log-backoff value* | **action** { *none* | *syslog* } }

no sysmon xbar error-count

threshold	Defines the failure threshold for the cross bar error-count event. The threshold is defined as N/W, where N is the number of events, and W is the number of consecutive polling periods. When the threshold is reached, actions configured for this event type will take place. Note that a threshold of 1/W implies that there is no threshold, and no event will be triggered.
events	Specifies the threshold in terms of the number of events. Valid values are 1 through 10.
polling-interval	Specifies the number of polling windows. The device polls the internal registers at the interval specified by the sysmon timer value. Valid values 1-32. However, the polling window number must be equal or greater than the number of events.
log-backoff	If an error condition persists, it will be continuously logged (internally and/or externally). The log back-off count skips configured number of logs before logging again. This avoids overflow of the internal log or of the syslog.
action	Specifies the action to take when the error count exceeds the specified threshold and log back-off values.
none	No action is taken.
syslog	The error is logged to syslog.

Global configuration mode.

This command is supported only on FSX devices.

The following example configures system monitoring for cross bar errors.

```
Brocade(config)# sysmon xbar error-count threshold 3 7
Brocade(config)# sysmon xbar error-count action syslog
Brocade(config)# sysmon xbar error-count log-backoff 15
```

sysmon xbar link

Configures the sysmon parameters for the crossbar link. The **no** form of this command resets the parameters to default values.

```
sysmon xbar link { threshold events polling-interval | log-backoff value | action { none | syslog } }
```

no sysmon xbar link

threshold Defines the failure threshold for the fabric adapter error-count event. The threshold is defined as N/W, where N is the number of events, and W is the number of consecutive polling periods. When the threshold is reached, actions configured for this event type will take place. Note that a threshold of 1/W implies that there is no threshold, and no event will be triggered.

events Specifies the threshold in terms of the number of events. Valid values are 1 through 10.

polling-interval Specifies the number of polling windows. The device polls the internal registers at the interval specified by the sysmon timer value. Valid values 1-32. However, the polling window number must be equal or greater than the number of events.

log-backoff If an error condition persists, it will be continuously logged (internally and/or externally). The log back-off count skips configured number of logs before logging again. This avoids overflow of the internal log or of the syslog.

action Specifies the action to take when the error count exceeds the specified threshold and log back-off values.

none No action is taken.

syslog The error is logged to syslog.

Global configuration mode.

This command is supported only on FSX devices.

The following example configures system monitoring for cross bar link errors:

```
Brocade(config)# sysmon xbar link threshold 3 7
Brocade(config)# sysmon xbar link action syslog
Brocade(config)# sysmon xbar link log-backoff 15
```

System monitoring for Packet Processors

On FSX devices, errors typically detected in packet processors include:

- Parity errors
- Error Checking Code (ECC) errors
- ConfigTable0 errors
- TCAM error
- TCAM action parity errors
- Token bucket priority parity errors
- State variable parity errors
- Link list RAM ECC errors
- FBUF RAM ECC errors
- Egress VLAN parity errors
- Ingress VLAN parity errors
- Layer 2 port isolation parity errors
- Layer 3 port isolation parity errors

- VIDX parity errors

Besides the error count, sysmon also checks for connectivity of SFM/XBAR links. This happens at the interval defined by the sysmon-timer command generally or specifically for cross bar.

Use the following commands to configure and display the statistics of cross bar or switch fabric module:

- [sysmon pp error-count](#) on page 261
- [show sysmon logs](#) on page 262
- [show sysmon counters](#) on page 263
- [show sysmon config](#) on page 266

sysmon pp error-count

Configures the sysmon monitoring parameters for error events in packet processors. The **no** form of this command resets the parameters to default values.

```
sysmon pp error-count { threshold eventspolling-interval | log-backoff value | action { none | syslog } }
```

```
no sysmon pp error-count
```

threshold	Defines the failure threshold for the fabric adapter error-count event. The threshold is defined as N/W, where N is the number of events, and W is the number of consecutive polling periods. When the threshold is reached, actions configured for this event type will take place. Note that a threshold of 1/W implies that there is no threshold, and no event will be triggered.
log-backoff	If an error condition persists, it will be continuously logged (internally and/or externally). The log back-off count skips configured number of logs before logging again. This avoids overflow of the internal log or of the syslog.
action	Specifies the action to take when the error count exceeds the specified threshold and log back-off values.
	none No action is taken. This is the default action.
	syslog The error is logged to syslog.

Global configuration mode.

This is a global configuration for all packet processors-- you cannot configure sysmon parameters for individual packet processors. However, you can display the logs for individual packet processors by specifying the packet processor identifier.

This command is supported only on FSX devices.

The following example configures system monitoring on packet processors:

```
Brocade(config)# sysmon pp error-count threshold 3 7
Brocade(config)# sysmon pp error-count action syslog
Brocade(config)# sysmon pp error-count log-backoff 15
```

clear sysmon counters

Clears sysmon counters for all or specific event types.

```
clear sysmon counters all
```

```
clear sysmon counters fa { error | link } { all | decimal }
```

```
clear sysmon counters pp error { all | decimal }
```

```
clear sysmon counters xbar { error | link } { all | decimal }
```

```
clear sysmon counters { ecc-error | link-error }
```

all	Clears all sysmon counters.
fa	Clears the fabric adaptor sysmon counters.
	error Clears the fabric adaptor error counters. You can specify all or a fabric adaptor, identified by the index.
	link Clears the fabric adaptor sysmon counters for links. You can specify all or a fabric adaptor identified by the index.
pp error	Clears packet processor sysmon counters. You can specify all or a packet processor identified by the index.
xbar	Clears cross bar sysmon counters for cross bar. You can specify all or a cross bar identified by the index.
	error Clears the cross bar sysmon error counters. You can specify all or a cross bar identified by the index.
	link Clears the cross bar sysmon counters for links. You can specify all or a cross bar identified by the index.
ecc-error	Clears the ECC error count on FCX and ICX devices. This option is not supported on FSX devices.
	stack-unit Specifies the stack unit on which errors to be cleared.
	all Specifies that all stack units are cleared of errors.
link-error	Clears the link error count on FCX and ICX devices. This option is not supported on FSX devices.
	stack-unit Specifies the stack unit on which errors to be cleared.
	all Specifies that all stack units are cleared of errors.

Global configuration mode.

The following example clears the fabric adaptor sysmon counters.

```
Brocade(config)# clear sysmon counters fa error all
```

show sysmon logs

Displays the entries written to syslog for all event types if the action specified is to log them into syslog. If the action specified is **none**, the sysmon logs display nothing.

show sysmon logs

Privileged EXEC mode.

Global configuration mode.

The following example displays the syslog entries that were made by sysmon if the action specified either at the global level or type level was to log the events to syslog. If the action specified was **none**, no syslog entries exist.

```
Brocade(config)# show sysmon logs
Aug 3 03:59:22:C:Sysmon:XBAR LINK: SFM1/XBAR1/FPOR0 -- NO SYNC
Aug 3 03:59:22:C:Sysmon:FA Link: SLOT9/FA16/Link0 -- HG.Link error
Aug 3 03:58:22:W:Sysmon:PP ERROR: SLOT4/PP6 error occurred
Aug 3 03:59:34:W:Sysmon:FA ERROR: SLOT1/FA0 error occurred
Aug 3 03:60:34:W:Sysmon:XBAR ERROR: SFM1/XBAR1/FPOR2 -error occurred
```

The following table describes the output of this command:

TABLE 44 show sysmon log s command output fields

Field	Description
Date and time	Aug 3 03:59:22
Critical or Warning	A 'C' indicates a critical error and a 'W' indicates a warning.
Sysmon	Message coming from Sysmon
Event type	Possible values are FA ERROR, FA Link, XBAR ERROR, XBAR LINK, or PP ERROR
Component identifier	Identifies the component of the system where the error was detected
Error	A brief description of the error

show sysmon counters

Displays sysmon counters for all or specific event types.

show sysmon counters type { error | link }

show sysmon counters { ecc-error | link-error }

type The event type for which sysmon counters are displayed. For FSX devices, the options are all, fa (fabric adapter), pp (packet processor), and xbar (cross bar). For FCX and ICX devices, the options are ecc-error and link-error. The default value is all.

error Displays the error counter for the specified event type.

link Displays the link error counters. You can specify either all or specific links.

ecc-error Displays the ECC error count on FCX and ICX devices. This option is not supported on FSX devices.

stack-unit Specifies the stack unit on which errors to be displayed.

all Displays errors for all stack units.

link-error Displays the link error count on FCX and ICX devices. This option is not supported on FSX devices.

stack-unit Specifies the stack unit on which errors to be displayed.

all Displays errors for all stack units.

Privileged EXEC mode.

Global configuration mode.

The following displays all fabric adaptor statistics on an FSX device:

```

Brocade# show sysmon counters fa link all
Sysmon FA HG.link error detected (number of times)
link2          FA-link0          FA-link1          FA-
SLOT          FA-dev          Sync/FC (RX, TX)  Sync/FC (RX, TX)  Sync/
FC (RX, TX) Sync/FC (RX, TX)
1             0             0/(0,0)           0/(0,0)           0/(0,0)
2             2             0/(0,0)           0/(0,0)           0/(0,0)
9             16            1751/(1750,1750) 0/(0,0)           0/(0,0) 0/
(0,0)
9             17             0/(0,0)           0/(0,0)           0/(0,0)

```

The following example displays the error events that sysmon has recorded for the fabric adaptor O.

```

Brocade# show sysmon counters fa error 0
Sysmon error detected on: SLOT 1, FA 0(number of times)
****PUMA Device 0 VOQUnit0 error detect
Set 0 EnQ Drop detect = 0
Set 1 EnQ Drop detect = 0
Set 2 EnQ Drop detect = 0
Set 3 EnQ Drop detect = 0
tail drop detect = 0 filter drop detect = 0, ecc drop detect = 0
****PUMA Device 0 VOQUnit1 error detect
Set 0 EnQ Drop detect = 0
Set 1 EnQ Drop detect = 0
Set 2 EnQ Drop detect = 0
Set 3 EnQ Drop detect = 0
tail drop detect = 0 filter drop detect = 0, ecc drop detect = 0
****PUMA Device 0 CRX error detect
CRC detect = 0, Lost SOP.EOP detect = 0, no egress Buf detect = 0
fifo full detect = 0, UC congest detect = 0, MC congest detect = 0
bad buf alloc detect = 0, e2e drop detect = 0

```

The following example shows the crossbar errors for the switch fabric module O.

```

Brocade# show sysmon counters xbar error 0
Sysmon SFM 1 xbar 0 HG.link Rx error detected (number of times)
HG.link  BadLen  BadHeader  ReformatErr
0          0          0          0
1          0          0          0
2          0          1          0
3          0          0          0
4          0          0          0
5          0          0          0
6          0          0          0
7          0          0          0
8          0          0          0
9          0          0          0
10         0          0          0
11         0          0          0

```

The following example displays the cross bar link errors for the SFM module O.

```

Brocade# show sysmon counters xbar link 0
Sysmon SFM 0 xbar 1 HG.link NO-SYNC detected (number of times)
HG.link  NO-SYNC
0          0
1          0
2          0
3          0
4          0
5          1757
6          0
7          0
8          0
9          0
10         0
11         0

```

The following example displays the error counter for the specified packet processor O.

```

Brocade# show sysmon counter pp error 0
Sysmon error detected on: SLOT 1, PP 0(number of times)
****PUMA Device 0 Buffer SRAM error detect
Ingress buffer error detect = 0
Egress buffer error detect = 1
****PUMA Device 0 Control SRAM error detect
CSU : Parity error detect = 0, ECC error detect = 0
LPM0: Parity error detect = 0, ECC error detect = 0
LPM1: Parity error detect = 0, ECC error detect = 0
LPM2: Parity error detect = 0, ECC error detect = 0
LPM3: Parity error detect = 0, ECC error detect = 0

```


The following example displays all error counter data on an FCX device:

```

Brocade(config)#show sysmon counters all
Sysmon error detected on: Stacking Unit 1 (number of times)
****Stacking unit 1 (FCX) Link error detect
Port 24
  Link error detect = 0 remote fault detect = 0 lane error detect = 0
Port 25
  Link error detect = 0 remote fault detect = 0 lane error detect = 0
Port 26
  Link error detect = 0 remote fault detect = 0 lane error detect = 0
Port 27
  Link error detect = 0 remote fault detect = 0 lane error detect = 0
=====
Sysmon error detected on: Stacking Unit 2 (number of times)
****Stacking unit 2 (FCX) Link error detect
Port 24
  Link error detect = 0 remote fault detect = 0 lane error detect = 0
Port 25
  Link error detect = 0 remote fault detect = 0 lane error detect = 0
Port 26
  Link error detect = 0 remote fault detect = 0 lane error detect = 0
Port 27
  Link error detect = 0 remote fault detect = 0 lane error detect = 0
=====
Sysmon error detected on: Stacking Unit 3 (number of times)
****Stacking unit 3 (FCX) Link error detect
Port 24
  Link error detect = 0 remote fault detect = 0 lane error detect = 0
Port 25
  Link error detect = 0 remote fault detect = 0 lane error detect = 0
Port 26
  Link error detect = 0 remote fault detect = 0 lane error detect = 0
Port 27
  Link error detect = 0 remote fault detect = 0 lane error detect = 0
=====
Sysmon error detected on: Stacking Unit 4 (number of times)
****Stacking unit 4 (FCX) Link error detect
Port 24
  Link error detect = 0 remote fault detect = 0 lane error detect = 0
Port 25
  Link error detect = 0 remote fault detect = 0 lane error detect = 0
Port 26
  Link error detect = 0 remote fault detect = 0 lane error detect = 0
Port 27
  Link error detect = 0 remote fault detect = 0 lane error detect = 0
=====
Sysmon error detected on: Stacking Unit 5 (number of times)
****Stacking unit 5 (FCX) Link error detect
Port 24
  Link error detect = 0 remote fault detect = 0 lane error detect = 0
Port 25
  Link error detect = 0 remote fault detect = 0 lane error detect = 0
Port 26
  Link error detect = 0 remote fault detect = 0 lane error detect = 0
Port 27
  Link error detect = 0 remote fault detect = 0 lane error detect = 0
=====
Sysmon ECC error detected on: Stacking Unit 1 (number of times)
****Stacking unit 1 (ICX) ecc error detect
ECC one-time error detect = 0 ECC two-time error detect = 0
=====
Sysmon ECC error detected on: Stacking Unit 2 (number of times)
****Stacking unit 2 (ICX) ecc error detect
ECC one-time error detect = 0 ECC two-time error detect = 0
=====
Sysmon ECC error detected on: Stacking Unit 3 (number of times)
****Stacking unit 3 (ICX) ecc error detect
ECC one-time error detect = 0 ECC two-time error detect = 0
=====
Sysmon ECC error detected on: Stacking Unit 4 (number of times)

```

```

****Stacking unit 4 (ICX) ecc error detect
ECC one-time error detect = 0 ECC two-time error detect = 0
=====
Sysmon ECC error detected on: Stacking Unit 5 (number of times)
****Stacking unit 5 (ICX) ecc error detect
ECC one-time error detect = 0 ECC two-time error detect = 0
=====

```

show sysmon config

Displays the complete sysmon configuration, including the global configuration and the event-specific configuration.

show sysmon config

User EXEC mode.

Privileged EXEC mode.

The following command displays the sysmon configuration on an FSX device. The global configuration is displayed first, followed by the configuration for specific events.

```

Brocade> show sysmon config
=====
System Monitoring (Sysmon) is: enabled
Sysmon timer = 3 minutes
=====
Threshold: Times error detected / Consecutive times event polling.
Log Backoff Number: Number of times skip log before log again.
=====
Sysmon Event: FA_ERROR_COUNT (Enabled)
Threshold:      2/10
Log Backoff Number: 10
Action: log(internal) /syslog
Sysmon Event: FA_LINK (Enabled)
Threshold:      2/10
Log Backoff Number: 10
Action: log(internal) /syslog
Sysmon Event: XBAR_ERROR_COUNT (Enabled)
Threshold:      2/10
Log Backoff Number: 10
Action: log(internal) /syslog
Sysmon Event: XBAR_LINK (Enabled)
Threshold:      2/10
Log Backoff Number: 10
Action: log(internal) /syslog
Sysmon Event: PP_ERROR_COUNT (Enabled)
Threshold:      2/10
Log Backoff Number: 10
Action: log(internal) /syslog

```

The following example displays the sysmon configuration on an FCX device:

```

Brocade(config)#show sysmon config
=====
System Monitoring (Sysmon) is: enabled
Sysmon timer = 3 minutes
=====
Threshold: Times error detected / Consecutive times event polling.
Log Backoff Numner: Number of times skip log before log again.
=====
Sysmon Event: LINK_STATUS (Enabled)
Threshold:      2/10
Log Backoff Number: 10
Action: log(internal) /syslog
Sysmon Event: ECC_STATS (Enabled)
Threshold:      2/10
Log Backoff Number: 10
Action: log(internal) /syslog

```

show sysmon system sfm

Displays the status of the switch fabric modules.

show sysmon system sfm { all | number }

all Displays the statistics for all SFMs on the device.

number Specifies the SFM ID for which the statistics is to be displayed.

User EXEC mode.

Privileged EXEC mode.

Global configuration mode.

This command is supported only on FSX devices.

The following command displays the statistics for all SFMs on the device.

```
Brocade(config)# show sysmon system sfm all
SFM= 1,Xbar= 2
 X-link  Status  FlowCtrl  FA-dev/Link  Status  FlowCtrl
  2      OK      0x0      19/0         OK      0x0
  3      OK      0x0      13/0         OK      0x0
  4      OK      0x0      0/1          OK      --
  5      OK      0x0      3/0          OK      0x0
  7      OK      0x0      10/1         OK      --
  8      OK      0x0      7/0          OK      0x0
  9      OK      0x0      17/0         OK      0x0
=====
SFM= 1,Xbar= 3
 X-link  Status  FlowCtrl  FA-dev/Link  Status  FlowCtrl
  1      OK      0x0      17/1         OK      0x0
  2      OK      0x0      3/1          OK      0x0
  4      OK      0x0      0/2          OK      --
  5      OK      0x0      19/1         OK      0x0
  7      OK      0x0      10/2         OK      --
 10      OK      0x0      7/1          OK      0x0
 11      OK      0x0      13/1         OK      0x0
=====
```


Syslog

- [About Syslog messages.....](#) 269
- [Displaying Syslog messages.....](#) 269
- [Syslog service configuration.....](#) 272

About Syslog messages

Brocade 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.

You also can specify the IP address or host name of up to six Syslog servers. When you specify a Syslog server, the Brocade 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 Brocade local Syslog buffer is cleared during a system reload or reboot, but the Syslog messages sent to the Syslog server remain on the server.

NOTE

To enable the Brocade device to retain Syslog messages after a soft reboot (**reload** command). Refer to [Retaining Syslog messages after a soft reboot](#) on page 280.

The Syslog service on a Syslog server receives logging messages from applications on the local host or from devices such as a Layer 2 Switch or Layer 3 Switch. 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 local buffer, enter the **show logging** command at any level of the CLI. The following shows an example display output.

```
device>#show logging
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
  Buffer logging: level ACDMEINW, 3 messages logged
  level code: A=alert C=critical D=debugging M=emergency E=error
```

```

        I=informational N=notification W=warning
Static Log Buffer:
Dec 15 19:04:14:A:Fan 1, fan on right connector, failed
Dynamic Log Buffer (50 entries):
Dec 15 18:46:17:I:Interface ethernet 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 273.

Enabling real-time display of Syslog messages

By default, to view Syslog messages generated by a Brocade device, you need to display the Syslog buffer or the log on a Syslog server used by the Brocade 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. However, to enable display of real-time Syslog messages in Telnet or SSH sessions, you also must enable display within the individual sessions.

To enable real-time display of Syslog messages, enter the following command at the global CONFIG level of the CLI.

```
device(config)#logging console
```

Syntax: [no] loggingconsole

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.

Enabling real-time display for a Telnet or SSH session

To also enable the real-time display for a Telnet or SSH session, enter the following command from the Privileged EXEC level of the session.

```
telnet@device#terminal monitor
Syslog trace was turned ON
```

Syntax: terminal monitor

Notice that the CLI displays a message to indicate the status change for the feature. To disable the feature in the management session, enter the **terminal monitor** command again. The command toggles the feature on and off.

```
telnet@device#terminal monitor
Syslog trace was turned OFF
```

Here is an example of how the Syslog messages are displayed.

```
telnet@device#terminal monitor
Syslog trace was turned ON
SYSLOG: <9>device, Power supply 2, power supply on left connector, failed
SYSLOG: <14>device, Interface ethernet 6, state down
SYSLOG: <14>device, Interface ethernet 2, state up
    
```

Broadcast, unknown unicast, and multicast suppression Syslog and SNMP notification

Rate limiting broadcast, unknown unicast, and multicast (BUM) traffic protects a switch, router node, or network from Denial of Service (DoS) attacks or unintentional traffic configurations. When an incoming packet exceeds the maximum number of bytes that you set with rate limiting, a Syslog notification is generated.

Restrictions and limitations

- All of the restrictions that are applicable while configuring ACLs on an interface apply to this feature. Refer to the *Brocade FastIron Security Configuration Guide* for the restrictions that apply to ACLs. The main restrictions are:
 - You cannot change the ports VLAN membership.
 - You cannot apply another ACL or MAC-filter to the interface.
- By default, the Syslog logs once a minute; however, you can configure Syslog notifications so that they log at a maximum interval of every 10 minutes.

Enabling BUM suppression logging

Follow these steps to enable logging.

Rate limiting must be enabled.

1. Enter configuration mode.

```
device# configure terminal
```

2. Enter Ethernet configuration mode.

```
device(config)# interface ethernet 1/1/1
```

3. Enable rate limiting.

```
device(config-if-e10000-1/1/1)# broadcast limit 8388607 kbps
```

Broadcast is used in this example, multicast and unknown unicast are the same with the command name switched to either **multicast** or **unknown-unicast**.

4. Enable logging when the limit exceeds kbps.

```
device(config-if-e10000-1/1/1)# broadcast limit 100 kbps log
```

Broadcast is used in this example, multicast and unknown unicast are the same with the command name switched to either **multicast** or **unknown-unicast**.

5. Globally configure the log interval.

```
device(config)# rate-limit-log 6
device(config)# exit
```

6. Verify the logging interval.

```
device(config)# show running-config | include rate-limit-log
rate-limit-log 6
```

7. Verify the configuration.

```
device# show logging | include 1/1/1
Jan 13 12:02:12:I:Security: Interface ethernet 1/1/1 reached the Broadcast traffic limit and 1434 kB
are dropped
```

Enabling BUM suppression logging configuration example

```
device# configure terminal
device(config)# interface ethernet 1/1/1
device(config-if-e10000-1/1/1)# broadcast limit 8388607
device(config-if-e10000-1/1/1)# broadcast limit 100 kbps log
device(config)# rate-limit-log 6
device(config)# show running-config | include rate-limit-log
device(config)# exit
device# show logging | include 1/1/1
```

Viewing BUM suppression Syslog notifications

Use these commands to display BUM suppression syslog notification information.

Use the **show logging** command to view the BUM suppression Syslog notifications for all interfaces.

```
device# show logging
Jan 13 12:02:12:I:Security: Interface ethernet 1/1/1 reached the Broadcast traffic limit and 11620 kB are
dropped
Jan 13 12:14:23:I:Security: Interface ethernet 1/3/12 reached the Multicast traffic limit and 870 kB are
dropped
Jan 13 12:45:38:I:Security: Interface ethernet 3/2/14 reached the Unknown-Unicast traffic limit and 2321 kB
are dropped
```

The first section of the output is `mmm dd hh:mm:ss:Info:System`.

To view the BUM suppression Syslog notifications for a specific interface use the following command.

```
device# show logging | include 1/1/1
Jan 13 12:02:12:I:Security: Interface ethernet 1/1/1 reached the Broadcast traffic limit and 11620 kB are
dropped
```

Displaying real-time Syslog messages

Any terminal logged on to a Brocade switch can receive real-time Syslog messages when the **terminal monitor** command is issued.

Syslog service configuration

The procedures in this section describe how to perform the following Syslog configuration tasks:

- Specify a Syslog server. You can configure the Brocade device to use up to six Syslog servers. (Use of a Syslog server is optional. The system can hold up to 1000 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 Brocade device, enter the following command from any level of the CLI.

```
device>#show logging
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
  Buffer logging: level ACDMEINW, 3 messages logged
  level code: A=alert C=critical D=debugging M=emergency E=error
               I=informational N=notification W=warning
Static Log Buffer:
Dec 15 19:04:14:A:Fan 1, fan on right connector, failed
Dynamic Log Buffer (50 entries):
Dec 15 18:46:17:I:Interface ethernet 1/1/4, 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 45 CLI display of Syslog buffer configuration

Field	Definition
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 278. 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 or equivalent Web Management Interface option. Refer to Clearing the Syslog messages from the local buffer on page 281.
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 buffers:

- Static - logs power supply failures, fan failures, and temperature warning or shutdown messages
- Dynamic - logs all other message types

In the static log, new messages replace older ones, so only the most recent message is displayed. For example, only the most recent temperature warning message will be present in the log. If multiple temperature warning messages are sent to the log, the latest one replaces the previous one. The static buffer is not configurable.

The message types that appear in the static buffer do not appear in the dynamic buffer. The dynamic buffer contains up to the maximum number of messages configured for the buffer (50 by default), then begins removing the oldest messages (at the bottom of the log) to make room for new ones.

The static and dynamic buffers are both displayed when you display the log.

```
device#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 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.

Clearing log entries

When you clear log entries, you can selectively clear the static or dynamic buffer, or you can clear both. For example, to clear only the dynamic buffer, enter the following command at the Privileged EXEC level.

```
device#clear logging dynamic-buffer
```

Syntax: clear logging [*dynamic-buffer* | *static-buffer*]

You can specify *dynamic-buffer* to clear the dynamic buffer or *static-buffer* to clear the static buffer. If you do not specify a buffer, both buffers are cleared.

Time stamps

The contents of the time stamp differ depending on whether you have set the time and date on the onboard system clock.

If you have set the time and date on the onboard system clock, the date and time are shown in the following format:

```
mm dd hh:mm:ss
```

where

- **mm** - abbreviation for the name of the month
- **dd** - day
- **hh** - hours
- **mm** - minutes
- **ss** - seconds

For example, "Oct 15 17:38:03" means October 15 at 5:38 PM and 3 seconds.

If you have not set the time and date on the onboard system clock, the time stamp shows the amount of time that has passed since the device was booted, in the following format.

num d num h num m num s

where

- *num d* - day
- *num h* - hours
- *num m* - minutes
- *num s* - 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 with the onboard clock set

The example shows the format of messages on a device where the onboard system clock has been set. Each time stamp shows the month, the day, and the time of the system clock when the message was generated. For example, the system time when the most recent message (the one at the top) was generated was October 15 at 5:38 PM and 3 seconds.

```
device#show logging
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 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 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 18
0000.001f.77ed) -> 10.99.4.69(http), 1 event(s)
```

Example of Syslog messages on a device with the onboard clock not set

The example shows the format of messages on a device where the onboard system clock is not set. Each time stamp shows the amount of time the device had been running when the message was generated. For example, the most recent message, at the top of the list of messages, was generated when the device had been running for 21 days, seven hours, two minutes, and 40 seconds.

```
device#show logging
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 1/4/18
0000.001f.77ed) -> 10.99.4.69(http), 1 event(s)
19d07h03m30s:warning:list 101 denied tcp 10.157.22.26(0) (Ethernet 1/4/18
0000.001f.77ed) -> 10.99.4.69(http), 1 event(s)
17d06h58m30s:warning:list 101 denied tcp 10.157.22.198(0) (Ethernet 1/4/18
0000.001f.77ed) -> 10.99.4.69(http), 1 event(s)
```

Generating the Syslog specific to RFC 5424

By default, Syslog is generated in accordance with RFC 3164. To provide the maximum amount of information in every Syslog in a structured format, you can enable Syslog logging specific to RFC 5424.

The Syslog that conforms to RFC 5424 has an enhanced Syslog header that helps to identify the type of Syslog, filter the Syslog message, identify the Syslog generation time with year and milliseconds with respect to the time zone, and other enhancements. The Syslog specific to RFC 5424 can be enabled using the **logging enable rfc5424** command. Logging buffer must be cleared before enabling Syslog specific to RFC 5424, otherwise the system throws an error.

NOTE

If the **logging cli-command** command is present in the running configuration, switching between Syslog functionality that follows the default RFC 3164 standard and Syslog specific to RFC 5424 standard is not supported.

The following table provides a comparison of the syslog header information available in the RFC 3164 and RFC 5424 Syslog logging.

TABLE 46 Syslog headers available for RFC 3164 and RFC 5424

Syslog RFC 3164	Syslog RFC 5424
PRIORITY	PRIORITY
	VERSION
TIMESTAMP	TIMESTAMP
HOSTNAME	HOSTNAME
	APP-NAME
	PROCID
	MSGID
	STRUCTURED-DATA
MSG	MSG

RFC 5424 provides the following Syslog headers:

- PRIORITY — This represents both Facility and Severity of the messages as described in RFC 3164.
- VERSION — This field denotes the version of the Syslog protocol specification.
- TIMESTAMP — This is a formalized timestamp. TIMESTAMP denotes the date and time when the event is logged and includes the syslog generation time with the year and milliseconds with respect to the time zone.

The following examples show the date and time format in RFC 5424.

NOTE

The suffix "Z", when applied to a time, denotes a Coordinated Universal Time (UTC) offset of 00:00.

For example, 2015-08-13T22:14:15.003Z represents August 13, 2015 at 10:14:15pm, 3 milliseconds into the next second. The timestamp is in UTC. The timestamp provides millisecond resolution.

- HOSTNAME — It identifies the machine that originally sent the Syslog message. The order of preference for the contents of the HOSTNAME field is as follows and may have one of the following values:
 - FQDN
 - Hostname
 - NILVALUE — A field used when the Syslog application is incapable of obtaining its host name.
- APP-NAME — This identifies the device or application from which the message is originated. The APP-NAME is intended for filtering messages on a relay or collector. The NILVALUE is used when the Syslog application is incapable of obtaining its APP-NAME.

- PROCID — This field is often used to provide the process name or process ID associated with a Syslog system. The NILVALUE is present when a process ID is not available.
- MSGID — It identifies the type of message. The NILVALUE is used when the Syslog application does not, or cannot, provide any value.
- STRUCTURED-DATA — This provides a mechanism to express information in a well-defined and interpretable data format as per RFC 5424. STRUCTURED-DATA can contain zero, one, or multiple SD elements. In case of zero structured data elements, the STRUCTURED-DATA field uses NILVALUE.
- MSG — It contains a free-form message that provides information about the event.

Displaying syslog messages generated as per RFC 5424

If Syslog logging specific to RFC 5424 is enabled, the **show logging** command displays the Syslog messages generated in the format as per RFC 5424.

```
device# show logging
Syslog logging: enabled (RFC: 5424, 0 messages dropped, 0 flushes, 0 overruns)
  Buffer logging: level ACDMEINW, 22 messages logged
  level code: A=alert C=critical D=debugging M=emergency E=error
              I=informational N=notification W=warning

Static Log Buffer:
Dec 19 01:36:18:I: brocade - - - [meta sequenceId=8] BOMSystem: Stack unit 1 Power supply 1 is up
Dec 19 01:36:24:I: brocade - - - [meta sequenceId=17] BOMSystem: Stack unit 3 POE Power supply 1 with
748000 mwatts capacity is up
Dec 19 01:36:24:A: brocade - - - [meta sequenceId=19] BOMSystem: Stack unit 3 POE Power supply 2 is down

Dynamic Log Buffer (50 lines):
2012-12-19T01:36:40.798Z:I: brocade - - - [meta sequenceId=23] BOMSystem: Interface ethernet 3/1/23, state
up
2012-12-19T01:36:40.797Z:I: brocade - - - [meta sequenceId=22] BOMSystem: Interface ethernet 3/1/13, state
up
2012-12-19T01:36:40.796Z:I: brocade - - - [meta sequenceId=21] BOMSystem: Interface ethernet 3/1/1, state up
2012-12-19T01:36:24.591Z:A: brocade - - - [meta sequenceId=20] BOMStack unit 3 Power supply 2 is down
2012-12-19T01:36:24.591Z:I: brocade - - - [meta sequenceId=18] BOMSystem: Stack unit 3 Power supply 1 with
748000 mwatts capacity is up
2012-12-19T01:36:23.406Z:I: brocade - - - [meta sequenceId=16] BOMSystem: Interface ethernet 3/3/1, state up
2012-12-19T01:36:22.526Z:I: brocade - - - [meta sequenceId=15] BOMStack: Stack unit 1 has been elected as
ACTIVE unit of the stack system
2012-12-19T01:36:21.297Z:I: brocade - - - [meta sequenceId=14] BOMSystem: Interface ethernet 1/4/1, state up
2012-12-19T01:36:20.858Z:I: brocade - - - [meta sequenceId=13] BOMStack: Stack unit 1 has been elected as
ACTIVE unit of the stack system
2012-12-19T01:36:20.822Z:I: brocade - - - [meta sequenceId=12] BOMStack: Stack unit 3 has been added to the
stack system
2012-12-19T01:36:20.500Z:I: brocade - - - [meta sequenceId=11] BOMSystem: Interface ethernet 1/4/1, state
down
2012-12-19T01:36:19.695Z:I: brocade - - - [meta sequenceId=10] BOMSystem: Interface ethernet 1/4/1, state up
2012-12-19T01:36:18.509Z:I: brocade - - - [meta sequenceId=9] BOMSystem: Stack unit 1 Power supply 1 is u
2012-12-19T01:36:17.865Z:I: brocade - - - [meta sequenceId=7] BOMSystem: Interface ethernet 1/3/1, state up
2012-12-19T01:36:16.466Z:I: brocade - - - [meta sequenceId=6] BOMSystem: Interface ethernet mgmt1, state up
2012-12-19T01:36:16.447Z:I: brocade - - - [meta sequenceId=5] BOMSystem: Warm start
2012-12-19T01:36:16.260Z:D: brocade - - - [meta sequenceId=4] BOMDHCPC: starting dhcp client service on 57
port(s)
2012-12-19T01:36:16.259Z:D: brocade - - - [meta sequenceId=3] BOMDHCPC: Found static IP address 10.20.15.15
subnet mask 255.255.255.0 on port mgmt1
2012-12-19T01:36:16.259Z:D: brocade - - - [meta sequenceId=2] BOMDHCPC: Found static IP address 20.20.20.3
subnet mask 255.255.255.0 on port 1/1/3
2012-12-19T01:36:16.259Z:D: brocade - - - [meta sequenceId=1] BOMDHCPC: Found static IP address 10.10.10.2
subnet mask 255.255.255.0 on port 1/1/1
```

Disabling or re-enabling Syslog

Syslog is enabled by default. To disable it, enter the **logging on** command at the global CONFIG level.

```
device(config)#no logging on
```

Syntax: **[no] logging on** [*udp-port*]

The *udp-port* parameter specifies the application port used for the Syslog facility. The default is 514.

To re-enable logging, re-enter the **logging on** command.

```
device(config)#logging on
```

This command enables local Syslog logging with the following defaults:

- Messages of all severity levels (Emergencies - Debugging) are logged.
- Up to 50 messages are retained in the local Syslog buffer.
- No Syslog server is specified.

Specifying a Syslog server

To specify a Syslog server, enter the **logging host** command.

```
device(config)#logging host 10.0.0.99
```

Syntax: **logginghost** *ip-addr* | *server-name*

Specifying an additional Syslog server

To specify an additional Syslog server, enter the **logging host** command again. You can specify up to six Syslog servers.

```
device(config)#logging host 10.0.0.99
```

Syntax: **logginghost** *ip-addr* | *server-name*

Disabling logging of a message level

To change the message level, disable logging of specific message levels. You must disable the message levels on an individual basis.

For example, to disable logging of debugging and informational messages, enter the following commands.

```
device(config)#no logging buffered debugging
device(config)#no logging buffered informational
```

Syntax: **[no] loggingbuffered** *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.

Changing the number of entries the local buffer can hold

You also can use the **logging buffered** command to change the number of entries the local Syslog buffer can store. For example.

```
device(config)#logging buffered 1000
device(config)#write memory
device(config)#exit
device#reload
```

Syntax: **[no] logging buffered** *num*

The default number of messages is 50. For FastIron devices, you can set the Syslog buffer limit from 1 - 1000 entries.

Local buffer configuration notes

- You must save the configuration and reload the software to place the change into effect.
- If you decrease the size of the buffer, the software clears the buffer before placing the change into effect.
- If you increase the size of the Syslog buffer, the software will clear some of the older locally buffered Syslog messages.

Changing the log facility

The Syslog daemon on the Syslog server uses a facility to determine where to log the messages from the Brocade device. The default facility for messages the Brocade 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 Brocade device to use two Syslog servers, the device uses the same facility on both servers.

```
device(config)#logging facility local0
```

Syntax: **loggingfacility** *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/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/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 interface names in Syslog messages

By default, an interface 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:

```
FastIron(config)# ip show-portname
```

This command is applied globally to all interfaces on Layer 2 Switches and Layer 3 Switches.

Syntax:[no] ip show-portname

By default, Syslog messages show the interface type, such as "ethernet", and so on. For example, you see the following

```
SYSLOG: <14>0d00h02m18s:ICX6610-48P Router System: Interface ethernet 1/1/5, state up
```

However, if ip show-portname is configured and a name has been assigned to the port, the port name replaces the interface type as in the example below, where "port5_name" is the name of the port.

```
SYSLOG: <14>0d00h02m18s:ICX6610-48P Router System: Interface port5_name 1/1/5, state up
```

Also, when you display the messages in the Syslog, you see 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 see "lab2" displayed as in the example below:

```
device# show logging
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
  Buffer logging: level ACDMEINW, 3 messages logged
  level code: A=alert C=critical D=debugging M=emergency E=error
               I=informational N=notification W=warning
Static Log Buffer:
Dec 15 19:04:14:A:Fan 1, fan on right connector, failed
Dynamic Log Buffer (50 entries):
Dec 15 18:46:17:I:Interface ethernet Lab2
, state up
Dec 15 18:45:15:I:Warm start
```

Retaining Syslog messages after a soft reboot

You can configure the device to save the System log (Syslog) after a soft reboot (**reload** command).

Syslog reboot configuration considerations

- If the Syslog buffer size was set to a different value using the CLI command **logging buffered**, the System log will be cleared after a soft reboot, even when this feature (**logging persistence**) is in effect. This will occur only with a soft reboot immediately following a Syslog buffer size change. A soft reboot by itself will not clear the System log. To prevent the system from clearing the System log, leave the number of entries allowed in the Syslog buffer unchanged.
- This feature does not save Syslog messages after a hard reboot. When the Brocade device is power-cycled, the Syslog messages are cleared.
- If *logging persistence* is enabled and you load a new software image on the device, you must first clear the log if you want to reload the device. (Refer to [Clearing the Syslog messages from the local buffer](#) on page 281.)

To configure the device to save the System log messages after a soft reboot, enter the following command.

```
device(config)#logging persistence
```

Syntax: [no] logging persistence

Enter **no logging persistence** to disable this feature after it has been enabled.

Clearing the Syslog messages from the local buffer

To clear the Syslog messages stored in the local buffer of the Brocade device, enter the **clear logging** command.

```
device#clear logging
```

Syntax: clear logging

Syslog messages

- [Brocade Syslog messages.....](#) 283

This section lists all of the Syslog messages. Note that some of the messages apply only to Layer 3 switches.

NOTE

This chapter does not list Syslog messages that can be displayed when a debug option is enabled.

The messages are listed by message level, in the following order, then by message type:

- Emergencies (none)
- Alerts
- Critical
- Errors
- Warnings
- Notifications
- Informational
- Debugging

Brocade Syslog messages

Message	<code>num-modules modules and 1 power supply, need more power supply!!</code>
Explanation	Indicates that the chassis needs more power supplies to run the modules in the chassis.
	The num-modules parameter indicates the number of modules in the chassis.
Message Level	Alert
Message	<code>Fan num , location , failed</code>
Explanation	A fan has failed.
	The num is the fan number.
	The location describes where the failed fan is in the chassis.
Message Level	Alert
Message	<code>MAC Authentication failed for mac-address on portnum</code>
Explanation	RADIUS authentication was successful for the specified mac-address on the specified portnum ; however, the VLAN returned in the RADIUS Access-Accept message did not refer to a valid VLAN or VLAN ID on the Brocade device. This is treated as an authentication failure.
Message Level	Alert
Message	<code>MAC Authentication failed for mac-address on portnum (Invalid User)</code>
Explanation	RADIUS authentication failed for the specified mac-address on the specified portnum because the MAC address sent to the RADIUS server was not found in the RADIUS server users database.
Message Level	Alert
Message	<code>MAC Authentication failed for mac-address on portnum (No VLAN Info received from RADIUS server)</code>

Explanation	RADIUS authentication was successful for the specified mac-address on the specified portnum ; however, dynamic VLAN assignment was enabled for the port, but the RADIUS Access-Accept message did not include VLAN information. This is treated as an authentication failure.
Message Level	Alert
Message	MAC Authentication failed for mac-address on portnum (Port is already in another radius given vlan)
Explanation	RADIUS authentication was successful for the specified mac-address on the specified portnum ; however, the RADIUS Access-Accept message specified a VLAN ID, although the port had previously been moved to a different RADIUS-assigned VLAN. This is treated as an authentication failure.
Message Level	Alert
Message	MAC Authentication failed for mac-address on portnum (RADIUS given vlan does not exist)
Explanation	RADIUS authentication was successful for the specified mac-address on the specified portnum ; however, the RADIUS Access-Accept message specified a VLAN that does not exist in the Brocade configuration. This is treated as an authentication failure.
Message Level	Alert
Message	MAC Authentication failed for mac-address on portnum (RADIUS given VLAN does not match with TAGGED vlan)
Explanation	Multi-device port authentication failed for the mac-address on a tagged port because the packet with this MAC address as the source was tagged with a VLAN ID different from the RADIUS-supplied VLAN ID.
Message Level	Alert
Message	Management module at slot slot-num state changed from module-state to module-state .
Explanation	Indicates a state change in a management module. The slot-num indicates the chassis slot containing the module. The module-state can be one of the following: <ul style="list-style-type: none"> • active • standby • crashed • coming-up • unknown
Message Level	Alert
Message	OSPF LSA Overflow, LSA Type = lsa-type
Explanation	Indicates an LSA database overflow. The lsa-type parameter indicates the type of LSA that experienced the overflow condition. The LSA type is one of the following: <ul style="list-style-type: none"> • 1 - Router • 2 - Network • 3 - Summary • 4 - Summary • 5 - External
Message Level	Alert
Message	OSPF Memory Overflow

Explanation	OSPF has run out of memory.
Message Level	Alert
Message	System: Module in slot slot-num encountered PCI config read error: Bus PCI-bus-number , Dev PCI-device-number , Reg Offset PCI-config-register-offse t .
Explanation	The module encountered a hardware configuration read error.
Message Level	Alert
Message	System: Module in slot slot-num encountered PCI config write error: Bus PCI-bus-number , Dev PCI-device-number , Reg Offset PCI-config-register-offset .
Explanation	The module encountered a hardware configuration write error.
Message Level	Alert
Message	System: Module in slot slot-num encountered PCI memory read error: Mem Addr memory-address
Explanation	The module encountered a hardware memory read error.
	The memory-address is in hexadecimal format.
Message Level	Alert
Message	System: Module in slot slot-num encountered PCI memory write error: Mem Addr memory-address .
Explanation	The module encountered a hardware memory write error.
	The memory-address is in hexadecimal format.
Message Level	Alert
Message	System: Module in slot slot-num encountered unrecoverable PCI bridge validation failure. Module will be deleted.
Explanation	The module encountered an unrecoverable (hardware) bridge validation failure. The module will be disabled or powered down.
Message Level	Alert
Message	System: Module in slot slot-num encountered unrecoverable PCI config read failure. Module will be deleted.
Explanation	The module encountered an unrecoverable hardware configuration read failure. The module will be disabled or powered down.
Message Level	Alert
Message	System: Module in slot slot-num encountered unrecoverable PCI config write failure. Module will be deleted.
Explanation	The module encountered an unrecoverable hardware configuration write failure. The module will be disabled or powered down.
Message Level	Alert
Message	System: Module in slot slot-num encountered unrecoverable PCI device validation failure. Module will be deleted.
Explanation	The module encountered an unrecoverable (hardware) device validation failure. The module will be disabled or powered down.
Message Level	Alert
Message	System: Module in slot slot-num encountered unrecoverable PCI memory read failure. Module will be deleted.
Explanation	The module encountered an unrecoverable hardware memory read failure. The module will be disabled or powered down.

Message Level	Alert
Message	System: Module in slot slot-num encountered unrecoverable PCI memory write failure. Module will be deleted.
Explanation	The module encountered an unrecoverable hardware memory write failure. The module will be disabled or powered down.
Message Level	Alert
Message	System: No Free Tcam Entry available. System will be unstable
Explanation	You must reboot the device.
Message Level	Alert
Message	System: Temperature is over shutdown level, system is going to be reset in num seconds
Explanation	The chassis temperature has risen above shutdown level. The system will be shut down in the amount of time indicated.
Message Level	Alert
Message	Temperature degrees C degrees, warning level warn-degrees C degrees, shutdown level shutdown-degrees C degrees
Explanation	Indicates an over temperature condition on the active module. The degrees value indicates the temperature of the module. The warn-degrees value is the warning threshold temperature configured for the module. The shutdown-degrees value is the shutdown temperature configured for the module.
Message Level	Alert
Message	Authentication shut down portnum due to DOS attack
Explanation	Denial of Service (DoS) attack protection was enabled for multi-device port authentication on the specified portnum , and the per-second rate of RADIUS authentication attempts for the port exceeded the configured limit. The Brocade device considers this to be a DoS attack and disables the port.
Message Level	Critical
Message	BGP4: Not enough memory available to run BGP4
Explanation	The device could not start the BGP4 routing protocol because there is not enough memory available.
Message Level	Debug
Message	DOT1X: Not enough memory
Explanation	There is not enough system memory for 802.1X authentication to take place. Contact Brocade Technical Support.
Message Level	Debug
Message	No of prefixes received from BGP peer ip-addr exceeds maximum prefix-limit...shutdown
Explanation	The Layer 3 switch has received more than the specified maximum number of prefixes from the neighbor, and the Layer 3 switch is therefore shutting down its BGP4 session with the neighbor.
Message Level	Error
Message	IPv6: IPv6 protocol disabled on the device from session-id
Explanation	IPv6 protocol was disabled on the device during the specified session.
Message Level	Informational
Message	IPv6: IPv6 protocol enabled on the device from session-id
Explanation	IPv6 protocol was enabled on the device during the specified session.
Message Level	Informational

Message	MAC Filter applied to port port-id by username from session-id (filter id= filter-ids)
Explanation	Indicates a MAC address filter was applied to the specified port by the specified user during the specified session. session-id can be console, telnet, ssh, or snmp. filter-ids is a list of the MAC address filters that were applied.
Message Level	Informational
Message	MAC Filter removed from port port-id by username from session-id (filter id= filter-ids)
Explanation	Indicates a MAC address filter was removed from the specified port by the specified user during the specified session. session-id can be console, telnet, ssh, or snmp. filter-ids is a list of the MAC address filters that were removed.
Message Level	Informational
Message	Security: Password has been changed for user username from session-id
Explanation	Password of the specified user has been changed during the specified session ID or type. session-id can be console, telnet, ssh, or snmp.
Message Level	Informational
Message	device-name : Logical link on interface ethernet slot#/port# is down.
Explanation	The specified ports were logically brought down while singleton was configured on the port.
Message Level	Informational
Message	device-name : Logical link on interface ethernet slot#/port# is up.
Explanation	The specified ports were logically brought up while singleton was configured on the port.
Message Level	Informational
Message	user-name login to PRIVILEGED mode
Explanation	A user has logged into the Privileged EXEC mode of the CLI. The user-name is the user name.
Message Level	Informational
Message	user-name login to USER EXEC mode
Explanation	A user has logged into the USER EXEC mode of the CLI. The user-name is the user name.
Message Level	Informational
Message	user-name logout from PRIVILEGED mode
Explanation	A user has logged out of Privileged EXEC mode of the CLI. The user-name is the user name.
Message Level	Informational
Message	user-name logout from USER EXEC mode
Explanation	A user has logged out of the USER EXEC mode of the CLI. The user-name is the user name.
Message Level	Informational
Message	ACL ACL id added deleted modified from console telnet ssh snmp session
Explanation	A user created, modified, deleted, or applied an ACL through an SNMP, console, SSH, or Telnet session.

Message Level	Informational
Message	Bridge is new root, vlan vlan-id , root ID root-id
Explanation	A Spanning Tree Protocol (STP) topology change has occurred, resulting in the Brocade device becoming the root bridge. The vlan-id is the ID of the VLAN in which the STP topology change occurred. The root-id is the STP bridge root ID.
Message Level	Informational
Message	Bridge root changed, vlan vlan-id , new root ID string , root interface portnum
Explanation	A Spanning Tree Protocol (STP) topology change has occurred. The vlan-id is the ID of the VLAN in which the STP topology change occurred. The root-id is the STP bridge root ID. The portnum is the number of the port connected to the new root bridge.
Message Level	Informational
Message	Bridge topology change, vlan vlan-id , interface portnum , changed state to stp-state
Explanation	A Spanning Tree Protocol (STP) topology change has occurred on a port. The vlan-id is the ID of the VLAN in which the STP topology change occurred. The portnum is the port number. The stp-state is the new STP state and can be one of the following: <ul style="list-style-type: none"> • disabled • blocking • listening • learning • forwarding • unknown
Message Level	Informational
Message	Cold start
Explanation	The device has been powered on.
Message Level	Informational
Message	DHCP: snooping on untrusted port portnum , type number, drop
Explanation	The device has indicated that the DHCP client receives DHCP server reply packets on untrusted ports, and packets are dropped.
Message Level	Informational
Message	DOT1X: port portnum - MAC mac address Cannot apply an ACL or MAC filter on a port member of a VE (virtual interface)
Explanation	The RADIUS server returned an IP ACL or MAC address filter, but the port is a member of a virtual interface (VE).
Message Level	Informational
Message	DOT1X: port portnum - MAC mac address cannot remove inbound ACL
Explanation	An error occurred while removing the inbound ACL.
Message Level	Informational

Message	DOT1X: port portnum - MAC mac address Downloading a MAC filter, but MAC filter have no effect on router port
Explanation	The RADIUS server returned an MAC address filter, but the portnum is a router port (it has one or more IP addresses).
Message Level	Informational
Message	DOT1X: port portnum - MAC mac address Downloading an IP ACL, but IP ACL have no effect on a switch port
Explanation	The RADIUS server returned an IP ACL, but the portnum is a switch port (no IP address).
Message Level	Informational
Message	DOT1X:port portnum - MAC mac address Error - could not add all MAC filters
Explanation	The Brocade device was unable to implement the MAC address filters returned by the RADIUS server.
Message Level	Informational
Message	DOT1X: port portnum - MAC mac address Invalid MAC filter ID - this ID doesn't exist
Explanation	The MAC address filter ID returned by the RADIUS server does not exist in the Brocade configuration.
Message Level	Informational
Message	DOT1X: port portnum - MAC mac address Invalid MAC filter ID - this ID is user defined and cannot be used
Explanation	The port was assigned a MAC address filter ID that had been dynamically created by another user.
Message Level	Informational
Message	DOT1X: port portnum - MAC mac address is unauthorized because system resource is not enough or the invalid information to set the dynamic assigned IP ACLs or MAC address filters
Explanation	802.1X authentication failed for the Client with the specified mac address on the specified portnum either due to insufficient system resources on the device, or due to invalid IP ACL or MAC address filter information returned by the RADIUS server.
Message Level	Informational
Message	DOT1X: port portnum - MAC mac address Port is already bound with MAC filter
Explanation	The RADIUS server returned a MAC address filter, but a MAC address filter had already been applied to the port.
Message Level	Informational
Message	DOT1X:port portnum - MAC mac address This device doesn't support ACL with MAC Filtering on the same port
Explanation	The RADIUS server returned a MAC address filter while an IP ACL was applied to the port, or returned an IP ACL while a MAC address filter was applied to the port.
Message Level	Informational
Message	DOT1X: Port portnum is unauthorized because system resource is not enough or the invalid information to set the dynamic assigned IP ACLs or MAC address filters
Explanation	802.1X authentication could not take place on the port. This happened because strict security mode was enabled and one of the following occurred: <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)
Message Level	Informational

Message	DOT1X: Port portnum currently used vlan-id changes to vlan-id due to dot1x-RADIUS vlan assignment
Explanation	A user has completed 802.1X authentication. The profile received from the RADIUS server specifies a VLAN ID for the user. The port to which the user is connected has been moved to the VLAN indicated by vlan-id .
Message Level	Informational
Message	DOT1X: Port portnum currently used vlan-id is set back to port default vlan-id vlan-id
Explanation	The user connected to portnum has disconnected, causing the port to be moved back into its default VLAN, vlan-id .
Message Level	Informational
Message	DOT1X: Port portnum , AuthControlledPortStatus change: authorized
Explanation	The status of the interface controlled port has changed from unauthorized to authorized.
Message Level	Informational
Message	DOT1X: Port portnum , AuthControlledPortStatus change: unauthorized
Explanation	The status of the interface controlled port has changed from authorized to unauthorized.
Message Level	Informational
Message	Enable super port-config read-only password deleted added modified from console telnet ssh snmp OR Line password deleted added modified from console telnet ssh snmp
Explanation	A user created, re-configured, or deleted an Enable or Line password through the SNMP, console, SSH, or Telnet session.
Message Level	Informational
Message	ERR_DISABLE: Interface ethernet portnum err-disable recovery timeout
Explanation	Errdisable recovery timer expired and the port has been reenabled.
Message Level	Informational
Message	ERR_DISABLE: Interface ethernet 16, err-disable recovery timeout
Explanation	If the wait time (port is down and is waiting to come up) expires and the port is brought up the following message is displayed.
Message Level	Informational
Message	ERR_DISABLE: Link flaps on port ethernet 16 exceeded threshold; port in err-disable state
Explanation	The threshold for the number of times that a port link toggles from "up" to "down" and "down" to "up" has been exceeded.
Message Level	Informational
Message	Interface portnum , line protocol down
Explanation	The line protocol on a port has gone down. The portnum is the port number.
Message Level	Informational
Message	Interface portnum , line protocol up
Explanation	The line protocol on a port has come up. The portnum is the port number.
Message Level	Informational
Message	Interface portnum , state down

Explanation	A port has gone down. The portnum is the port number.
Message Level	Informational
Message	Interface portnum , state up
Explanation	A port has come up. The portnum is the port number.
Message Level	Informational
Message	MAC Based Vlan Disabled on port port id
Explanation	A MAC Based VLAN has been disabled on a port
Message Level	Informational
Message	MAC Based Vlan Enabled on port port id
Explanation	A MAC Based VLAN has been enabled on a port.
Message Level	Informational
Message	MAC Filter added deleted modified from console telnet ssh snmp session filter id = MAC filter ID , src MAC = Source MAC address any, dst MAC = Destination MAC address any
Explanation	A user created, modified, deleted, or applied this MAC address filter through the SNMP, console, SSH, or Telnet session.
Message Level	Informational
Message	MSTP: BPDU-guard interface ethernet port-number detect (Received BPDU), putting into err-disable state.
Explanation	BPDU guard violation occurred in MSTP.
Message Level	Informational
Message	OPTICAL MONITORING: port port-number is not capable.
Explanation	The optical transceiver is qualified by Brocade, but the transceiver does not support digital optical performance monitoring.
Message Level	Informational
Message	Port p priority changed to n
Explanation	A port priority has changed.
Message Level	Informational
Message	Port portnum , srcip-security max-ipaddr-per-int reached.Last IP= ipaddr
Explanation	The address limit specified by the srcip-security max-ipaddr-per-interface command has been reached for the port.
Message Level	Informational
Message	Port portnum , srcip-security max-ipaddr-per-int reached.Last IP= ipaddr
Explanation	The address limit specified by the srcip-security max-ipaddr-per-interface command has been reached for the port.
Message Level	Informational
Message	Security: console login by username to USER PRIVILEGE EXEC mode
Explanation	The specified user logged into the device console into the specified EXEC mode.
Message Level	Informational
Message	Security: console logout by username
Explanation	The specified user logged out of the device console.

Message Level	Informational
Message	Security: telnet SSH login by username from src IP ip-address , src MAC mac-address to USER PRIVILEGE EXEC mode
Explanation	The specified user logged into the device using Telnet or SSH from either or both the specified IP address and MAC address. The user logged into the specified EXEC mode.
Message Level	Informational
Message	Security: telnet SSH logout by username from src IP ip-address, src MAC mac-address to USER PRIVILEGE EXEC mode
Explanation	The specified user logged out of the device. The user was using Telnet or SSH to access the device from either or both the specified IP address and MAC address. The user logged out of the specified EXEC mode.
Message Level	Informational
Message	SNMP read-only community read-write community contact location user group view engineid trap [host] [value -str] deleted added modified from console telnet ssh snmp session
Explanation	A user made SNMP configuration changes through the SNMP, console, SSH, or Telnet session. [value-str] does not appear in the message if SNMP community or engineid is specified.
Message Level	Informational
Message	SNMP Auth. failure, intruder IP: ip-addr
Explanation	A user has tried to open a management session with the device using an invalid SNMP community string. The ip-addr is the IP address of the host that sent the invalid community string.
Message Level	Informational
Message	SSH telnet server enabled disabled from console telnet ssh snmp session [by user username]
Explanation	A user enabled or disabled an SSH or Telnet session, or changed the SSH enable/disable configuration through the SNMP, console, SSH, or Telnet session.
Message Level	Informational
Message	startup-config was changed or startup-config was changed by user-name
Explanation	A configuration change was saved to the startup-config file. The user-name is the user ID, if they entered a user ID to log in.
Message Level	Informational
Message	STP: Root Guard Port port-number, VLAN vlan-ID consistent (Timeout).
Explanation	Root guard unblocks a port.
Message Level	Informational
Message	STP: Root Guard Port port-number , VLAN vlan-ID inconsistent (Received superior BPDU) .
Explanation	Root guard blocked a port.
Message Level	Informational
Message	STP: VLAN vlan id BPDU-Guard on Port port id triggered (Received BPDU), putting into err-disable state
Explanation	The BPDU guard feature has detected an incoming BPDU on {vlan-id, port-id}
Message Level	Informational
Message	STP: VLAN vlan id Root-Protect Port port id , Consistent (Timeout)
Explanation	The root protect feature goes back to the consistent state.

Message Level	Informational
Message	STP: VLAN vlan id Root-Protect Port port id , Inconsistent (Received superior BPDU)
Explanation	The root protect feature has detected a superior BPDU and goes into the inconsistent state on { vlan-id , port-id }.
Message Level	Informational
Message	STP: VLAN vlan-id BPDU-guard port port-number detect (Received BPDU), putting into err-disable state
Explanation	STP placed a port into an errdisable state for BPDU guard.
Message Level	Informational
Message	STP: VLAN 1 BPDU-guard port port-number detect (Received BPDU), putting into err-disable state.
Explanation	BPDU guard violation in occurred in STP or RSTP.
Message Level	Informational
Message	Syslog server IP-address deleted added modified from console telnet ssh snmp OR Syslog operation enabled disabled from console telnet ssh snmp
Explanation	A user made Syslog configuration changes to the specified Syslog server address, or enabled or disabled a Syslog operation through the SNMP, console, SSH, or Telnet session.
Message Level	Informational
Message	SYSTEM: Optic is not Brocade-qualified (port-number)
Explanation	Brocade does not support the optical transceiver.
Message Level	Informational
Message	System: Fan fan id (from left when facing right side), ok
Explanation	The fan status has changed from fail to normal.
Message Level	Informational
Message	System: Fan speed changed automatically to fan speed
Explanation	The system automatically changed the fan speed to the speed specified in this message.
Message Level	Informational
Message	System: No free TCAM entry. System will be unstable
Explanation	There are no TCAM entries available.
Message Level	Informational
Message	System: Static MAC entry with MAC Address mac-address is added from the unit / slot / port to unit / slot / port on VLANs vlan-id to vlan-id
Explanation	A MAC address is added to a range of interfaces, which are members of the specified VLAN range.
Message Level	Informational
Message	System: Static MAC entry with MAC Address mac-address is added to the unit / slot / port to unit / slot / port on vlan-id
Explanation	A MAC address is added to a range of interfaces, which are members of the specified VLAN.
Message Level	Informational
Message	System: Static MAC entry with MAC Address mac-address is added to portnumber unit / slot / port on VLAN vlan-id
Explanation	A MAC address is added to an interface and the interface is a member of the specified VLAN.
Message Level	Informational
Message	System: Static MAC entry with MAC Address mac-address is deleted from the unit/slot/ port to unit / slot / port on vlan-id
Explanation	A MAC address is deleted from a range of interfaces, which are members of the specified VLAN.

Message Level	Informational
Message	System: Static MAC entry with MAC Address mac-address is deleted from et he unit / slot / port to unit / slot / port on VLANs vlan-id to vlan-id
Explanation	A MAC address is deleted from a range of interfaces, which are members of the specified VLAN range.
Message Level	Informational
Message	System: Static MAC entry with MAC Address mac-address is deleted from portnumber unit / slot / port on vlan-id
Explanation	A MAC address is deleted from an interface and the interface is a member of the specified VLAN.
Message Level	Informational
Message	System: Static MAC entry with MAC Address mac-address is deleted from portnumber unit / slot / port on VLANs vlan-id to vlan-id
Explanation	A MAC address is deleted from an interface and the interface is a member of the specified VLAN range.
Message Level	Informational
Message	telnet SSH access [by username] from src IP source ip address , src MAC source MAC address rejected, n attempts
Explanation	There were failed SSH, or Telnet login access attempts from the specified source IP and MAC address. <ul style="list-style-type: none"> • [by user username] does not appear if telnet or SSH clients are specified. • n is the number of times this SNMP trap occurred in the last five minutes, or other configured number of minutes.
Message Level	Informational
Message	Trunk group (ports) created by 802.3ad link-aggregation module.
Explanation	802.3ad link aggregation is configured on the device, and the feature has dynamically created a trunk group (aggregate link). The ports variable is a list of the ports that were aggregated to make the trunk group.
Message Level	Informational
Message	user username added deleted modified from console telnet ssh snmp
Explanation	A user created, modified, or deleted a local user account through the SNMP, console, SSH, or Telnet session.
Message Level	Informational
Message	vlan vlan id added deleted modified from console telnet ssh snmp session
Explanation	A user created, modified, or deleted a VLAN through the SNMP, console, SSH, or Telnet session.
Message Level	Informational
Message	Warm start
Explanation	The system software (flash code) has been reloaded.
Message Level	Informational
Message	Stack: Stack unit unit# has been deleted to the stack system
Explanation	The specified unit has been deleted from the stacking system.
Message Level	Informational
Message	Stack unit unitNumber has been elected as ACTIVE unit of the stack system
Explanation	The specified unit in a stack has been elected as the Master unit for the stacking system.
Message Level	Informational
Message	Stack: Stack unit unit# has been added to the stack system
Explanation	The specified unit has been added to the stacking system.
Message Level	Informational

Message	System: Management MAC address changed to mac_address
Explanation	The management MAC address of a stacking system has been changed
Message Level	Informational
Message	System: Stack unit unit# Fan fan# (description), failed
Explanation	The operational status of a fan in the specified unit in a stack changed from normal to failure.
Message Level	Informational
Message	System: Stack unit unit# Power supply power-supply# is down
Explanation	The operational status of a power supply of the specified unit in a stack changed from normal to failure.
Message Level	Informational
Message	System: Stack unit unit# Power supply power-supply# is up
Explanation	The operational status of a power supply of the specified unit in a stack changed from failure to normal.
Message Level	Informational
Message	System: Stack unit unit# Fan fan# (description), ok
Explanation	The operational status of a fan in the specified unit in a stack changed from failure to normal.
Message Level	Informational
Message	System: Stack unit unitNumber Temperature actual-temp C degrees, warning level warning-temp C degrees, shutdown level shutdown-temp C degrees
Explanation	The actual temperature reading for a unit in a stack is above the warning temperature threshold.
Message Level	Informational
Message	vlan vlan-id Bridge is RootBridge mac-address (MgmtPriChg)
Explanation	802.1W changed the current bridge to be the root bridge of the given topology due to administrative change in bridge priority.
Message Level	Informational
Message	vlan vlan-id Bridge is RootBridge mac-address (MsgAgeExpiry)
Explanation	The message age expired on the Root port so 802.1W changed the current bridge to be the root bridge of the topology.
Message Level	Informational
Message	vlan vlan-id interface portnum Bridge TC Event (DOT1wTransition)
Explanation	802.1W recognized a topology change event in the bridge. The topology change event is the forwarding action that started on a non-edge Designated port or Root port.
Message Level	Informational
Message	vlan vlan-id interface portnum STP state - state (DOT1wTransition)
Explanation	802.1W changed the state of a port to a new state: forwarding, learning, blocking. If the port changes to blocking, the bridge port is in discarding state.
Message Level	Informational
Message	vlan vlan-id New RootBridge mac-address RootPort portnum (BpduRcvd)
Explanation	802.1W selected a new root bridge as a result of the BPDUs received on a bridge port.
Message Level	Informational
Message	vlan vlan-id New RootPort portnum (RootSelection)
Explanation	802.1W changed the port role to Root port, using the root selection computation.
Message Level	Informational
Message	ACL exceed max DMA L4 cam resource, using flow based ACL instead
Explanation	The port does not have enough Layer 4 CAM entries for the ACL.

To correct this condition, allocate more Layer 4 CAM entries. To allocate more Layer 4 CAM entries, enter the following command at the CLI configuration level for the interface:

ip access-group max-l4-cam num

Message Level Notification

Message ACL insufficient L4 cam resource, using flow based ACL instead

Explanation The port does not have a large enough CAM partition for the ACLs

Message Level Notification

Message ACL insufficient L4 session resource, using flow based ACL instead

Explanation The device does not have enough Layer 4 session entries.

To correct this condition, allocate more memory for sessions. To allocate more memory, enter the following command at the global CONFIG level of the CLI interface:

system-max session-limit num

Message Level Notification

Message ACL port fragment packet inspect rate rate exceeded on port portnum

Explanation The fragment rate allowed on an individual interface has been exceeded.

The *rate* indicates the maximum rate allowed.

The *portnum* indicates the port.

This message can occur if fragment throttling is enabled.

Message Level Notification

Message ACL system fragment packet inspect rate rate exceeded

Explanation The fragment rate allowed on the device has been exceeded.

The *rate* indicates the maximum rate allowed.

This message can occur if fragment throttling is enabled.

Message Level Notification

Message Authentication Disabled on portnum

Explanation The multi-device port authentication feature was disabled on the on the specified portnum .

Message Level Notification

Message Authentication Enabled on portnum

Explanation The multi-device port authentication feature was enabled on the on the specified portnum .

Message Level Notification

Message BGP Peer ip-addr DOWN (IDLE)

Explanation Indicates that a BGP4 neighbor has gone down.

The *ip-addr* is the IP address of the neighbor BGP4 interface with the Brocade device.

Message Level Notification

Message BGP Peer ip-addr UP (ESTABLISHED)

Explanation Indicates that a BGP4 neighbor has come up.

The *ip-addr* is the IP address of the neighbor BGP4 interface with the Brocade device.

Message Level Notification

Message DHCP: snooping on untrusted port portnum , type number, drop

Explanation Indicates that the DHCP client receives DHCP server reply packets on untrusted ports, and packets are dropped.

Message Level	Notification
Message	<code>DOT1X issues software but not physical port down indication of Port portnum to other software applications</code>
Explanation	The device has indicated that the specified is no longer authorized, but the actual port may still be active.
Message Level	Notification
Message	<code>DOT1X issues software but not physical port up indication of Port portnum to other software applications</code>
Explanation	The device has indicated that the specified port has been authenticated, but the actual port may not be active.
Message Level	Notification
Message	<code>DOT1X: Port port_id Mac mac_address -user user_id - RADIUS timeout for authentication</code>
Explanation	The RADIUS session has timed out for this 802.1x port.
Message Level	Notification
Message	<code>ISIS L1 ADJACENCY DOWN system-id on circuit circuit-id</code>
Explanation	The Layer 3 switch adjacency with this Level-1 IS-IS has gone down. The system- <i>id</i> is the system ID of the IS-IS. The circuit- <i>id</i> is the ID of the circuit over which the adjacency was established.
Message Level	Notification
Message	<code>ISIS L1 ADJACENCY UP system-id on circuit circuit-id</code>
Explanation	The Layer 3 switch adjacency with this Level-1 IS-IS has come up. The system- <i>id</i> is the system ID of the IS-IS. The circuit- <i>id</i> is the ID of the circuit over which the adjacency was established.
Message Level	Notification
Message	<code>ISIS L2 ADJACENCY DOWN system-id on circuit circuit-id</code>
Explanation	The Layer 3 switch adjacency with this Level-2 IS-IS has gone down. The system- <i>id</i> is the system ID of the IS-IS. The circuit- <i>id</i> is the ID of the circuit over which the adjacency was established.
Message Level	Notification
Message	<code>ISIS L2 ADJACENCY UP system-id on circuit circuit-id</code>
Explanation	The Layer 3 switch adjacency with this Level-2 IS-IS has come up. The system- <i>id</i> is the system ID of the IS-IS. The circuit- <i>id</i> is the ID of the circuit over which the adjacency was established.
Message Level	Notification
Message	<code>Local ICMP exceeds burst-max burst packets, stopping for lockup seconds!!</code>
Explanation	The number of ICMP packets exceeds the burst-max threshold set by the <code>ip icmp burst</code> command. The Brocade device may be the victim of a Denial of Service (DoS) attack. All ICMP packets will be dropped for the number of seconds specified by the lockup value. When the lockup period expires, the packet counter is reset and measurement is restarted.
Message Level	Notification
Message	<code>Local TCP exceeds burst-max burst packets, stopping for lockup seconds!!</code>

Explanation	<p>The number of TCP SYN packets exceeds the burst-max threshold set by the <code>ip tcp burst</code> command. The Brocade device may be the victim of a TCP SYN DoS attack.</p> <p>All TCP SYN packets will be dropped for the number of seconds specified by the <code>lockup p</code> value. When the lockup period expires, the packet counter is reset and measurement is restarted.</p>
Message Level	Notification
Message	<code>Local TCP exceeds num burst packets, stopping for num seconds!!</code>
Explanation	<p>Threshold parameters for local TCP traffic on the device have been configured, and the maximum burst size for TCP packets has been exceeded.</p> <p>The first num is the maximum burst size (maximum number of packets allowed).</p> <p>The second num is the number of seconds during which additional TCP packets will be blocked on the device.</p>
	NOTE
	This message can occur in response to an attempted TCP SYN attack.
Message Level	Notification
Message	<code>MAC Authentication RADIUS timeout for mac_address on port port_id</code>
Explanation	The RADIUS session has timed out for the MAC address for this port.
Message Level	Notification
Message	<code>MAC Authentication succeeded for mac-address on portnum</code>
Explanation	RADIUS authentication was successful for the specified mac-address on the specified portnum .
Message Level	Notification
Message	<code>Module was inserted to slot slot-num</code>
Explanation	<p>Indicates that a module was inserted into a chassis slot.</p> <p>The slot-num is the number of the chassis slot into which the module was inserted.</p>
Message Level	Notification
Message	<code>Module was removed from slot slot-num</code>
Explanation	<p>Indicates that a module was removed from a chassis slot.</p> <p>The slot-num is the number of the chassis slot from which the module was removed.</p>
Message Level	Notification
Message	<code>OSPF interface state changed,rid router-id , intf addr ip-addr , state ospf-state</code>
Explanation	<p>Indicates that the state of an OSPF interface has changed.</p> <p>The router-id is the router ID of the Brocade device.</p> <p>The ip-addr is the interface IP address.</p> <p>The ospf-state 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

Message Level	Notification
Message	OSPF intf authen failure, rid router-id , intf addr ip-addr , pkt src addr src-ip-addr , error type error-type , pkt type pkt-type
Explanation	<p>Indicates that an OSPF interface authentication failure has occurred.</p> <p>The <i>router-id</i> is the router ID of the Brocade device.</p> <p>The <i>ip-addr</i> is the IP address of the interface on the Brocade device.</p> <p>The <i>src-ip-addr</i> is the IP address of the interface from which the Brocade 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
Message Level	Notification
Message	OSPF intf config error, rid router-id , intf addr ip-addr , pkt src addr src-ip-addr , error type error-type , pkt type pkt-type
Explanation	<p>Indicates that an OSPF interface configuration error has occurred.</p> <p>The <i>router-id</i> is the router ID of the Brocade device.</p> <p>The <i>ip-addr</i> is the IP address of the interface on the Brocade device.</p> <p>The <i>src-ip-addr</i> is the IP address of the interface from which the Brocade 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

The packet-type can be one of the following:

- hello
- database description
- link state request
- link state update
- link state ack
- unknown

Message Level

Notification

Message

OSPF intf rcvd bad pkt, rid router-id , intf addr ip-addr , pkt src addr src-ip-addr, pkt type pkt-type

Explanation

Indicates that an OSPF interface received a bad packet.

The router-id is the router ID of the Brocade device.

The ip-addr is the IP address of the interface on the Brocade device.

The src-ip-addr is the IP address of the interface from which the Brocade device received the authentication failure.

The packet-type can be one of the following:

- hello
- database description
- link state request
- link state update
- link state ack
- unknown

Message Level

Notification

Message

OSPF intf rcvd bad pkt: Bad Checksum, rid ip-addr , intf addr ip-addr , pkt size num , checksum num , pkt src addr ip-addr , pkt type type

Explanation

The device received an OSPF packet that had an invalid checksum.

The rid ip-addr is the Brocade router ID.

The intf addr ip-addr is the IP address of the Brocade interface that received the packet.

The pkt size num is the number of bytes in the packet.

The checksum num is the checksum value for the packet.

The pkt src addr ip-addr is the IP address of the neighbor that sent the packet.

The pkt type type is the OSPF packet type and can be one of the following:

- hello
- database description
- link state request
- link state update
- link state acknowledgement
- unknown (indicates an invalid packet type)

Message Level

Notification

Message

OSPF intf rcvd bad pkt: Bad Packet type, rid ip-addr, intf addr ip-addr , pkt size num , checksum num , pkt src addr ip-addr , pkt type type

Explanation

The device received an OSPF packet with an invalid type.

The parameters are the same as for the Bad Checksum message. The pkt type type value is "unknown", indicating that the packet type is invalid.

Message Level

Notification

Message

OSPF intf rcvd bad pkt: Invalid packet size, rid ip-addr, intf addr ip-addr, pkt size num , checksum num , pkt src addr ip-addr , pkt type type

Explanation

The device received an OSPF packet with an invalid packet size.

The parameters are the same as for the Bad Checksum message.

Message Level

Notification

Message

OSPF intf rcvd bad pkt: Unable to find associated neighbor, rid ip-addr, intf addr ip-addr, pkt size num , checksum num , pkt src addr ip-addr , pkt type type

Explanation

The neighbor IP address in the packet is not in the list of OSPF neighbors in the Brocade device.

The parameters are the same as for the Bad Checksum message.

Message Level

Notification

Message

OSPF intf retransmit, rid router-id, intf addr ip-addr, nbr rid nbr-router-id , pkt type is pkt-type, LSA type lsa-type , LSA id lsa-id, LSA rid lsa-router-id

Explanation

An OSPF interface on the Brocade device has retransmitted a Link State Advertisement (LSA).

The router-id is the router ID of the Brocade device.

The ip-addr is the IP address of the interface on the Brocade device.

The nbr-router-id is the router ID of the neighbor router.

The packet-type can be one of the following:

- hello
- database description
- link state request
- link state update
- link state ack
- unknown

The lsa-type is the type of LSA.

The lsa-id is the LSA ID.

	The lsa-router-id is the LSA router ID.
Message Level	Notification
Message	OSPF LSDB approaching overflow, rid router-id , limit num
Explanation	The software is close to an LSDB condition. The router-id is the router ID of the Brocade device. The num is the number of LSAs.
Message Level	Notification
Message	OSPF LSDB overflow, rid router-id, limit num
Explanation	A Link State Database Overflow (LSDB) condition has occurred. The router-id is the router ID of the Brocade device. The num is the number of LSAs.
Message Level	Notification
Message	OSPF max age LSA, rid router-id , area area-id , LSA type lsa-type , LSA id lsa-id , LSA rid lsa-router-id
Explanation	An LSA has reached its maximum age. The router-id is the router ID of the Brocade device. The area-id is the OSPF area. The lsa-type is the type of LSA. The lsa-id is the LSA ID. The lsa-router-id is the LSA router ID.
Message Level	Notification
Message	OSPF nbr state changed, rid router-id , nbr addr ip-addr , nbr rid nbr-router-Id , state ospf-state
Explanation	Indicates that the state of an OSPF neighbor has changed. The router-id is the router ID of the Brocade device. The ip-addr is the IP address of the neighbor. The nbr-router-id is the router ID of the neighbor. The ospf-state indicates the state to which the interface has changed and can be one of the following:
	<ul style="list-style-type: none"> • down • attempt • initializing • 2-way • exchange start • exchange • loading • full • unknown
Message Level	Notification

Message	OSPF originate LSA, rid router-id , area area-id , LSA type lsa-type , LSA id lsa-id , LSA router id lsa-router-id
Explanation	An OSPF interface has originated an LSA. The router-id is the router ID of the Brocade device. The area-id is the OSPF area. The lsa-type is the type of LSA. The lsa-id is the LSA ID. The lsa-router-id is the LSA router ID.
Message Level	Notification
Message	OSPF virtual intf authen failure, rid router-id , intf addr ip-addr , pkt src addr src-ip-addr , error type error-type , pkt type pkt-type
Explanation	Indicates that an OSPF virtual routing interface authentication failure has occurred. The router-id is the router ID of the Brocade device. The ip-addr is the IP address of the interface on the Brocade device. The src-ip-addr is the IP address of the interface from which the Brocade device received the authentication failure. The error-type can be one of the following: <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 The packet-type 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
Message Level	Notification
Message	OSPF virtual intf config error, rid router-id , intf addr ip-addr , pkt src addr src-ip-addr , error type error-type , pkt type pkt-type
Explanation	Indicates that an OSPF virtual routing interface configuration error has occurred.

The router-id is the router ID of the Brocade device.

The ip-addr is the IP address of the interface on the Brocade device.

The src-ip-addr is the IP address of the interface from which the Brocade device received the error packet.

The error-type can be one of the following:

- bad version
- area mismatch
- unknown NBMA neighbor
- unknown virtual neighbor
- authentication type mismatch
- authentication failure
- network mask mismatch
- hello interval mismatch
- dead interval mismatch
- option mismatch
- unknown

The packet-type can be one of the following:

- hello
- database description
- link state request
- link state update
- link state ack
- unknown

Message Level

Notification

Message

OSPF virtual intf rcvd bad pkt, rid router-id , intf addr ip-addr , pkt src addr src-ip-addr , pkt type pkt-type

Explanation

Indicates that an OSPF interface received a bad packet.

The router-id is the router ID of the Brocade device.

The ip-addr is the IP address of the interface on the Brocade device.

The src-ip-addr is the IP address of the interface from which the Brocade device received the authentication failure.

The packet-type can be one of the following:

- hello
- database description
- link state request
- link state update
- link state ack
- unknown

Message Level

Notification

Message OSPF virtual intf retransmit, rid router-id , intf addr ip-addr , nbr rid nbr-router-id , pkt type is pkt-type , LSA type lsa-type , LSA id lsa-id , LSA rid lsa-router-id

Explanation An OSPF interface on the Brocade device has retransmitted a Link State Advertisement (LSA).

The router-id is the router ID of the Brocade device.

The ip-addr is the IP address of the interface on the Brocade device.

The nbr-router-id is the router ID of the neighbor router.

The packet-type can be one of the following:

- hello
- database description
- link state request
- link state update
- link state ack
- unknown

The lsa-type is the type of LSA.

The lsa-id is the LSA ID.

The lsa-router-id is the LSA router ID.

Message Level Notification

Message OSPF virtual intf state changed, rid router-id , area area-id , nbr ip-addr , state ospf-state

Explanation Indicates that the state of an OSPF virtual routing interface has changed.

The router-id is the router ID of the router the interface is on.

The area-id is the area the interface is in.

The ip-addr is the IP address of the OSPF neighbor.

The ospf-state indicates the state to which the interface has changed and can be one of the following:

- down
- loopback
- waiting
- point-to-point
- designated router
- backup designated router
- other designated router
- unknown

Message Level Notification

Message OSPF virtual nbr state changed, rid router-id , nbr addr ip-addr , nbr rid nbr-router-id , state ospf-state

Explanation Indicates that the state of an OSPF virtual neighbor has changed.

The router-id is the router ID of the Brocade device.

The ip-addr is the IP address of the neighbor.

The nbr-router-id is the router ID of the neighbor.

The ospf-state indicates the state to which the interface has changed and can be one of the following:

- down
- attempt
- initializing
- 2-way
- exchange start
- exchange
- loading
- full
- unknown

Message Level Notification

Message Transit ICMP in interface portnum exceeds num burst packets, stopping for num seconds!!

Explanation Threshold parameters for ICMP transit (through) traffic have been configured on an interface, and the maximum burst size for ICMP packets on the interface has been exceeded.

The portnum is the port number.

The first num is the maximum burst size (maximum number of packets allowed).

The second num is the number of seconds during which additional ICMP packets will be blocked on the interface.

NOTE

This message can occur in response to an attempted Smurf attack.

Message Level Notification

Message Transit TCP in interface portnum exceeds num burst packets, stopping for num seconds!

Explanation Threshold parameters for TCP transit (through) traffic have been configured on an interface, and the maximum burst size for TCP packets on the interface has been exceeded.

The portnum is the port number.

The first num is the maximum burst size (maximum number of packets allowed).

The second num is the number of seconds during which additional TCP packets will be blocked on the interface.

NOTE

This message can occur in response to an attempted TCP SYN attack.

Message Level Notification

Message VRRP intf state changed, intf portnum , vrid virtual-router-id , state vrrp-state
 VRRP (IPv6) intf state changed, intf portnum , vrid virtual-router-id , state vrrp-state

Explanation A state change has occurred in a Virtual Router Redundancy Protocol (VRRP) or VRRP-E IPv4 or IPv6 interface.

The portnum is the port or interface where VRRP or VRRP-E is configured.

The virtual-router-id is the virtual router ID (VRID) configured on the interface.

The vrrp-state can be one of the following:

- init
- master
- backup
- unknown

Message Level Notification

Message DOT1X security violation at port portnum , malicious MAC address detected: mac-address

Explanation A security violation was encountered at the specified port number.

Message Level Warning

Message Dup IP ip-addr detected, sent from MAC mac-addr interface portnum

Explanation Indicates that the Brocade device received a packet from another device on the network with an IP address that is also configured on the Brocade device.

The ip-addr is the duplicate IP address.

The mac-addr is the MAC address of the device with the duplicate IP address.

The portnum is the Brocade port that received the packet with the duplicate IP address. The address is the packet source IP address.

Message Level Warning

Message IGMP/MLD no hardware vidx, broadcast to the entire vlan. rated limited number

Explanation IGMP or MLD snooping has run out of hardware application VLANs. There are 4096 application VLANs per device. Traffic streams for snooping entries without an application VLAN are switched to the entire VLAN and to the CPU to be dropped. This message is rate-limited to appear a maximum of once every 10 minutes. The rate-limited number shows the number on non-printed warnings.

Message Level Warning

Message IGMP/MLD: vlanId(portId) is V1 but rcvd V2 from nbr ipAddr

Explanation Port has received a query with a MLD version that does not match the port MLD version. This message is rate-limited to appear a maximum of once every 10 hours.

Message Level Warning

Message Latched low RX Power | TX Power | TX Bias Current | Supply Voltage | Temperature warning alarm | warning, port port-number

Explanation The optical transceiver on the given port has risen above or fallen below the alarm or warning threshold.

Message Level Warning

Message list ACL-num denied ip-proto src-ip-addr (src-tcp / udp-port) (Ethernet portnum mac-addr) - dst-ip-addr (dst-tcp / udp-port), 1 event(s)

Explanation Indicates that an Access Control List (ACL) denied (dropped) packets.

The ACL-num indicates the ACL number. Numbers 1 - 99 indicate standard ACLs. Numbers 100 - 199 indicate extended ACLs.

The ip-proto indicates the IP protocol of the denied packets.

The src-ip-addr is the source IP address of the denied packets.

The src-tcp / udp-port is the source TCP or UDP port, if applicable, of the denied packets.

The portnum indicates the port number on which the packet was denied.

The mac-addr indicates the source MAC address of the denied packets.

The dst-ip-addr indicates the destination IP address of the denied packets.

The dst-tcp / udp-port indicates the destination TCP or UDP port number, if applicable, of the denied packets.

Message Level

Warning

Message

`MAC filter group denied packets on port portnum, src macaddr mac-addr , num packets`

Explanation

Indicates that a MAC address filtergroup configured on a port has denied packets.

The portnum is the port on which the packets were denied.

The mac-addr is the source MAC address of the denied packets.

The num indicates how many packets matching the values above were dropped during the five-minute interval represented by the log entry.

Message Level

Warning

Message

`multicast no software resource: resource-name , rate-limited number`

Explanation

IGMP or MLD snooping has run out of software resources. This message is rate-limited to appear a maximum of once every 10 minutes. The rate-limited number shows the number of non-printed warnings.

Message Level

Warning

Message

`No global IP! cannot send IGMP msg.`

Explanation

The device is configured for **ip multicast active** but there is no configured IP address and the device cannot send out IGMP queries.

Message Level

Warning

Message

`No of prefixes received from BGP peer ip-addr exceeds warning limit num`

Explanation

The Layer 3 switch has received more than the allowed percentage of prefixes from the neighbor.

The ip-addr is the IP address of the neighbor.

The num is the number of prefixes that matches the percentage you specified. For example, if you specified a threshold of 100 prefixes and 75 percent as the warning threshold, this message is generated if the Layer 3 switch receives a 76th prefix from the neighbor.

Message Level

Warning

Message

`rip filter list list-num direction V1 | V2 denied ip-addr , num packets`

Explanation

Indicates that a RIP route filter denied (dropped) packets.

The list-num is the ID of the filter list.

The direction indicates whether the filter was applied to incoming packets or outgoing packets. The value can be one of the following:

- in
- out

The V1 or V2 value specifies the RIP version (RIPv1 or RIPv2).

The ip-addr indicates the network number in the denied updates.

The num indicates how many packets matching the values above were dropped during the five-minute interval represented by the log entry.

Message Level

Warning

Message

`Temperature is over warning level.`

Explanation

The chassis temperature has risen above the warning level.

Message Level Warning

Power over Ethernet

- Power over Ethernet overview..... 311
- Enabling and disabling Power over Ethernet..... 322
- Disabling support for PoE legacy power-consuming devices..... 323
- Enabling the detection of PoE power requirements advertised through CDP..... 324
- Setting the maximum power level for a PoE power-consuming device..... 325
- Setting the power class for a PoE power-consuming device..... 326
- Setting the power budget for a PoE interface module on an FSX device..... 327
- Setting the inline power priority for a PoE port 327
- Resetting PoE parameters..... 329
- Displaying Power over Ethernet information..... 329
- Inline power on PoE LAG ports..... 340
- Decouple PoE and datalink operations on PoE ports..... 342

Power over Ethernet overview

This section provides an overview of the requirements for delivering power over the LAN as defined by the Institute of Electrical and Electronics Engineers Inc. (IEEE) in specifications 802.3af (PoE) and 802.3at (PoE+ and High PoE).

Brocade PoE devices provide Power over Ethernet, compliant with the standards described in the IEEE 802.3af specification for delivering inline power. Brocade devices are compliant with both the 802.3af and 802.3at specifications. The 802.3af specification defined the original standard for delivering power over existing network cabling infrastructure, enabling multicast-enabled full streaming audio and video applications for converged services, such as Voice over IP (VoIP), Wireless Local Area Access (WLAN) points, IP surveillance cameras, and other IP technology devices. The 802.3at specification expands the standards to support higher power levels for more demanding powered devices, such as video IP phones, pan-tilt-zoom cameras, and high-power outdoor antennas for wireless access points. Except where noted, this document uses the term PoE to refer to PoE, PoE+, and High PoE.

For a list of the FastIron devices and modules that support PoE, PoE+, High PoE, Power over HDBaseT (PoH), or a combination, refer to the *FastIron Ethernet Switch Feature and Standards Support Matrix*.

PoE technology eliminates the need for an electrical outlet and dedicated UPS near IP powered devices. With power-sourcing equipment such as a BrocadeFastIron PoE device, power is consolidated and centralized in wiring closets, improving the reliability and resilience of the network.

Power over Ethernet terms used in this chapter

The following terms are introduced in this chapter:

- **High PoE** - Covered by IEEE 802.3at 2009, provides up to 60 Watts of power.
- **IP powered device (PD) or power-consuming device** - The Ethernet device that requires power. It is situated on the end of the cable opposite the power-sourcing equipment.
- **PoE+** - Covered by IEEE 802.at, provides up to 25.5 Watts of power.
- **PoH** - Covered by IEEE 802.3at 2009 and sometimes called power over HDBaseT, provides up to 95 Watts of power to power-consuming devices.
- **Power-sourcing device or Power-sourcing equipment (PSE)** - The source of the power, or the device that integrates the power onto the network. Power sourcing devices and equipment have embedded PoE technology. The Brocade FastIron PoE device is a power sourcing device.

Methods for delivering Power over Ethernet

There are two methods for delivering Power over Ethernet (PoE) as defined in the 802.3af and 802.3at specifications:

- **Endspan** - Power is supplied through the Ethernet ports on a power-sourcing device. With the Endspan solution, power can be carried over the two data pairs (Alternative A) or the two spare pairs (Alternative B).
- **Midspan** - Power is supplied by an intermediate power-sourcing device placed between the switch and the PD. With the Midspan solution, power is carried over the two spare pairs (Alternative B).

With both methods, power is transferred over four conductors, between the two pairs. 802.3af- and 802.3at-compliant PDs are able to accept power from either set of pairs.

Brocade PoE devices use the Endspan method, compliant with the 802.3af and 802.3at standards.

The Endspan and Midspan methods are described in more detail in the following sections.

NOTE


All 802.3af- and 802.3at-compliant power-consuming devices are required to support both application methods defined in the 802.3af and 802.3at specification.

PoE endspan method

The PoE Endspan method uses the Ethernet switch ports on power-sourcing equipment, such as a BrocadeFastIron PoE switch, which has embedded PoE technology to deliver power over the network.

With the Endspan solution, there are two supported methods of delivering power. In Alternative A, four wires deliver data and power over the network. Specifically, power is carried over the live wire pairs that deliver data as illustrated in the following figure. In Alternative B, the four wires of the spare pairs are used to deliver power over the network. Brocade PoE devices support Alternative A.

The Endspan method is shown in the following illustration.




PoE midspan method

The PoE Midspan method uses an intermediate device, usually another PD, to inject power into the network. The intermediate device is positioned between the switch and the PD and delivers power over the network using the spare pairs of wires (Alternative B). The intermediate device has multiple channels (typically 6 to 24), and each of the channels has data input and a data-plus-power RJ-45 output connector.

The Midspan method is illustrated in the following figure.

PoE Midspan Delivery Method



PoE autodiscovery

PoE autodiscovery is a detection mechanism that identifies whether an installed device is 802.3af- or 802.3at-compatible. When you plug a device into an Ethernet port that is capable of providing inline power, the autodiscovery mechanism detects whether the device requires power and how much power is needed. The autodiscovery mechanism also has a disconnect protection mechanism that shuts down the power once a PD has been disconnected from the network or when a faulty PD has been detected. This feature enables safe installation and prevents high-voltage damage to equipment.

PoE autodiscovery is achieved by periodically transmitting current or test voltages that can detect when a PD is attached to the network. When an 802.3af- or 802.3at-compatible device is plugged into a PoE, PoE+, or PoH port, the PD reflects test voltage back to the power-sourcing device (the Brocade device), ultimately causing the power to be switched on. Devices not compatible with 802.3af do not reflect test voltage back to the power-sourcing device.

Power class

A power class determines the amount of power a PD receives from power-sourcing equipment. When a valid PD is detected, the Brocade PoE device performs power classification by inducing a specific voltage and measuring the current consumption of the PD. Depending on the measured current, the appropriate class is assigned to the PD. PDs that do not support classification are assigned a class of 0 (zero). The following table shows the different power classes and their respective power consumption needs.

TABLE 47 Power classes for PDs

Class	Usage	Power (watts) from Power-Sourcing Device			
		Standard PoE	PoE+	High PoE	Power over HDBaseT (PoH)
0	default	15.4	15.4	15.4	15.4
1	optional	4	4	4	4
2	optional	7	7	7	7
3	optional	15.4	15.4	15.4	15.4
4	optional	N/A	30 ¹	60 ^{1 2}	95

Power specifications

The 802.3af (PoE) standard limits power to 15.4 watts (44 to 50 volts) from the power-sourcing device, in compliance with safety standards and existing wiring limitations. Though limited by the 802.3af standard, 15.4 watts of power was ample for most PDs, which consumed an average of 5 to 12 watts of power (IP phones, wireless LAN access points, and network surveillance cameras each consume an average of 3.5 to 9 watts of power). The 802.3at 2008 (PoE+) standard nearly doubles the power, providing 30 watts (52 to 55 volts) from the power-sourcing device. The 802.3at 2009 (High PoE) standard increases available power again, to 60 watts for High PoE and 95 watts for Power over HDBase-T (PoH).

NOTE

PoH ports on Brocade devices allocate 95 watts for PoE+, High PoE, and PoH PDs.

The PoE power supply provides power to the PoE circuitry block and ultimately to PoE power-consuming devices. The number of PoE power-consuming devices that one PoE power supply can support depends on the number of watts required by each power-consuming

¹ First eight ports of Brocade ICX 7450-24P or ICX 7450-48P supply 95w unless PD negotiates lower power requirement through LLDP protocol messages.

² Maximum power required for High PoE is 60 watts.

device and the capacity of the power supply or power supplies. Each PoE or PoE+ port supports a maximum of 15.4 or 30 watts of power per power-consuming device. Each PoH port supports a maximum of 95 watts of power (lower wattage can be negotiated through LLDP messages).

As an example, if each PoE power-consuming device attached to a FastIron PoE device is budgeted to consume 30 watts of power, one 720- or 748-watt power supply can power up to 24 PoE ports. With the exception of the ICX6430-C12 and the ICX6450-C12-PD, FastIron platforms support either a second power supply or an external power supply (EPS) to augment PoE power budget, depending on the product. Refer to the power supply specifications in the Brocade FastIron hardware installation guide for the appropriate FastIron device.

By default, a FastIron device pre-allocates power of 15.4 for a physically operational PoE configured port, 30 watts for a PoE+ configured port, and 95 watts for a PoH port. However, in an ICX 6450-C12 device that is operational without a direct power supply and has pass-through power, there is no pre-allocation. Instead, power is allocated only when a powered device is connected to the port. By default, the amount of power allocated depends on the power class of the powered device.

Dynamic upgrade of PoE power supplies

NOTE

This section applies to the FSX 800 and FSX 1600 chassis with PoE power supplies.

PoE+ requires higher power levels than standard PoE. In a chassis running software release 07.2.00 or higher, POE power supplies (SX-ACPWR-POE) are upgraded dynamically to 52 or 54 volts, depending on the maximum operating voltage the power supplies are capable of. The preferred voltage mode for PoE+ is 54 volts.

For safety reasons, all PoE power supplies installed in the chassis must operate at the same voltage mode, either 52 volts or 54 volts. The system selects the voltage mode of the power supply with the lowest supported voltage as the voltage mode for all PoE power supplies installed in the chassis. For example, in an FSX 800 chassis with one 52-volt capable PoE power supply and one 54-volt capable PoE power supply, both power supplies are configured dynamically to operate at 52 volts.

PoE+ voltage selection occurs during each of the following events:

- When the device is powered ON or is rebooted
- When a PoE power supply is installed in the chassis
- When a PoE power supply is removed from the chassis

These events are described in detail in the following sections.

NOTE

A PoE power supply upgrade does not persist beyond a single power cycle. An upgrade occurs automatically each time a power supply is re-inserted in the chassis.

You can use the **show inline power detail** command to display detailed information about the PoE power supplies installed in a FastIron PoE device. For more information, refer to section [Displaying detailed information about PoE power supplies](#) on page 334.



CAUTION

The SX-POE-AC-PWR power supply is designed exclusively for use with the BrocadeFSX PoE devices. The power supply produces extensive power to support 802.3af and 802.3at applications. Installing the power supply in a device other than the BrocadeFSX PoE device will cause extensive damage to your equipment.

Voltage selection during bootup

During bootup, the system selects the voltage mode (either 52 volts or 54 volts) of the power supply with the lowest supported voltage as the voltage mode for all PoE power supplies installed in the chassis. For example, if there is at least one power supply that supports 52 volts maximum, then all power supplies are configured to operate at 52 volts, even if other supplies are 54 volts-capable. Once the operating voltage is applied, the system displays and logs a warning message similar to the following:

```
device(config)#
Power supply 1 (from left when facing front side) detected.
Power supply 1 (from left when facing front side) is up.
WARNING: PoE power supplies in slots 1 are down rev. PoE/PoE+ function will work,
but output power may be less than 50V under worst case load.
```

If all power supplies are 54 volts-capable, then all power supplies are configured to operate at 54 volts. In this case, the system does not display or log a warning message.

Voltage selection when a PoE power supply is installed

When a PoE power supply is hotswapped into the chassis, the system automatically adjusts the voltage to match that of the PoE power supply or supplies that are currently installed in the chassis.

The following examples describe how the voltage is selected when a PoE power supply is installed:

- If a 54 volt-capable power supply is installed in a chassis that is operating with 52 volt-capable power supplies, the newly installed power supply is set to operate at 52 volts.
- If a 54 volt-capable power supply is installed in a chassis that is operating with 54 volt-capable power supplies, the newly installed power supply is set to operate at 54 volts.
- If a 52 volt-capable power supply is installed in a chassis that is operating with 54 volt-capable power supplies that are actively providing power, the system rejects the newly installed power supply since it cannot safely operate with the 54 volt-capable power supplies. In this case, the 52-volt power supply is powered OFF, and an error message similar to the following is displayed on the console.

```
device(config)#
Power supply 1 (from left when facing front side) detected.
Power supply 1 (from left when facing front side) is up.
Shutting down power supply in slot 1 because it is not compatible with the existing PoE power supplies.
Please remove and replace.
```

When the system is next reloaded, the power supply voltage will be selected as described in the section [Voltage selection during bootup](#) on page 316.

- If a 52 volt-capable power supply is installed in a chassis that is operating with 54 volt-capable power supplies that are *not* actively providing power, the system configures the power supplies to operate at 52 volts. In this case, the newly installed 52-volt power supply is *not* powered OFF, and a message similar to the following is displayed on the console.

NOTE: Automatically downgraded all PoE power supplies to 52V.

Voltage selection when a PoE power supply is removed

If a 52 volt PoE power supply is removed from the chassis, the system surveys the remaining power supplies to determine if they are 54 volts-capable. If the remaining supplies are 54 volts-capable and the system is not currently providing power to any PDs, the software upgrades the voltage of all supplies to 54 volts. The system displays and log a message similar to the following:

NOTE: Automatically upgraded all PoE power supplies to 54V.

However, if the system is currently providing power to one or more PDs, the system does not upgrade the voltage level. When the system is next reloaded, the power supply voltage is selected as described in the section [Voltage selection during bootup](#) on page 316.

Power over Ethernet cabling requirements

The 802.3af and 802.3at standards currently support PoE and PoE+ on 10/100/1000-Mbps Ethernet ports operating over standard Category 5 unshielded twisted pair (UTP) cable or better. If your network uses cabling categories less than Category 5, you cannot implement PoE without first upgrading your cables to Category 5 UTP cable or better. PoE has the following cabling requirements based on distance:

- Cat 5e - 25 meters
- Cat 6/6a - 55 meters
- Cat 7 - 100 meters.

Supported powered devices

Brocade PoE devices support a wide range of IP powered devices, including the following:

- Voice over IP (VoIP) phones
- Wireless LAN access points
- IP surveillance cameras

The following sections briefly describe these IP powered devices.

VoIP

Voice over IP (VoIP) is the convergence of traditional telephony networks with data networks. VoIP uses the existing data network infrastructure as the transport system for both services. Voice is traditionally transported on a network that uses circuit-switching technology, but data networks are built on packet-switching technology. To achieve this convergence, technology has been developed to take a voice signal, which originates as an analog signal, and transport it within a digital medium. This is done by devices such as VoIP telephones that receive the originating tones and place them in UDP packets. The size and frequency of these UDP packets depends on the coding / decoding (CODEC) technology that has been implemented in the VoIP telephone or device. The VoIP control packets use TCP/IP format.

IP surveillance cameras

IP surveillance technology provides digital streaming of video over Ethernet, providing real-time, remote access to video feeds from cameras.

The main benefit of using IP surveillance cameras on the network is that you can view surveillance images from any computer on the network. If you have access to the Internet, you can securely connect from anywhere in the world to view a chosen facility or even a single camera from your surveillance system. By using a Virtual Private Network (VPN) or the company intranet, you can manage password-protected access to images from the surveillance system. Similar to secure payment over the Internet, images and information are kept secure and can be viewed only by approved personnel.

Installing PoE firmware

PoE Firmware download can be initiated on one stack unit at a time on the FSX and FCX devices. On ICX 7250 and ICX 7450 devices, PoE Firmware download can be initiated on one stack unit at a time or on all PoE units or multiple stacks simultaneously. You can initiate Firmware download on different stack units even if Firmware download is in progress on other units. This reduces the time for Firmware upgrade and makes the process more efficient. PoE Firmware download can be initiated only from active units. However, switchover cannot be initiated when Firmware download is in progress. Firmware download status is reset or aborted after 20 minutes (non-configurable) from the time Firmware download is initiated. Firmware download initiated on a unit will be aborted if it is not completed within 20 minutes, making provisions to trigger switchover or initiate Firmware download again for the unit.

Firmware image file types

The following table lists PoE firmware files. The firmware files are specific to each device and cannot be used in any other device.

TABLE 48 PoE Firmware files

Product	PoE Firmware
FSX Gen 1 & 2 modules	fsx_poe_06.0.6.fw
FSX Gen 3 modules	fsx_poeplus_02.1.0.fw
FCX	fcx_poeplus_02.1.0.b004.fw
ICX 64xx	icx64xx_poeplus_02.1.0.b004.fw
ICX 6450-C12	icx64xxc12_poeplus_02.03.09.fw
ICX 6610	fcx_poeplus_02.1.0.b004.fw
ICX 7250	icx72xx_poeplus_01.8.8.b001.fw
ICX 7450	icx74xx_poh_01.8.8.b001.fw

Installing PoE firmware with TFTP

PoE firmware is stored in the PoE controller of the FastIron switch. You can install PoE firmware from the TFTP server on a FastIron switch using CLI commands. To do so, you should have a valid firmware image on the TFTP server.

NOTE

The PoE firmware upgrade feature is not supported in FIPS mode on Brocade devices.

NOTE

The CLI syntax to install PoE firmware is different on FSX, FCX, and ICX platforms.

NOTE

Installation of PoE firmware interrupts PoE services on the individual device or module as it is upgraded. PoE service restarts once PoE firmware installation is complete.

1. Place the PoE firmware on a TFTP server to which the Brocade device has access.

- Copy the PoE firmware from the TFTP server into the switch as shown in the following examples. Be sure to use the correct file image for the platform.

To install PoE firmware on FCX and ICX platforms, use the **inline power install-firmware stack-unit** command as shown in the following example.

```
device# inline power install-firmware stack-unit 1 tftp 10.120.54.161 icx74xx_poh_01.8.8.b001.fw
```

The process of PoE installation begins. You should see output similar to the following.

```
Family_Stack# Flash Memory Write (8192 bytes per dot) .....
tftp download successful stackId = 3 file name = poe-fw
Sending PoE Firmware to Stack Unit 3.
Flash Memory Write (8192 bytes per dot) .....
PoE: Power disabled on port 3/1/1 because of power management.
PoE: Power disabled on port 3/1/2 because of power management.
PoE: Power disabled on port 3/1/3 because of power management.
PoE: Power disabled on port 3/1/4 because of power management.
PoE: Power disabled on port 3/1/5 because of power management.
PoE: Power disabled on port 3/1/6 because of power management.
PoE: Power disabled on port 3/1/7 because of power management.
PoE: Power disabled on port 3/1/8 because of power management.
PoE: Power disabled on port 3/1/9 because of power management.
PoE: Power disabled on port 3/1/10 because of power management.
PoE: Power disabled on port 3/1/11 because of power management.
PoE: Power disabled on port 3/1/12 because of power management.
PoE: Power disabled on port 3/1/13 because of power management.
PoE: Power disabled on port 3/1/14 because of power management.
PoE: Power disabled on port 3/1/15 because of power management.
PoE: Power disabled on port 3/1/16 because of power management.
PoE: Power disabled on port 3/1/17 because of power management.
PoE: Power disabled on port 3/1/18 because of power management.
PoE: Power disabled on port 3/1/19 because of power management.
PoE: Power disabled on port 3/1/20 because of power management.
PoE: Power disabled on port 3/1/21 because of power management.
PoE: Power disabled on port 3/1/22 because of power management.
PoE: Power disabled on port 3/1/23 because of power management.
PoE: Power disabled on port 3/1/24 because of power management.
U3-MSG: PoE Warning: Upgrading firmware in slot 1...DO NOT HOTSWAP OR POWER DOWN THE MODULE.
U3-MSG: PoE Info: FW Download on slot 1...sending download command...
U3-MSG: PoE Info: FW Download on slot 1...TPE response received.
U3-MSG: PoE Info: FW Download on slot 1...sending erase command...
U3-MSG: PoE Info: FW Download on slot 1...erase command...accepted.
U3-MSG: PoE Info: FW Download on slot 1...erasing firmware memory...
U3-MSG: PoE Info: FW Download on slot 1...erasing firmware memory...completed
U3-MSG: PoE Info: FW Download on slot 1...sending program command...
U3-MSG: PoE Info: FW Download on slot 1...sending program command...accepted.
U3-MSG: PoE Info: FW Download on slot 1...programming firmware...takes around 12 minutes....
U3-MSG: PoE Info: Firmware Download on slot 1....10 percent completed.
U3-MSG: PoE Info: Firmware Download on slot 1....20 percent completed.
U3-MSG: PoE Info: Firmware Download on slot 1....30 percent completed.
U3-MSG: PoE Info: Firmware Download on slot 1....40 percent completed.
U3-MSG: PoE Info: Firmware Download on slot 1....50 percent completed.
U3-MSG: PoE Info: Firmware Download on slot 1....60 percent completed.
U3-MSG: PoE Info: Firmware Download on slot 1....70 percent completed.
U3-MSG: PoE Info: Firmware Download on slot 1....80 percent completed.
U3-MSG: PoE Info: Firmware Download on slot 1....90 percent completed.
U3-MSG: PoE Info: Firmware Download on slot 1....100 percent completed.
U3-MSG: PoE Info: FW Download on slot 1...programming firmware...completed.
U3-MSG: PoE Info: FW Download on slot 1...upgrading firmware...completed. Module will be reset.
U3-MSG: PoE Info: Resetting module in slot 1....completed.
PoE: Failed power allocation of 30000 mwatts on port 3/1/13. Will retry when more power budget.
PoE: Failed power allocation of 30000 mwatts on port 3/1/14. Will retry when more power budget.
PoE: Failed power allocation of 30000 mwatts on port 3/1/15. Will retry when more power budget.
PoE: Failed power allocation of 30000 mwatts on port 3/1/16. Will retry when more power budget.
PoE: Failed power allocation of 30000 mwatts on port 3/1/17. Will retry when more power budget.
PoE: Failed power allocation of 30000 mwatts on port 3/1/18. Will retry when more power budget.
PoE: Failed power allocation of 30000 mwatts on port 3/1/19. Will retry when more power budget.
PoE: Failed power allocation of 30000 mwatts on port 3/1/20. Will retry when more power budget.
```

```
PoE: Failed power allocation of 30000 mwatts on port 3/1/21. Will retry when more power budget.
PoE: Failed power allocation of 30000 mwatts on port 3/1/22. Will retry when more power budget.
PoE: Failed power allocation of 30000 mwatts on port 3/1/23. Will retry when more power budget.
PoE: Failed power allocation of 30000 mwatts on port 3/1/24. Will retry when more power budget.
```

3. After the firmware is downloaded into the controller, the controller resets and reboots with the new PoE firmware, You should see output similar to the following.

```
[MEMBER]local-3@ICX7450-24P Router>Download request from active unit 1 mac = 748e.f8dc.b39c
Downloading - poe.fw
Done.
PoE Info: Programming Brocade defaults.....
PoE Info: Programming Brocade defaults. Step 1: Writing port defaults on module in slot 1....
PoE Info: Programming Brocade Defaults: Step 2: Writing PM defaults on module in slot 1.
PoE Info: Programming Brocade defaults. Step 3: Writing user byte 0xf0 on module in slot 1.
PoE Info: Programming Brocade defaults. Step 4: Saving settings on module in slot 1.
PoE Info: Programming Brocade defaults....completed.

[MEMBER]local-3@ICX7450-24P Router>
```

NOTE

If you are attempting to transfer a file using TFTP but have received an error message, refer to [Firmware image file types](#) on page 318.

Upgrading the PoE firmware file using SCP

To use the PoE feature, download the PoE firmware file. You can then install it using SCP as shown in the following procedure.

NOTE

In a stack, you must install the PoE firmware on each individual member unit.

1. Place the PoE firmware file on an SCP-enabled host to which the Brocade device has access.

- Copy the PoE firmware file from the SCP-enabled host into the switch by entering the following command on the SCP-enabled host.

For FCX, ICX 6430, ICX 6450, ICX 6610, ICX 7250, and ICX 7450 devices:

```
pscp firmware hostname@management-ip:firmware:stackid:stack-id
```

For FSX devices:

```
pscp firmware hostname@management-ip:firmware:moduleid:module-id
```

For example:

```
C:/>pscp fsx_poe_07400.fw host1@10.10.1.1:firmware:stackid:1
```

The process of PoE firmware installation begins. In the FastIron device CLI, you should see output similar to the following.

```
Brocade(config)# scp download successful stackId = 1 file name = poe-fw
Sending PoE Firmware to Stack Unit 1.
PoE Warning: Upgrading firmware in slot 1....DO NOT SWITCH OVER OR POWER DOWN
THE UNIT.
PoE Info: FW Download on slot 1...sending download command...
PoE Info: FW Download on slot 1...TPE response received.
PoE Info: FW Download on slot 1...sending erase command...
PoE Info: FW Download on slot 1...erase command...accepted.
PoE Info: FW Download on slot 1...erasing firmware memory...
PoE Info: FW Download on slot 1...erasing firmware memory...completed
PoE Info: FW Download on slot 1...sending program command...
PoE Info: FW Download on slot 1...sending program command...accepted.
PoE Info: FW Download on slot 1...programming firmware...takes around 6
minutes....
Brocade(config)# U1-MSG: PoE Info: Firmware Download on slot 1.....10 percent
completed.

U1-MSG: PoE Info: Firmware Download on slot 1.....20 percent completed.
U1-MSG: PoE Info: Firmware Download on slot 1.....30 percent completed.
U1-MSG: PoE Info: Firmware Download on slot 1.....40 percent completed.

U1-MSG: PoE Info: Firmware Download on slot 1.....50 percent completed.
U1-MSG: PoE Info: Firmware Download on slot 1.....60 percent completed.
U1-MSG: PoE Info: Firmware Download on slot 1.....70 percent completed.
U1-MSG: PoE Info: Firmware Download on slot 1.....80 percent completed.
U1-MSG: PoE Info: Firmware Download on slot 1.....90 percent completed.
U1-MSG: PoE Info: Firmware Download on slot 1.....100 percent completed.
PoE Info: FW Download on slot 1...programming firmware...completed.
PoE Info: FW Download on slot 1...upgrading firmware...completed. Module will
be reset.
```

- After the firmware file is loaded into the device, the device resets and reboots with the new PoE firmware. You should see output similar to the following.

```
PoE Info: Resetting in slot 1....
PoE Info: Resetting module in slot 1....completed.
PoE Info: Programming Brocade defaults.....
PoE Info: Programming Brocade defaults. Step 1: Writing port defaults on
module in slot 1....
PoE Info: Programming Brocade Defaults: Step 2: Writing PM defaults on module
in slot 1.
PoE Info: Programming Brocade defaults. Step 3: Writing user byte 0xf0 on
module in slot 1.
PoE Info: Programming Brocade defaults. Step 4: Saving settings on module in
slot 1.
PoE Info: Programming Brocade defaults....completed
```

PoE and CPU utilization

Depending on the number of PoE-configured ports that have active power devices, there may be a slight and noticeable increase of up to 15 percent in CPU utilization. This is normal behavior for PoE and in typical scenarios does not affect the functionality of other features on the switch.

Enabling and disabling Power over Ethernet

To enable a port to receive inline power for power-consuming devices, use the **inline power** command for the appropriate port. Here is an example.

```
device# configure terminal
device(config)# interface ethernet 1/1/1
device(config-if-e1000-1/1/1)# inline power
```

Once you have entered the commands to enable inline power, the console displays the following message.

```
device(config-if-e1000-1/1/1)# PoE Info: Power enabled on port 1/1/1.
```

The following example disables inline power on a range of ports.

```
ICX7250-48P Router# configure terminal
ICX7250-48P Router(config)# interface ethernet 1/1/1 to 1/1/48
ICX7250-48P Router(config-mif-1/1/1-1/1/48)# no inline power
PoE: Power disabled on port 1/1/1 because of admin off.
PoE: Power disabled on port 1/1/2 because of admin off.
PoE: Power disabled on port 1/1/3 because of admin off.
PoE: Power disabled on port 1/1/4 because of admin off.
PoE: Power disabled on port 1/1/5 because of admin off.
PoE: Power disabled on port 1/1/6 because of admin off.
PoE: Power disabled on port 1/1/7 because of admin off.
PoE: Power disabled on port 1/1/8 because of admin off.
PoE: Power disabled on port 1/1/9 because of admin off.
PoE: Power disabled on port 1/1/10 because of admin off.
PoE: Power disabled on port 1/1/11 because of admin off.
PoE: Power disabled on port 1/1/12 because of admin off.
PoE: Power disabled on port 1/1/13 because of admin off.
PoE: Power disabled on port 1/1/14 because of admin off.
PoE: Power disabled on port 1/1/15 because of admin off.
PoE: Power disabled on port 1/1/16 because of admin off.
PoE: Power disabled on port 1/1/17 because of admin off.
PoE: Power disabled on port 1/1/18 because of admin off.
PoE: Power disabled on port 1/1/19 because of admin off.
PoE: Power disabled on port 1/1/20 because of admin off.
PoE: Power disabled on port 1/1/21 because of admin off.
PoE: Power disabled on port 1/1/22 because of admin off.
PoE: Power disabled on port 1/1/23 because of admin off.
PoE: Power disabled on port 1/1/24 because of admin off.
PoE: Power disabled on port 1/1/25 because of admin off.
PoE: Power disabled on port 1/1/26 because of admin off.
PoE: Power disabled on port 1/1/27 because of admin off.
PoE: Power disabled on port 1/1/28 because of admin off.
PoE: Power disabled on port 1/1/29 because of admin off.
PoE: Power disabled on port 1/1/30 because of admin off.
PoE: Power disabled on port 1/1/31 because of admin off.
PoE: Power disabled on port 1/1/32 because of admin off.
PoE: Power disabled on port 1/1/33 because of admin off.
PoE: Power disabled on port 1/1/34 because of admin off.
PoE: Power disabled on port 1/1/35 because of admin off.
PoE: Power disabled on port 1/1/36 because of admin off.
PoE: Power disabled on port 1/1/37 because of admin off.
PoE: Power disabled on port 1/1/38 because of admin off.
PoE: Power disabled on port 1/1/39 because of admin off.
PoE: Power disabled on port 1/1/40 because of admin off.
PoE: Power disabled on port 1/1/41 because of admin off.
```

```
PoE: Power disabled on port 1/1/42 because of admin off.
PoE: Power disabled on port 1/1/43 because of admin off.
PoE: Power disabled on port 1/1/44 because of admin off.
PoE: Power disabled on port 1/1/45 because of admin off.
PoE: Power disabled on port 1/1/46 because of admin off.
PoE: Power disabled on port 1/1/47 because of admin off.
PoE: Power disabled on port 1/1/48 because of admin off.
```

```
ICX7250-48P Router(config-mif-1/1/1-1/1/48) #
```

Syntax: `[no] inline power [device/slot/port] [to device/slot/port]`

Use the **no** form of the command to disable the port from receiving inline power.

NOTE

Inline power should not be configured between two switches, as it may cause unexpected behavior.

NOTE

FastIron PoE and PoE+ devices can automatically detect whether a power-consuming device is 802.3af- or 802.3at-compliant.

Disabling support for PoE legacy power-consuming devices

Brocade PoE devices automatically support most legacy power-consuming devices (devices not compliant with 802.3af 802.3at), as well as all 802.3af- and 802.3at-compliant devices. If desired, you can disable and re-enable support for legacy PoE power-consuming devices on a global basis (on the entire device) or on individual slots (FSX 800 and FSX 1600 chassis devices only). When you disable legacy support, 802.3af- and 802.3at-compliant devices are not affected.

To disable support for legacy power-consuming devices on a non-stackable device, enter the following command at the global CONFIG level of the CLI.

```
device(config) # no legacy-inline-power
```

To disable support for legacy power-consuming devices on a stackable device, enter the following command at the stack unit CONFIG level of the CLI.

```
device(config-unit-2) # no legacy-inline-power
```

On chassis devices, you can disable support for legacy power-consuming devices per slot. To disable legacy support on all ports in slot 2, enter the following command at the global CONFIG level of the CLI.

```
device(config) # no legacy-inline-power 2
```

NOTE

The **no legacy-inline-power** command does not require a software reload if it is entered prior to connecting the PDs. If the command is entered after the PDs are connected, the configuration must be saved (**write memory**) and the software reloaded after the change is placed into effect.

Syntax: `[no] legacy-inline-power [slotnum]`

NOTE

By default, the inline-power command reserves 30 watts. On PoH ports, inline-power reserves 95 watts.

To re-enable support for legacy power-consuming devices after it has been disabled, enter the **legacy-inline-power** command (without the **no** parameter).

The `slotnum` variable is required for chassis devices when you disable or re-enable legacy support on a slot.

Use the **show run** command to view whether support for PoE legacy power-consuming devices is enabled or disabled.

The following example turns off support for legacy inline power on an FSX 800.

```
SX800-3J31-u7(config)# leg
  legacy-inline-power          set legacy (capacitance-based) PD detection -
                              default
SX800-3J31-u7(config)# legacy-inline-power
  DECIMAL Slot number
  <cr>
SX800-3J31-u7(config)# legacy-inline-power 8
SX800-3J31-u7(config)# no leg
  legacy-inline-power          set legacy (capacitance-based) PD detection -
                              default
SX800-3J31-u7(config)# no legacy-inline-power 8
```

The following example turns off legacy inline power for a single stack unit.

```
ICX7250-24-3J32-u12(config)# no legacy-inline-power
```

The following example turns off legacy inline power support on the entire stack.

```
ICX7250-24-3J32# configure terminal
ICX7250-24-3J32 (config)# stack unit 12
ICX7250-24-3J32-u12(config)# no legacy-inline-power
```

Enabling the detection of PoE power requirements advertised through CDP

Many power-consuming devices, such as Cisco VoIP phones and other vendors' devices, use the Cisco Discovery Protocol (CDP) to advertise their power requirements to power-sourcing devices, such as Brocade PoE devices. Brocade power-sourcing equipment is compatible with Cisco and other vendors' power consuming devices and can detect and process power requirements for these devices automatically.

NOTE

If you configure a port with a maximum power level or a power class for a power-consuming device, the power level or power class takes precedence over the CDP power requirement. If you want a device to adhere to the CDP power requirement, do not configure a power level or power class on the associated port.

Command syntax for PoE power requirements

To enable the Brocade device to detect CDP power requirements, enter the following commands.

```
device# configure terminal
device(config)# cdp run
```

Syntax: `[no] cdp run`

Use the **no** form of the command to disable the detection of CDP power requirements.

Setting the maximum power level for a PoE power-consuming device

When PoE is enabled on a port to which a power-consuming device, or PD, is attached, by default, a Brocade PoE device supplies 15.4 watts of power at the RJ-45 jack, minus any power loss through the cables. A PoE+ device supplies either 15.4 or 30 watts of power (depending on the type of PD connected to the port), minus any power loss through the cables. A PoH device supplies 15.4, 30, or 95 watts of power (depending on the type of PD connected to the port), minus any power loss through the cables.

As an example, a PoE port with a default maximum power level of 15.4 watts receives a maximum of 12.95 watts of power after 2.45 watts of power loss through the cable. This is compliant with the IEEE 802.3af and 802.3at specifications for delivering inline power. Devices that are configured to receive less PoE power, for example, 4.0 watts of power, experience a lower rate of power loss through the cable.

If desired, you can manually configure the maximum amount of power that the Brocade PoE device supplies at the RJ-45 jack.

Considerations for setting power levels

Consider the following when enabling this feature:

- There are two ways to configure the power level for a PoE, PoE+, or High PoE power-consuming device. The first method is discussed in this section. The other method is provided in the section [Setting the power class for a PoE power-consuming device](#) on page 326. For each PoE port, you can configure either a maximum power level or a power class. You cannot configure both. You can, however, configure a maximum power level on one port and a power class on another port.
- The Brocade PoE, PoE+, or High PoE device adjusts the power on a port only if there are available power resources. If power resources are not available, the following message is displayed on the console and in the Syslog:


```
PoE: Failed power allocation of 30000 mwatts on port 1/1/21. Will retry when more power budget.
```
- If you are not using High PoE or PoH devices in any of the first 8 ports of the ICX7450-48P or ICX7450-24P, Brocade recommends that you limit the power on those ports using the **inline power power-limit** command. Limiting power with the **inline power power-by-class 4** command does not work for the ICX7450 because Class 4 encompasses 30-95W. However, Class 4 on units that do not support PoH or High Power is still 30W.
- FastIron devices pre-allocate power as per the configured maximum power for a physically operational PoE, PoE+, or High PoE configured port. However, in an ICX 6450-C12 device that is operational without direct power supply and that has pass-through power, there is no pre-allocation of power. Instead, power is allocated only when a powered device is connected to the port.

Configuring power levels command syntax

To configure the maximum power level for a power-consuming device, use the **inline power power-limit** command as shown in the following example.

```
device# configure terminal
device(config)# interface ethernet 1/1/1
device(config-if-e1000-1/1/1)# inline power power-limit 14000
```

These commands enable inline power on interface ethernet 1 in slot 1 of unit 1 and set the PoE power level to 14,000 milliwatts (14 watts).

Syntax: `inline power power-limit power-level`

The *power level* variable is the maximum power level in number of milliwatts. The following values are supported:

- **PoE** - Enter a value from 1000 through 15,400. The default is 15,400.
- **PoE+** - Enter a value from 1000 through 30,000. The default is 30,000.
- **PoH** - Enter a value from 1000 through 95,000. The default is 95,000. Value is always adjusted to nearest multiple of 5.

NOTE

Do not configure a power level higher than the default listed. Setting the power level higher than the default could damage the PD.

For information about resetting the maximum power level, refer to [Resetting PoE parameters](#) on page 329.

Setting the power class for a PoE power-consuming device

A power class specifies the maximum amount of power that a Brocade PoE, PoE+, or PoH device supplies to a power-consuming device. The following table shows the different power classes and their respective maximum power allocations.

TABLE 49 Power classes for PDs

Class	Usage	Power (watts) from Power-Sourcing Device		
		Standard PoE	PoE+	Power over HDBaseT (PoH)
0	default	15.4	15.4	15.4
1	optional	4	4	4
2	optional	7	7	7
3	optional	15.4	15.4	15.4
4	optional	15.4	30	95

Refer to [Considerations for setting power levels](#) on page 325 for essential information. Consider the following points when setting the power class for a PoE power-consuming device.

- The power class includes any power loss through the cables. For example, a PoE port with a power class of 3 (15.4 watts) receives a maximum of 12.95 watts of power after 2.45 watts of power loss through the cable. This is compliant with the IEEE 802.3af and 802.3at specifications for delivering inline power. Devices that are configured to receive less PoE power, for example, class 1 devices (4.0 watts), experience a lower rate of power loss through the cable.
- The Brocade PoE, PoE+, or PoH device adjusts the power on a port only if there are available power resources. If power resources are not available, the following message is displayed on the console and in the Syslog:

```
PoE: Failed power allocation of 30000 mwatts on port 1/1/21. Will retry when more power budget.
```

Setting the power class command syntax

To configure the power class for a PoE power consuming device, enter commands such as the following.

```
ICX7250-48P Switch# configure terminal
ICX7250-48P Switch(config)# interface ethernet 1/1/1
ICX7250-48P Switch(config-if-e1000-1/1/1)# inline power power-by-class 4
Warning: Inline power configuration on port 1/1/1 has been modified.
ICX7250-48P Switch(config-if-e1000-1/1/1)# show inline power 1
```

```
Power Capacity:      Total is 720000 mWatts. Current Free is 690000 mWatts.
Power Allocations:  Requests Honored 3 times
```

Port	Admin State	Oper State	---Power(mWatts)---		PD Type	PD Class	Pri	Fault/Error
			Consumed	Allocated				
1/1/1	On	On	14460	30000	802.3af	Class 3	3	n/a

These commands enable inline power on interface ethernet 1 in slot 1 of unit 1 and set the power class to 2.

Syntax: `inline power power-by-class class value`

The *class value* variable is the power class. Enter a value between 0 and 4. The default is 0. The table in [Setting the power class for a PoE power-consuming device](#) on page 326 shows the different power classes and their respective maximum power allocations.

For information about resetting the power class, refer to [Resetting PoE parameters](#) on page 329.

Setting the power budget for a PoE interface module on an FSX device

By default, each PoE and PoE+ interface module has a maximum power budget of 65535 watts.

On an FSX device, you can change the amount of power allocated to each PoE or PoE+ interface module installed in the chassis.

NOTE

FSX does not support PoH.

To change the power allocation on an FSX device, use the **inline power budget** command as shown in the following example.

```
device(config)# inline power budget 150000 module 7
```

In the previous example, the command allocates 150000 milliwatts (150 watts) to the PoE interface module in slot 7. The command takes effect immediately. The results are displayed in the "power budget" column in the **show inline power detail** output. The configuration (inline power budget 150000 module 7) is displayed in the **show running-config** output.

Syntax: `inline power budget num module slot`

The *num* variable is the number of milliwatts to allocate to the module. Enter a value from 0 through 65535000.

The *slot* variable specifies where the PoE or PoE+ module resides in the chassis.

Setting the inline power priority for a PoE port

In a configuration where PoE power-consuming devices collectively have a greater demand for power than the PoE power supply or supplies can provide, the FastIron PoE device must place the PoE ports that it cannot power in *standby* or *denied* mode (waiting for power) until the available power increases. The available power increases when one or more PoE ports are powered down, or, if applicable, when an additional PoE power supply is installed in the FastIron PoE device.

When PoE ports are in *standby* or *denied* mode (waiting for power) and the FastIron PoE device receives additional power resources, by default, the device allocates newly available power to the standby ports in priority order, with the highest priority ports first, followed by the next highest priority ports, and so on. Within a given priority, standby ports are considered in ascending order, by slot number and then by port number, provided enough power is available for the ports. For example, PoE port 1/1/11 should receive power before PoE port 1/2/1. However, if PoE port 1/1/11 needs 12 watts of power and PoE port 1/2/1 needs 10 watts of power, but only 11 watts of

power become available on the device, the FastIron PoE device allocates the power to port 1/2/1 because it does not have sufficient power for port 1/1/11.

You can configure an *inline power priority* on PoE ports, so that ports with a higher inline power priority take precedence over ports with a low inline power priority. For example, if a new PoE port comes online and the port is configured with a high priority, if necessary (if power is already fully allocated to power consuming devices), the FastIron PoE device removes power from a PoE port or ports that have a lower priority and allocates the power to the PoE port that has the higher value.

Ports that are configured with the same inline power priority are given precedence based on the slot number and port number in ascending order, provided enough power is available for the port. For example, if both PoE port 1/1/2 and PoE port 1/2/1 have a high inline power priority value, PoE port 1/1/2 receives power before PoE port 1/2/1. However, if PoE port 1/1/2 needs 12 watts of power and PoE port 1/2/1 needs 10 watts of power, but only 11 watts of power become available on the device, the FastIron PoE device allocates the power to PoE port 1/2/1 because it does not have sufficient power for port 1/1/2. By default, all ports are configured with a low inline power priority.

Command syntax for setting the inline power priority for a PoE port

To configure an inline power priority for a PoE port on a FastIron PoE device, use the **inline power priority** command as shown in the following example.

```
ICX7250-48P Switch# configure terminal
ICX7250-48P Switch(config)# interface ethernet 1/1/1
ICX7250-48P Switch(config-if-e1000-1/1/1)# inline power
  decouple-datalink    Decouple PoE from data link operational behavior
  power-by-class       Allocate power based on class of the power devices
  power-limit          Allocate power based on specified limit
  priority              Priority class for the purpose of power management
  <cr>
ICX7250-48P Switch(config-if-e1000-1/1/1)# inline power priority
  DECIMAL    Priority value 1..3 (highest..lowest)
  <cr>
ICX7250-48P Switch(config-if-e1000-1/1/1)# inline power priority 1
Warning: Inline power configuration on port 1/1/1 has been modified.
```

In the previous example, the command enables inline power on interface ethernet 1 in slot 1 of unit 1 and sets the inline power priority level to high.

Syntax: **[no] inline power priority** *priority num*

The *priority num* parameter is the inline power priority number. The default is 3 (low priority). You can specify one of the following values:

- 3 - Low priority
- 2 - High priority
- 1 - Critical priority

Use the **inline power** command without a priority number to reset a port priority to the default (low) priority.

Use the **no inline power** command to disable the port from receiving inline power.

For information about resetting the inline power priority, refer to "Resetting PoE parameters."

To view the inline power priority for all PoE ports, issue the **show inline power** command at the Privileged EXEC level of the CLI. Refer to "Displaying PoE operational status."

Resetting PoE parameters

You can override or reset PoE port parameters including power priority, power class, and maximum power level. To do so, you must specify each PoE parameter in the CLI command line. This section provides some CLI examples.

NOTE

When you reset PoE parameters on an FSX 800 or FSX 1600, you reset the parameters for the entire PoE chassis.

1--Changing a PoE port power priority from low to high

To change a PoE port power priority from low (the default value) to high and keep the current maximum configured power level of 3000, enter commands such as the following.

```
device# configure terminal
device(config)# interface ethernet 1/1/1
device(config-if-e1000-1/1/1)# inline power priority 2 power-limit 3000
```

You must specify both the inline power priority and the maximum power level (**power-limit** command), even though you are keeping the current configured maximum power level at 3000. If you do not specify the maximum power level, the device will apply the default value. Also, you must specify the inline power priority before specifying the power limit.

2--Changing a port power class from 2 to 3

To change a port power class from 2 (7 watts maximum) to 3 (15.4 watts maximum) and keep the current configured power priority of 2, enter commands such as the following.

```
device#configure terminal
device(config)# interface ethernet 1/1/1
device(config-if-e1000-1/1/1)# inline power priority 2 power-by-class 3
```

You must specify both the power class and the inline power priority, even though you are not changing the power priority. If you do not specify the power priority, the device will apply the default value of 3 (low priority). Also, you must specify the inline power priority before specifying the power class.

The following example sets PoE parameters on interface 2/1/1 in stack unit 12.

```
ICX7250-24-3J32# configure terminal
ICX7250-24-3J32(config)# stack unit 12
ICX7250-24-3J32-u12(config)# interface ethernet 2/1/1
ICX7250-24-3J32-u12(config-if-e1000-2/1/1)# inline power
  decouple-datalink    Decouple PoE from data link operational behavior
  power-by-class       Allocate power based on class of the power devices
  power-limit          Allocate power based on specified limit
  priority              Priority class for the purpose of power management
  <cr>
ICX7250-24-3J32-u12(config-if-e1000-2/1/1)# inline power priority
  DECIMAL    Priority value 1..3 (highest..lowest)
  <cr>
ICX7250-24-3J32-u12(config-if-e1000-2/1/1)# inline power priority 3
  power-by-class    Allocate power based on class of the power devices
  power-limit       Allocate power based on specified limit
  <cr>
ICX7250-24-3J32-u12(config-if-e1000-2/1/1)# inline power priority 3 power-limit 14000
ICX7250-24-3J32-u12(config-if-e1000-2/1/1)#
```

Displaying Power over Ethernet information

The show commands described in this section are available for viewing PoE operational status, PD data, and PoE power supply status.

Displaying PoE operational status

The **show inline power** command displays operational information about Power over Ethernet.

You can view the PoE operational status for the entire device, for a specific PoE module only, or for a specific interface only. In addition, you can use the **show inline power detail** command to display in-depth information about PoE power supplies. To display PoE data specific to PD ports, use the **show inline power pd** command.

The following example displays **show inline power** command output for a PoE device.

```
ICX7250-48P Router# show inline power

Power Capacity:          Total is 720000 mWatts. Current Free is 384000 mWatts.

Power Allocations:      Requests Honored 146 times
```

Port	Admin State	Oper State	---Power(mWatts)---		PD Type	PD Class	Pri	Fault/Error
			Consumed	Allocated				
1/1/1	On	On	6385	7000	802.3af	Class 2	3	n/a
1/1/2	On	On	6479	7000	802.3af	Class 2	3	n/a
1/1/3	On	On	6479	7000	802.3af	Class 2	3	n/a
1/1/4	On	On	6573	7000	802.3af	Class 2	3	n/a
1/1/5	On	On	6479	7000	802.3af	Class 2	3	n/a
1/1/6	On	On	6479	7000	802.3af	Class 2	3	n/a
1/1/7	On	On	6385	7000	802.3af	Class 2	3	n/a
1/1/8	On	On	6385	7000	802.3af	Class 2	3	n/a
1/1/9	On	On	6385	7000	802.3af	Class 2	3	n/a
1/1/10	On	On	6479	7000	802.3af	Class 2	3	n/a
1/1/11	On	On	6385	7000	802.3af	Class 2	3	n/a
1/1/12	On	On	6385	7000	802.3af	Class 2	3	n/a
1/1/13	On	On	6291	7000	802.3af	Class 2	3	n/a
1/1/14	On	On	6385	7000	802.3af	Class 2	3	n/a
1/1/15	On	On	5915	7000	802.3af	Class 2	3	n/a
1/1/16	On	On	6385	7000	802.3af	Class 2	3	n/a
1/1/17	On	On	6479	7000	802.3af	Class 2	3	n/a
1/1/18	On	On	6573	7000	802.3af	Class 2	3	n/a
1/1/19	On	On	6479	7000	802.3af	Class 2	3	n/a
1/1/20	On	On	6573	7000	802.3af	Class 2	3	n/a
1/1/21	On	On	6479	7000	802.3af	Class 2	3	n/a
1/1/22	On	On	6479	7000	802.3af	Class 2	3	n/a
1/1/23	On	On	6479	7000	802.3af	Class 2	3	n/a
1/1/24	On	On	6479	7000	802.3af	Class 2	3	n/a
1/1/25	On	On	6385	7000	802.3af	Class 2	3	n/a
1/1/26	On	On	6385	7000	802.3af	Class 2	3	n/a
1/1/27	On	On	6385	7000	802.3af	Class 2	3	n/a
1/1/28	On	On	6385	7000	802.3af	Class 2	3	n/a
1/1/29	On	On	6385	7000	802.3af	Class 2	3	n/a
1/1/30	On	On	6385	7000	802.3af	Class 2	3	n/a
1/1/31	On	On	6385	7000	802.3af	Class 2	3	n/a
1/1/32	On	On	6385	7000	802.3af	Class 2	3	n/a
1/1/33	On	On	6291	7000	802.3af	Class 2	3	n/a
1/1/34	On	On	6291	7000	802.3af	Class 2	3	n/a
1/1/35	On	On	6291	7000	802.3af	Class 2	3	n/a
1/1/36	On	On	6291	7000	802.3af	Class 2	3	n/a
1/1/37	On	On	6291	7000	802.3af	Class 2	3	n/a
1/1/38	On	On	6385	7000	802.3af	Class 2	3	n/a
1/1/39	On	On	6291	7000	802.3af	Class 2	3	n/a
1/1/40	On	On	6291	7000	802.3af	Class 2	3	n/a
1/1/41	On	On	6385	7000	802.3af	Class 2	3	n/a
1/1/42	On	On	6479	7000	802.3af	Class 2	3	n/a
1/1/43	On	On	6385	7000	802.3af	Class 2	3	n/a
1/1/44	On	On	6479	7000	802.3af	Class 2	3	n/a
1/1/45	On	On	6291	7000	802.3af	Class 2	3	n/a
1/1/46	On	On	6385	7000	802.3af	Class 2	3	n/a
1/1/47	On	On	6385	7000	802.3af	Class 2	3	n/a
1/1/48	On	On	6385	7000	802.3af	Class 2	3	n/a
Total			306950	336000				

Syntax: `show inline power [device/slot/port]`

TABLE 50 Field definitions for the show inline power command

Column	Definition
Power Capacity	The total PoE power supply capacity and the amount of available power (current free) for PoE power consuming devices. Both values are shown in milliwatts.
Power Allocations	The number of times the device fulfilled PoE requests for power.
Port	The slot number and port number.
Admin State	Specifies whether or not Power over Ethernet has been enabled on the port. This value can be one of the following: <ul style="list-style-type: none"> On - The inline power command was issued on the port. Off - The inline power command has not been issued on the port.
Oper State	Shows the status of inline power on the port. This value can be one of the following: <ul style="list-style-type: none"> On - The PoE power supply is delivering inline power to the PD. Off - The PoE power supply is not delivering inline power to the PD. Denied - The port is in standby mode (waiting for power) because the device does not currently have enough available power for the port. <p>NOTE When you enable a port using the CLI, it may take 12 or more seconds before the operational state of that port is displayed correctly in the show inline power output.</p>
Power Consumed	The number of current, actual milliwatts that the PD is consuming.
Power Allocated	The number of milliwatts allocated to the port. This value is either the default or configured maximum power level, or the power class that was automatically detected by the device.
PD Type	The type of PD connected to the port. This value can be one of the following: <ul style="list-style-type: none"> 802.3at - The PD connected to this port is 802.3at-compliant. 802.3af - The PD connected to this port is 802.3af-compliant. Legacy - The PD connected to this port is a legacy product (not 802.3af-compliant). N/A - Power over Ethernet is configured on this port, and one of the following is true: <ul style="list-style-type: none"> The device connected to this port is a non-powered device. No device is connected to this port. The port is in <i>standby</i> or <i>denied</i> mode (waiting for power).
PD Class	Determines the maximum amount of power a PD receives. The table in the section Setting the power class for a PoE power-consuming device on page 326 shows the different power classes and their respective maximum power allocations. <p>This field can also be "Unknown" when the device attached to the port cannot advertise its power class.</p> <p>NOTE If an 802.3at PD with a class 4 value is connected to a Brocade FastIron switch, the switch must be running FastIron release 08.0.20 or later to be able to perform the necessary power negotiations.</p>

TABLE 50 Field definitions for the show inline power command (continued)

Column	Definition
Pri	<p>The port <i>in-line power priority</i>, which determines the order in which the port will receive power while in <i>standby</i> mode (waiting for power). Ports with a higher priority will receive power before ports with a low priority. This value can be one of the following:</p> <ul style="list-style-type: none"> • 3 - Low priority • 2 - High priority • 1 - Critical priority
Fault/Error	<p>If applicable, this is the fault or error that occurred on the port. This value can be one of the following:</p> <ul style="list-style-type: none"> • critical temperature - The PoE chip temperature limit rose above the safe operating level, thereby powering down the port. • detection failed - discharged capacitor - The port failed capacitor detection (legacy PD detection) because of a discharged capacitor. This can occur when connecting a non-PD on the port. • detection failed - out of range capacitor - The port failed capacitor detection (legacy PD detection) because of an out-of-range capacitor value. This can occur when connecting a non-PD on the port. • internal h/w fault - A hardware problem has hindered port operation. • lack of power - The port has shut down due to lack of power. • main supply voltage high - The voltage was higher than the maximum voltage limit, thereby tripping the port. • main supply voltage low - The voltage was lower than the minimum voltage limit, thereby tripping the port. • overload state - The PD consumes more power than the maximum limit configured on the port, based on the default configuration, user configuration, or CDP configuration. • over temperature - The port temperature rose above the temperature limit, thereby powering down the port. • PD DC fault - A succession of underload and overload states, or a PD DC/DC fault, caused the port to shutdown. • short circuit - A short circuit was detected on the port delivering power. • underload state - The PD consumes less power than the minimum limit specified in the 802.3af standard. • voltage applied from ext src - The port failed capacitor detection (legacy PD detection) because the voltage applied to the port was from an external source.
Total	<p>The total power in milliwatts being <i>consumed</i> by all PDs connected to the Interface module, and the total power in milliwatts <i>allocated</i> to all PDs connected to the Interface module.</p>
Grand Total	<p>The total number of current, actual milliwatts being <i>consumed</i> by all PDs connected to the FastIron PoE device, and the total number of milliwatts <i>allocated</i> to all PDs connected to the FastIron PoE device.</p>

Displaying PoE data specific to PD ports

The **show inline power pd** command displays operational information specific to the PD ports.

This command displays information about the number of PD ports available, how much PD power is available to PSE, how much PD power is currently switched to PSE, and the PD port level status.

If a PD module is present, then the command displays the following global power information for the PD ports:

- Total PD power available to PSE
- Total PD power switched to PSE

In the absence of valid PSU power, the total PD power switched is equal to that available to PSE, as shown in the following example.

```
device# show inline power pd
Number of PD Ports: 2
Total PD Power Available to PSE: 22400
Total PD Power Switched to PSE: 22400
Port   Oper   Oper   Fault/
      State Mode   Error
-----
1/2/1  On     802.3at n/a
1/2/2  On     802.3at n/a
```

The following shows an example of the **show inline power pd** display output on a PoE device with the internal PSU up and no PD ports on.

```
device# show inline power pd
Number of PD Ports: 2
Total PD Power Available to PSE: 0
Total PD Power Switched to PSE: 0
Port   Oper   Oper   Fault/
      State Mode   Error
-----
1/2/1  Off    n/a    n/a
1/2/2  Off    n/a    n/a
```

The following shows an example of the **show inline power pd** display output on a PoE device with the internal PSU up and one PD port on in the AT mode.

```
device# show inline power pd
Number of PD Ports: 2
Total PD Power Available to PSE: 0
Total PD Power Switched to PSE: 0
Port   Oper   Oper   Fault/
      State Mode   Error
-----
1/2/1  On     802.3at n/a
1/2/2  Off    n/a    n/a
```

The following shows an example of the **show inline power pd** display output on a PoE device with the internal PSU down and two PD ports on in the AT mode.

```
device# show inline power pd
Number of PD Ports: 2
Total PD Power Available to PSE: 22400
Total PD Power Switched to PSE: 22400
Port   Oper   Oper   Fault/
      State Mode   Error
-----
1/2/1  On     802.3at n/a
1/2/2  On     802.3at n/a
```

The following example of the **show inline power pd** command output is displayed for devices such as the ICX 6430-C12 that do not support PD ports.

```
device# show inline power pd
The cli is not applied to this platform
```

Syntax: show inline power pd

TABLE 51 Field definitions for the show inline power pd command

Column	Definition
Number of PD Ports	The number of PD ports in the system.
Total PD Power Available to PSE	Total PD power available to PSE.
Total PD Power Switched to PSE	Total PD power switched to PSE. It is either 0 or the total available power.
Port	The port number of the PD port.
Oper State	The operational state of the PD port. This value can be one of the following: <ul style="list-style-type: none"> On - The PD port is linked to a PSE port and is consuming power. Off - The PD port is not linked to a PSE port and is not consuming power.
Oper Mode	The operational mode of the PD port. This value is meaningful if Oper State is On. This value can be one of the following: <ul style="list-style-type: none"> 802.3af - The PD port is operating in the AF mode. 802.3at - The PD port is operating in the AT mode. n/a - The PD port is not operational.
Fault/Error	Shows the error or fault conditions affecting the PD port. This value can be one of the following: <ul style="list-style-type: none"> An error or fault code is displayed. n/a - No error or fault condition on the PD port.

Displaying detailed information about PoE power supplies

The **show inline power detail** command displays detailed operational information about the PoE power supplies in FastIron PoE switches. The command output differs on FCX POE+ switches compared to FSX Series and ICX switches.

The following is an example of the **show inline power detail** command output on an FCX POE+ switch.

```

device# show inline power detail
Power Supply Data On stack 1:
+++++
Power Supply #1:
    Max Curr:      7.5 Amps
    Voltage:       54.0 Volts
    Capacity:      410 Watts
POE Details Info. On Stack 1 :
General PoE Data:
+++++
Firmware
Version
-----
02.1.0
Cumulative Port State Data:
+++++
#Ports  #Ports  #Ports  #Ports  #Ports  #Ports  #Ports
Admin-On Admin-Off Oper-On  Oper-Off Off-Denied Off-No-PD Off-Fault
-----
45      3       0       48      0       45      0
Cumulative Port Power Data:
+++++
#Ports  #Ports  #Ports  Power      Power
Pri: 1  Pri: 2  Pri: 3  Consumption Allocation
-----
0       0       45     0.0 W      0.0 W
Power Supply Data On stack 2:

```

```

+++++
Power Supply Data:
+++++
Power Supply #1:
    Max Curr:      7.5 Amps
    Voltage:       54.0 Volts
    Capacity:      410 Watts
POE Details Info. On Stack 2 :
General PoE Data:
+++++
Firmware
Version
-----
02.1.0
... continued on next page...
Slot #Ports #Ports #Ports Power Power Power
Pri: 1 Pri: 2 Pri: 3 Consumption Allocation Budget
-----
3 0 0 48 513.468 W 739.200 W 65535.0 W
4 0 0 48 1349.320 W 1440.0 W 65535.0 W
-----
Total:0 0 96 1862.788 W 2179.200 W 131070.0 W
... continued from previous page...
Cumulative Port State Data:
+++++
#Ports #Ports #Ports #Ports #Ports #Ports #Ports
Admin-On Admin-Off Oper-On Oper-Off Off-Denied Off-No-PD Off-Fault
-----
20      4      0      24      0      20      0
Cumulative Port Power Data:
+++++
#Ports #Ports #Ports      Power      Power
Pri: 1 Pri: 2 Pri: 3 Consumption Allocation
-----
20      0      0      0.0 W      0.0 W
Power Supply Data On stack 3:
+++++
Power Supply #1:
    Max Curr:      7.5 Amps
    Voltage:       54.0 Volts
    Capacity:      410 Watts
POE Details Info. On Stack 3 :
General PoE Data:
+++++
Firmware
Version
-----
02.1.0
Cumulative Port State Data:
+++++
#Ports #Ports #Ports #Ports #Ports #Ports #Ports
Admin-On Admin-Off Oper-On Oper-Off Off-Denied Off-No-PD Off-Fault
-----
22      2      0      24      0      22      0
Cumulative Port Power Data:
+++++
#Ports #Ports #Ports      Power      Power
Pri: 1 Pri: 2 Pri: 3 Consumption Allocation
-----
0      10     12      0.0 W      0.0 W

```

The following is an example of the **show inline power detail** command output on a FSX Series PoE switch.

```

device# show inline power detail
Power Supply Data:
+++++
PoE+ Max Operating Voltage: 54 V
Power Supply #1:
    Model Number:   32004000
    Serial Number:  093786124716
    Firmware Ver:   1.6

```

```

Test Date:      9/12/09 (mm/dd/yy)
H/W Status:    807
Max Curr:      50.0 Amps
Voltage:       54.0 Volts
Capacity:      2500 Watts
PoE Capacity:  2260 Watts
Consumption:   2095 Watts

General PoE Data:
+++++
Slot  Firmware
      Version
-----
3   Device 1: 02.1.0  Device 2: 02.1.0
4   Device 1: 02.1.0  Device 2: 02.1.0
6   02.1.0
7   Device 1: 02.1.0  Device 2: 02.1.0
8   02.1.0

Cumulative Port State Data:
+++++
Slot  #Ports  #Ports  #Ports  #Ports  #Ports  #Ports  #Ports
      Admin-On Admin-Off Oper-On Oper-Off Off-Denied Off-No-PD Off-Fault
-----
3     48      0        48      0        0          0          0
4     48      0        48      0        0          0          0
6     24      0         0       24       0         24         0
7     48      0         4       44       44         0          0
8     24      0         0       24       0         24         0
-----
Total:192      0        100     92        44         48         0
... continued on next page...
... continued from previous page...
Cumulative Port Power Data:
+++++
Slot  #Ports  #Ports  #Ports      Power      Power      Power
      Pri: 1  Pri: 2  Pri: 3  Consumption  Allocation  Budget
-----
3     0      0      48      513.90 W    739.200 W    65535.0 W
4     0      0      48     1346.497 W  1440.0 W     65535.0 W
6     0      0      24         0.0 W      0.0 W        65535.0 W
7     0      0      48      43.72 W    61.600 W    65535.0 W
8     0      0      24         0.0 W      0.0 W        65535.0 W
-----
Total:0      0      192     1902.659 W  2240.800 W   327675.0 W

```

The following is an example of **show inline power detail** command output for an ICX 7250 stack.

```
ICX7250-48p-3J32-u10#show inline power detail
```

```
Power Supply Data On stack 1:
+++++
```

```
Power Supply Data:
+++++
```

```
Power Supply #1:
  Max Curr:      13.3 Amps
  Voltage:       54.0 Volts
  Capacity:      720 Watts
Power Supply #2:
  Max Curr:      6.6 Amps
  Voltage:       54.0 Volts
  Capacity:      360 Watts
Power Supply #3:
  Max Curr:      6.6 Amps
  Voltage:       54.0 Volts
  Capacity:      360 Watts
```


POE Details Info. On Stack 1 :

General PoE Data:
 ++++++

Firmware
 Version

 01.2.1 Build 003

Cumulative Port State Data:
 ++++++

#Ports Admin-On	#Ports Admin-Off	#Ports Oper-On	#Ports Oper-Off	#Ports Off-Denied	#Ports Off-No-PD	#Ports Off-Fault
48	0	0	48	0	47	1

Cumulative Port Power Data:
 ++++++

#Ports Pri: 1	#Ports Pri: 2	#Ports Pri: 3	Power Consumption	Power Allocation
0	0	48	0.0 W	0.0 W

ICX7250-48p-3J32-u10#

The following example provides details on an ICX 7250 connected to an EPS.

ICX7250-24P Router# show chassis
 The stack unit 1 chassis info:

Power supply 1 (NA - AC - PoE) present, status ok
 Power supply 1 Fan Air Flow Direction: Front to Back
 Power supply 2 (NA - DC - PoE) present, status ok

Fan 1 ok, speed (manual): [[1]]<->2
 Fan 2 ok, speed (manual): [[1]]<->2

Fan controlled temperature:
 Rule 1/2 (MGMT THERMAL PLANE): 49.0 deg-C
 Rule 2/2 (PoE THERMAL PLANE): 40.5 deg-C

Fan speed switching temperature thresholds:
 Rule 1/2 (MGMT THERMAL PLANE):
 Speed 1: NM<----->93 deg-C
 Speed 2: 82<----->105 deg-C (shutdown)
 Rule 2/2 (PoE THERMAL PLANE):
 Speed 1: NM<----->58 deg-C
 Speed 2: 49<----->105 deg-C (shutdown)

Fan 1 Air Flow Direction: Front to Back
 Fan 2 Air Flow Direction: Front to Back
 Slot 1 Current Temperature: 49.0 deg-C (Sensor 1), 39.5 deg-C (Sensor 2)
 Slot 2 Current Temperature: NA
 Warning level.....: 100.0 deg-C
 Shutdown level.....: 105.0 deg-C
 Boot Prom MAC : cc4e.24b4.906c
 Management MAC: cc4e.24b4.906c

ICX7250-24P Router# show inline power

Power Capacity: Total is 720000 mWatts. Current Free is 0 mWatts.

Power Allocations: Requests Honored 82 times

Port	Admin State	Oper State	---Power(mWatts)---		PD Type	PD Class	Pri	Fault/Error
			Consumed	Allocated				
1/1/1	On	On	28264	30000	802.3at	Class 4	3	n/a
1/1/2	On	On	28921	30000	802.3at	Class 4	3	n/a
1/1/3	On	On	28170	30000	802.3at	Class 4	3	n/a
1/1/4	On	On	28170	30000	802.3at	Class 4	3	n/a
1/1/5	On	On	28452	30000	802.3at	Class 4	3	n/a
1/1/6	On	On	28170	30000	802.3at	Class 4	3	n/a
1/1/7	On	On	28452	30000	802.3at	Class 4	3	n/a
1/1/8	On	On	28358	30000	802.3at	Class 4	3	n/a
1/1/9	On	On	28170	30000	802.3at	Class 4	3	n/a
1/1/10	On	On	28170	30000	802.3at	Class 4	3	n/a
1/1/11	On	On	28170	30000	802.3at	Class 4	3	n/a
1/1/12	On	On	28170	30000	802.3at	Class 4	3	n/a
1/1/13	On	On	28264	30000	802.3at	Class 4	3	n/a
1/1/14	On	On	28264	30000	802.3at	Class 4	3	n/a
1/1/15	On	On	26010	30000	802.3at	Class 4	3	n/a
1/1/16	On	On	28358	30000	802.3at	Class 4	3	n/a
1/1/17	On	On	28546	30000	802.3at	Class 4	3	n/a
1/1/18	On	On	28640	30000	802.3at	Class 4	3	n/a
1/1/19	On	On	28640	30000	802.3at	Class 4	3	n/a
1/1/20	On	On	28640	30000	802.3at	Class 4	3	n/a
1/1/21	On	On	28640	30000	802.3at	Class 4	3	n/a
1/1/22	On	On	28640	30000	802.3at	Class 4	3	n/a
1/1/23	On	On	28452	30000	802.3at	Class 4	3	n/a
1/1/24	On	On	28640	30000	802.3at	Class 4	3	n/a
Total			679371	720000				

ICX7250-24P Router# show inline power detail

Power Supply Data On stack 1:

+++++

Power Supply Data:

+++++

Power Supply #1:

Max Curr: 6.6 Amps
 Voltage: 54.0 Volts
 Capacity: 360 Watts

Power Supply #2:

Max Curr: 6.6 Amps
 Voltage: 54.0 Volts
 Capacity: 360 Watts

POE Details Info. On Stack 1 :

General PoE Data:

+++++

Firmware
 Version

 01.6.1 Build 009

Cumulative Port State Data:

+++++

```

#Ports   #Ports   #Ports   #Ports   #Ports   #Ports   #Ports
Admin-On Admin-Off Oper-On  Oper-Off Off-Denied Off-No-PD Off-Fault
-----
24       0        24      0        0        0        0
    
```

Cumulative Port Power Data:
 ++++++

```

#Ports #Ports #Ports      Power      Power
Pri: 1  Pri: 2  Pri: 3  Consumption Allocation
-----
0       0      24      679.371 W  720.0 W
    
```

ICX7250-24P Router#

Syntax: show inline power detail

TABLE 52 Field definitions for the show inline power detail command

Column	Definition
Power supply data	
PoE+ Max Operating Voltage	This field is applicable to FSX 800 and FSX 1600 PoE+ chassis devices only. It displays the maximum operating voltage supported by the PoE power supply. Possible values are: <ul style="list-style-type: none"> • 52 V • 54 V
Model Number	The manufacturing part number of the PoE power supply. Values are eight digits in length and begin with "32" or "30" for example: <ul style="list-style-type: none"> • 32016000 • 32007000
Serial Number	The serial number of the PoE power supply, for example, AA100730213.
Firmware Ver	The PoE power supply firmware version.
Test Date	The PoE power supply firmware test date in the format mm/dd/yyyy.
H/W Status	The PoE power supply hardware status code. This field is used by Brocade Technical Support for troubleshooting.
Max Curr	The PoE power supply maximum current capacity.
Voltage	The PoE power supply current input voltage.
Capacity	The PoE power supply total power capacity (in watts).
PoE Capacity	The PoE power supply PoE power capacity (in watts).
Consumption	The total number of watts consumed by PoE power consuming devices and PoE modules in the system, plus any internal or cable power loss. <p>NOTE Under the lower total inline power consumption level by Powered Devices (PDs) on FSX 800 and FSX 1600 devices, the power consumption displayed by the power supply units (PSUs) is inaccurately displayed as lower than the actual power consumption of the PSUs due to the sensitivity limitations of power supply measurements.</p>
General PoE data	
Slot	The device/slot/port number.
Firmware Version	The firmware version for the device/slot/port.
Cumulative port state data	

TABLE 52 Field definitions for the show inline power detail command (continued)

Column	Definition
<p>NOTE When you enable a port using the CLI, it may take 12 or more seconds before the operational state of that port is displayed correctly in the show inline power output.</p>	
Slot	The Interface device/slot/port number.
#Ports Admin-On	The number of ports on the Interface module on which the inline power command was issued.
#Ports Admin-Off	The number of ports on the Interface module on which the inline power command was not issued.
#Ports Oper-On	The number of ports on the Interface module that are receiving inline power from the PoE power supply.
#Ports Oper-Off	The number of ports on the Interface module that are not receiving inline power from the PoE power supply.
#Ports Off-Denied	The number of ports on the Interface module that were denied power because of insufficient power.
#Ports Off-No-PD	The number of ports on the Interface module to which no PDs are connected.
#Ports Off-Fault	The number of ports on the Interface module that are not receiving power because of a subscription overload.
Total	The totals for all of the fields in the Cumulative Port State Data report.
Cumulative port power data	
Slot	The device/slot/port number.
#Ports Pri: 1	The number of PoE ports on the Interface module that have a PoE port priority of 1.
#Ports Pri: 2	The number of PoE ports on the Interface module that have a PoE port priority of 2.
#Ports Pri: 3	The number of PoE ports on the Interface module that have a PoE port priority of 3.
Power Consumption	The total number of watts consumed by PoE power consuming devices, plus any cable loss.
Power Allocation	The number of watts allocated to the Interface module PoE ports. This value is the sum of the ports' default or configured maximum power levels, or power classes automatically detected by the FastIron PoE device.
Power Budget	The power budget allocated to the slot. The default value is 65535 watts. Any other value indicates that the power budget was configured using the CLI command inline power budget .
Total	The totals for all of the fields in the Cumulative Port Power Data report.

Inline power on PoE LAG ports

The inline power on Power over Ethernet (PoE) LAG ports feature allows you to enable inline power on PoE LAG ports with the introduction of a new **inline power ethernet** command, available in global configuration mode.

Without the **inline power ethernet** command, you cannot enable inline power on any secondary LAG ports because the interface configuration mode is not available for LAG secondary ports to run the **inline power** command.

You can configure inline power in interface configuration mode on a port that is not a member of a LAG. If that port then becomes part of a LAG, you can use the **inline power ethernet** command to configure inline power parameters on any other port in that LAG.

LAG operational changes can affect the PoE power state unless the **decouple-datalink** keyword is used as a command option when configuring inline power on the LAG ports. For more information, refer to "Decouple the PoE and datalink operations on PoE ports."

After configuring inline power on PoE ports, you can verify the configuration using the **show running-config** command. If you have configured inline power on a regular PoE port in either global configuration or interface configuration mode, the inline power configuration commands display under the interface configuration level. If a regular PoE port becomes a PoE LAG port, or a PoE LAG port is configured under global configuration mode, the inline power configuration commands display under the global configuration level. If a LAG is removed, the inline power configuration commands for all ports display under the interface configuration level.



WARNING

If you downgrade to a release earlier than 08.0.01, there is no backwards compatibility for the **inline power ethernet** command or the **decouple-datalink** keyword.

Restriction

If you want to keep decoupling in place on a PoE port when you configure the **inline power ethernet** command to change its other parameters, for example, priority, you must also configure the **decouple-datalink** keyword.

Configuring inline power on PoE ports in a LAG

Perform the following steps to configure and deploy a link aggregation group (LAG) on the required PoE ports on both the Brocade power sourcing equipment (PSE) and the PD. This task also enables inline power on the PoE ports.

1. Configure a LAG.

The following example configures a static LAG named mylag with an ID of 5.

```
Device(config)# lag "mylag" static id 5
```

2. Configure ports into the LAG membership.

The following example configures the four ports, 1/1/1, 1/1/2, 1/1/3, and 1/1/4, into the LAG membership.

```
Device(config-lag-mylag)# ports ethernet 1/1/1 to 1/1/4
```

3. Configure a primary port for the LAG.

The following example configures port 1/1/1 as the primary port.

```
Device(config-lag-mylag)# primary-port 1/1/1
```

4. Deploy the LAG.

The following example deploys the mylag LAG.

```
Device(config-lag-mylag)# deploy
```

5. Configure inline power on the primary port with the power-by-class option.

The following example configures inline power on the primary port, 1/1/1, with power-by-class option 3.

```
Device(config)# inline power ethernet 1/1/1 power-by-class 3
```

- Configure inline power on a secondary port with the default option.

The following example configures inline power on port 1/1/2 with the default option.

```
Device(config)# inline power ethernet 1/1/2
```

- Configure inline power on a secondary port with the power management option.

The following example configures inline power on port 1/1/3 with power management option 2. The range is 1 (lowest) to 3 (highest). The default is 1.

```
Device(config)# inline power ethernet 1/1/3 priority 2
```

- Configure inline power on a secondary port, specifying the actual power value.

The following example configures inline power on the port 1/1/4, specifying an actual power value of 12000 mWatts.

```
Device(config)# inline power ethernet 1/1/4 power-limit 12000
```

Decouple PoE and datalink operations on PoE ports

Although PoE and datalink operations are functionally independent of each other, some datalink operations affect the operational behavior of PoE ports. The Decoupling of PoE and Datalink Operations feature allows you to override the current default behavior.

The following are some example datalink operations that can affect the operational state of the PoE on PoE ports:

- Using disable or enable CLI on the power sourcing equipment (PSE) port interface
- Adding or deleting a tagged PSE port from a VLAN or VLAN group
- The PSE port enters an ErrDisable state
- Adding or deleting a PSE port from a LAG and deploying it

When the optional **decouple-datalink** keyword is configured using the **inline power** or **inline power ethernet** command, the datalink operational behavior on a PoE port does not affect the power state of the powered device (PD) that is connecting to the port. You can also configure the power limits and power-management priority. The **inline power** command is available in interface configuration mode for most PoE ports, and the **inline power ethernet** command is available in global configuration mode for LAG ports.

The feature Decoupling of Inline Power and Datalink Operations on PoE Ports is useful when a PoE port is powering a PD that serves a PSE device such as the ICX 6450-C12-PD.

NOTE

The **decouple-datalink** keyword was introduced in Release 08.0.01 to support the Decoupling of PoE and Datalink Operations feature. Decoupling of inline power and datalinks is not supported in releases earlier than Release 08.0.01.



WARNING

If you downgrade to a release earlier than 08.0.01, there is no backwards compatibility for the **decouple-datalink** keyword or the **inline power ethernet** command.

Restriction

If you want to keep decoupling in place on a PoE port when you configure the **inline power ethernet** command to change its other parameters, for example, priority, you must also configure the **decouple-datalink** keyword.

Decoupling of PoE and datalink operations on PoE LAG ports

Perform the following steps to decouple the behavior of the Power over Ethernet (PoE) and the datalink operations for PoE Link Aggregation Group (LAG) ports.

This task provides a method of overriding the current default behavior of datalink operations that affect the operation of PoE ports. If you use the optional **decouple-datalink** keyword when enabling inline power with the **inline power ethernet** command, the datalink operational behavior on a PoE port does not affect the power state of the powered device (PD) that is connecting to the port.

Configure this task on the Brocade PSE for any PoE ports that require the decoupling of inline power and datalink operations. Any Layer 2 features can then be configured and deployed on these PoE ports. To avoid the disruption of inline power after the LAG ports are powered up, perform the following configuration steps in order.

1. Configure inline power on the primary port with the power-by-class option.

The following example configures inline power on the primary port, 1/1/1, with power-by-class option 3 and decouples the datalink operations and the inline power for this port.

```
Device(config)# inline power ethernet 1/1/1 decouple-datalink power-by-class 3
```

2. Configure inline power on a secondary port with the default option.

The following example configures inline power on port 1/1/2 and decouples the datalink operations and the inline power for this port.

```
Device(config)# inline power ethernet 1/1/2 decouple-datalink
```

3. Configure inline power on a secondary port with the power-management priority option.

The following example configures inline power on port 1/1/3 with power-management priority 2 and decouples the datalink operations and the inline power for this port.

```
Device(config)# inline power ethernet 1/1/3 decouple-datalink priority 2
```

4. Configure inline power on a secondary port, specifying the actual power value.

The following example configures inline power on the port 1/1/4, specifying an actual power value of 12000 mWatts, and decouples the datalink operations and the inline power for this port.

```
Device(config)# inline power ethernet 1/1/4 decouple-datalink power-limit 12000
```

5. Configure a LAG.

The following example configures a static LAG named mylag with an ID of 5.

```
Device(config)# lag "mylag" static id 5
```

6. Configure ports into the LAG membership.

The following example configures the four ports, 1/1/1, 1/1/2, 1/1/3, and 1/1/4, into the LAG membership.

```
Device(config-lag-mylag)# ports ethernet 1/1/1 to 1/1/4
```

7. Configure a primary port for the LAG.

The following example configures port 1/1/1 as the primary port.

```
Device(config-lag-mylag)# primary-port 1/1/1
```

8. Deploy the LAG.

The following example deploys the mylag LAG.

```
Device(config-lag-mylag)# deploy
LAG mylag deployed successfully!
```

Decoupling of PoE and datalink operations on regular PoE ports

Use this procedure to decouple PoE and datalink operations on regular PoE ports.

While PoE and datalink operations are functionally independent of each other, some datalink operations affect the operational behavior of PoE ports. When the optional **decouple-datalink** keyword is configured using the **inline power** command, the datalink operational behavior on a PoE port does not affect the power state of the powered device (PD) that is connecting to the port. You can also configure the power limits and power-management priority. The **inline power** command is available in interface configuration mode for most PoE ports and the **inline power ethernet** command is available in global configuration mode for LAG ports.

Perform the following steps to enable inline power and decouple the behavior of the Power over Ethernet (PoE) and the datalink operations for regular PoE ports. This task provides a method of overriding the current default behavior of datalink operations that affect the operation of PoE ports. If you use the optional **decouple-datalink** keyword when enabling inline power using the **inline power** command, the datalink operational behavior on a PoE port does not affect the power state of the powered device (PD) that is connecting to the port.

NOTE

To enable inline power and decouple PoE and datalink operations on PoE LAG ports, refer to “Decoupling of PoE and datalink operations on PoE LAG ports.”

Perform this task on the Brocade PSE for any PoE ports that require the decoupling of PoE operations and datalink operations. Any Layer 2 features can then be configured and deployed on these PoE ports.

1. Enable interface configuration for a PoE port.

The following example enters interface configuration mode for Ethernet port 1/1/1.

```
Device(config)# interface ethernet 1/1/1
```

2. Configure inline power on the Ethernet 1/1/1 port with the power-by-class option.

The following example configures inline power on the PoE port, Ethernet 1/1/1, with power-by-class option 3 and decouples the datalink operations from the PoE operations for this port.

```
Device(config-if-e1000-1/1/1)# inline power decouple-datalink power-by-class 3
```

3. Enable interface configuration for Ethernet 1/1/2 port.

The following example enters interface configuration mode for Ethernet port 1/1/2.

```
Device(config-if-e1000-1/1/1)# interface ethernet 1/1/2
```

4. Configure inline power on Ethernet 1/1/2 port with the default option.

The following example configures inline power on Ethernet 1/1/2 port and decouples the datalink operations from the PoE operations for this port.

```
Device(config-if-e1000-1/1/2)# inline power decouple-datalink
```


5. Enable interface configuration for Ethernet 1/1/3 port.

The following example enters interface configuration mode for Ethernet port 1/1/3.

```
Device(config-if-e1000-1/1/2)# interface ethernet 1/1/3
```

6. Configure inline power on Ethernet port 1/1/3 with the power-management priority option.

The following example configures inline power on port 1/1/3 with power-management priority 2 and decouples the datalink operations from the PoE operations for this port.

```
Device(config-if-e1000-1/1/3)# inline power decouple-datalink priority 2
```

7. Enables interface configuration for Ethernet 1/1/4 port.

The following example enters interface configuration mode for Ethernet port 1/1/4.

```
Device(config-if-e1000-1/1/3)# interface ethernet 1/1/4
```

8. Configure inline power on Ethernet 1/1/4 port, specifying the actual power value.

The following example configures inline power on Ethernet port 1/1/4 port, specifies an actual power value of 12000 mWatts, and decouples the datalink operations and the PoE operations for this port.

```
Device(config-if-e1000-1/1/4)# inline power decouple-datalink power-limit 12000
```


40 Gbps Breakout Ports

- Overview of 40 Gbps breakout ports..... 347
- Configuring 40 Gbps breakout ports..... 347
- Configuring sub-ports..... 348
- Removing breakout configuration..... 350
- Displaying information for breakout ports..... 352

Overview of 40 Gbps breakout ports

A 40 Gbps breakout cable can be used on ICX 7750 standalone units to break out certain 40 Gbps ports into four 10 Gbps sub-ports. The 40 Gbps breakout cable is available for use on ICX 7750-48C, ICX 7750-48F, and ICX 7750-26Q models.

NOTE

Any interface-level configuration must be removed from a 40 Gbps port before it can be broken out into sub-ports. Refer to [Configuring 40 Gbps breakout ports](#) on page 347 for more information.

NOTE

Breakout can be configured only when the device is in store-and-forward mode. Breakout is not supported in cut-through mode.

NOTE

Stacking cannot be enabled on ICX 7750 units that have breakout configuration on any 40 Gbps ports.

Ports available for breakout are shown for each model in the following table. Refer to the *ICX 7750 Ethernet Switch Hardware Installation Guide* for information on installing breakout cables.

TABLE 53 ICX 7750 ports available for breakout

	ICX 7750-48C	ICX 7750-48F	ICX 7750-26Q
Module 1	N/A	N/A	1/1/5 through 1/1/16 (12 ports)
Module 2	1/2/1 through 1/2/6 (6 ports)	1/2/1 through 1/2/6 (6 ports)	1/2/1 through 1/2/6 (6 ports)
Module 3	1/3/1 through 1/3/6 (6 ports)	1/3/1 through 1/3/6 (6 ports)	1/3/1 through 1/3/6 (6 ports)

Configuring 40 Gbps breakout ports

Use the **breakout ethernet** command to divide available ICX 7750 40 Gbps ports into four 10 Gbps sub-ports when a breakout cable is attached.

By default, all main 40 Gbps ports are configured to come up in 40 Gbps mode. Once ports are cabled for breakout, configure the ports using the **breakout ethernet** command at the global configuration level.

NOTE

You should remove any interface-level configuration before configuring breakout.

NOTE

If the device is in cut-through mode and you attempt to configure breakout, an error is returned. Cut-through must be disabled to return the unit to store-and-forward mode before breakout is configured.

The **breakout ethernet** command first checks for existing configuration on the port. If existing configuration is detected, an error message similar to the following is displayed to indicate that prior configuration must be removed.

```
Device# configure terminal
Device(config)# breakout ethernet 1/1/11
Error: Port 1/1/11 is tagged
```

Once any previous configuration is removed, the **breakout ethernet** command must be reissued. The resulting configuration must be saved, and the unit must then be reloaded before the four 10 Gbps sub-ports are created and accessible.

For example, to configure ports 1/3/1 through 1/3/6 for breakout, issue the following commands:

```
Device# configure terminal
Device(config)# breakout ethernet 1/3/1 to ethernet 1/3/6
```

The following example configures breakout on port 1/1/5. On the first configuration attempt, an error is returned. The interface-level configuration is removed. Then the **write-memory** command is issued, followed by the **reload** command, to successfully configure the port for breakout.

```
Device# configure terminal
Device(config)# breakout ethernet 1/1/5
Error: Port 1/1/5 has sflow forwarding
Device(config)# interface ethernet 1/1/5
Device(config-if-e40000-1/1/5)# no sflow forwarding
Device(config-if-e40000-1/1/5)# end
Device# write memory
Write startup-config done.
Device# configure terminal
Device(config)# breakout ethernet 1/1/5
Reload required. Please write memory and then reload or power cycle.
Device(config)# write memory
Write startup-config done.
Device(config)# Flash Memory Write (8192 bytes per dot) .
Copy Done.
Device(config)# end
Device# reload
```

Configuring sub-ports

After 40 Gbps ports are successfully configured and activated for breakout, the sub-ports are available for configuration.

NOTE

Sub-port configuration persists only as long as the original 40 Gbps port is configured for breakout. Once breakout is removed and the device is reloaded, the sub-ports and their configuration are also removed.

NOTE

When a breakout cable is removed, the breakout configuration still exists. The user should manually issue the **no breakout** command to change a breakout port to a regular port.

Once a 40 Gbps port is broken out, the configuration is saved (with the **write memory** command), and the unit is reloaded with the updated configuration, four sub-ports are available for detailed configuration.

The sub-ports are configured like any other port; however, special four-tuple notation is required to reference them. Regular ports are identified by three-tuple notation; that is, by three numbers separated by a forward slash to indicate unit, slot, and port. For example 1/2/3 designates unit 1/slot 2/port 3. To designate sub-ports, you must add a fourth identification number, for example, 1/2/3:4. The four 10 Gbps sub-ports for port 1/2/3 can be represented as 1/2/3:1, 1/2/3:2, 1/2/3:3, and 1/2/3:4.

The following example shows no breakout on port 1/2/4, a 40 Gbps port that is up.

```
device# show interface brief
```

Port	Link	State	Dupl	Speed	Trunk	Tag	Pvid	Pri	MAC	Name
1/1/1	Down	None	None	None	None	No	1	0	cc4e.2439.3700	
1/1/2	Down	None	None	None	None	No	1	0	cc4e.2439.3701	
1/1/3	Down	None	None	None	None	No	1	0	cc4e.2439.3702	
1/1/4	Down	None	None	None	None	No	1	0	cc4e.2439.3703	
1/1/5	Down	None	None	None	None	No	1	0	cc4e.2439.3704	
1/1/6	Down	None	None	None	None	No	1	0	cc4e.2439.3708	
1/1/7	Down	None	None	None	None	No	1	0	cc4e.2439.370c	
1/1/8	Down	None	None	None	None	No	1	0	cc4e.2439.3710	
1/1/9	Down	None	None	None	None	No	1	0	cc4e.2439.3714	
1/1/10	Down	None	None	None	None	No	1	0	cc4e.2439.3718	
1/1/11	Down	None	None	None	None	No	1	0	cc4e.2439.371c	
1/1/12	Down	None	None	None	None	No	1	0	cc4e.2439.3720	
1/1/13	Down	None	None	None	None	No	1	0	cc4e.2439.3724	
1/1/14	Down	None	None	None	None	No	1	0	cc4e.2439.3728	
1/1/15	Down	None	None	None	None	No	1	0	cc4e.2439.372c	
1/1/16	Down	None	None	None	None	No	1	0	cc4e.2439.3730	
1/1/17	Down	None	None	None	None	No	1	0	cc4e.2439.3734	
1/1/18	Down	None	None	None	None	No	1	0	cc4e.2439.3735	
1/1/19	Down	None	None	None	None	No	1	0	cc4e.2439.3736	
1/1/20	Down	None	None	None	None	No	1	0	cc4e.2439.3737	
1/2/1	Down	None	None	None	None	No	1	0	cc4e.2439.3715	
1/2/2	Down	None	None	None	None	No	1	0	cc4e.2439.3719	
1/2/3	Down	None	None	None	None	No	1	0	cc4e.2439.371d	
1/2/4	Up	Forward	Full	40G	None	No	1	0	cc4e.2439.3721	
1/2/5	Down	None	None	None	None	No	1	0	cc4e.2439.3725	
1/2/6	Down	None	None	None	None	No	1	0	cc4e.2439.3729	
mgmt1	Up	None	Full	1G	None	No	None	0	cc4e.2439.3700	

The following example breaks out port 1/2/4.

```
device(config)# breakout ethernet 1/2/4
Reload required. Please write memory and then reload or power cycle.
device(config)# end
device# write memory
Write startup-config done.

device# Flash Memory Write (8192 bytes per dot) .
Copy Done.
device# reload
```

The following example shows that port 1/2/4 has been configured for breakout into four 10 Gbps sub-ports.

```
device# show interface brief
```

Port	Link	State	Dupl	Speed	Trunk	Tag	Pvid	Pri	MAC	Name
1/1/1	Down	None	None	None	None	No	1	0	cc4e.2439.3700	
1/1/2	Down	None	None	None	None	No	1	0	cc4e.2439.3701	
1/1/3	Down	None	None	None	None	No	1	0	cc4e.2439.3702	
1/1/4	Down	None	None	None	None	No	1	0	cc4e.2439.3703	
1/1/5	Down	None	None	None	None	No	1	0	cc4e.2439.3704	
1/1/6	Down	None	None	None	None	No	1	0	cc4e.2439.3708	
1/1/7	Down	None	None	None	None	No	1	0	cc4e.2439.370c	
1/1/8	Down	None	None	None	None	No	1	0	cc4e.2439.3710	
1/1/9	Down	None	None	None	None	No	1	0	cc4e.2439.3714	
1/1/10	Down	None	None	None	None	No	1	0	cc4e.2439.3718	
1/1/11	Down	None	None	None	None	No	1	0	cc4e.2439.371c	
1/1/12	Down	None	None	None	None	No	1	0	cc4e.2439.3720	
1/1/13	Down	None	None	None	None	No	1	0	cc4e.2439.3724	
1/1/14	Down	None	None	None	None	No	1	0	cc4e.2439.3728	
1/1/15	Down	None	None	None	None	No	1	0	cc4e.2439.372c	
1/1/16	Down	None	None	None	None	No	1	0	cc4e.2439.3730	
1/1/17	Down	None	None	None	None	No	1	0	cc4e.2439.3734	
1/1/18	Down	None	None	None	None	No	1	0	cc4e.2439.3735	
1/1/19	Down	None	None	None	None	No	1	0	cc4e.2439.3736	
1/1/20	Down	None	None	None	None	No	1	0	cc4e.2439.3737	

1/2/1	Down	None	None	None	None	No	1	0	cc4e.2439.3715
1/2/2	Down	None	None	None	None	No	1	0	cc4e.2439.3719
1/2/3	Down	None	None	None	None	No	1	0	cc4e.2439.371d
1/2/4:1	Up	Forward	Full	10G	None	No	1	0	cc4e.2439.3721
1/2/4:2	Up	Forward	Full	10G	None	No	1	0	cc4e.2439.3722
1/2/4:3	Up	Forward	Full	10G	None	No	1	0	cc4e.2439.3723
1/2/4:4	Up	Forward	Full	10G	None	No	1	0	cc4e.2439.3724
1/2/5	Down	None	None	None	None	No	1	0	cc4e.2439.3725
1/2/6	Down	None	None	None	None	No	1	0	cc4e.2439.3729
mgmt1	Up	None	Full	1G	None	No	None	0	cc4e.2439.3700

The following example configures names for port 1/2/4 sub-ports.

```

device> enable
device# configure terminal
device(config)# interface ethernet 1/2/4:1
device(config-if-e10000-1/2/2:1)# port-name subport1
device(config-if-e10000-1/2/2:1)# interface ethernet 1/2/4:2
device(config-if-e10000-1/2/2:2)# port-name subport2
device(config-if-e10000-1/2/2:2)# interface ethernet 1/2/4:3
device(config-if-e10000-1/2/2:3)# port-name subport3
device(config-if-e10000-1/2/2:3)# interface ethernet 1/2/4:4
device(config-if-e10000-1/2/2:4)# port-name subport4
device(config-if-e10000-1/2/2:4)# end
device(config)# end
device# end
device>
    
```

Removing breakout configuration

Use the **no breakout** command as described to remove 40 Gbps breakout configuration.

Removing 4X10 Gbps sub-ports and restoring the original 40 Gbps port requires the same steps as configuring breakout.

Enter the **no breakout** command for an individual port or port range as shown in the following examples. However, for the restored 40 Gbps port configuration to take effect, you must also execute the **write memory** command and then use the **reload** command to update the unit's configuration.

The following example checks for ports with active breakout configuration and then removes breakout from ports 1/3/1 through 1/3/6.

```
Device# show breakout

Unit-Id: 1

Port      Module Exist  Module Conf  breakout_conf  breakout_oper
1/1/5     Yes           No           Yes            Yes
1/1/6     Yes           No           Yes            Yes
1/1/7     Yes           No           Yes            Yes
1/1/8     Yes           No           Yes            Yes
1/1/9     Yes           No           Yes            Yes
1/1/10    Yes           No           Yes            Yes
1/1/11    Yes           No           Yes            Yes
1/1/12    Yes           No           Yes            Yes
1/1/13    Yes           No           Yes            Yes
1/1/14    Yes           No           Yes            Yes
1/1/15    Yes           No           Yes            Yes
1/1/16    Yes           No           Yes            Yes
1/2/1     Yes           No           Yes            Yes
1/2/2     Yes           No           Yes            Yes
1/2/3     Yes           No           Yes            Yes
1/2/4     Yes           No           Yes            Yes
1/2/5     Yes           No           Yes            Yes
1/2/6     Yes           No           Yes            Yes
1/3/1     Yes           No           Yes            Yes
1/3/2     Yes           No           Yes            Yes
1/3/3     Yes           No           Yes            Yes
1/3/4     Yes           No           Yes            Yes
1/3/5     Yes           No           Yes            Yes
1/3/6     Yes           No           Yes            Yes
```

```
Device# configure terminal
Device(config)# no breakout ethernet 1/3/1 to 1/3/6
Reload required. Please write memory and then reload or power cycle.
Device(config)# write memory
Write startup-config done.

Device(config)# Flash Memory Write (8192 bytes per dot) .
Copy Done.
Device(config)# end
Device# reload
```

NOTE

If there had been any configuration on any sub-ports (1/3/1:1 to 1/3/6:4), the **no breakout** command would have returned an error. The configuration would then have to be removed from the sub-ports before breakout configuration could be removed.

The following example shows a failed attempt to remove breakout from port 1/1/5 as indicated by the error message. Configuration is then removed from sub-port 1/1/5:1 before the breakout configuration is successfully removed.

Once the updated configuration is loaded, the ports are restored as full 40 Gbps ports. The former sub-port configuration is not retained in memory.

```
device(config)# no breakout ethernet 1/1/5
Error: Port 1/1/5:1 is tagged

device(config)# vlan 200
device(config-vlan-200)# no tagged ethernet 1/1/5:1
Deleted tagged port(s) to port-vlan 200.
device(config)# end
device# configure terminal
device(config)# no breakout ethernet 1/1/5
Reload required. Please write memory and then reload or power cycle.
device(config)# end
device# write memory
Write startup-config done.
```

```
device# Flash Memory Write (8192 bytes per dot) .  
Copy Done.
```

Displaying information for breakout ports

Use the **show breakout** command to display breakout port status.

The **show breakout** command indicates which ports are configured for breakout and which breakout ports are in operation. The command also displays ports that have been configured for breakout but that are not yet broken out into sub-ports, pending reload.

The following example displays breakout port information for an ICX 7750-48F. Port 1/2/1 is the only port with active sub-ports; however, ports 1/2/2 and 1/2/4 are configured for breakout, pending reload.

```
Device# show breakout  
Unit-Id: 1  
Port      Module Exist  Module Conf  Breakout-config  Breakout-oper  
1/2/1     yes           no           yes              yes  
1/2/2     yes           no           yes              no  
1/2/3     yes           no           no               no  
1/2/4     yes           no           yes              no  
1/2/5     yes           no           no               no  
1/2/6     yes           no           no               no  
1/3/1     yes           no           no               no  
1/3/2     yes           no           no               no  
1/3/3     yes           no           no               no  
1/3/4     yes           no           no               no  
1/3/5     yes           no           no               no  
1/3/6     yes           no           no               no
```


OpenSSL License

- [OpenSSL license](#)..... 353

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 .
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.

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 Netscapes 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 and/or other materials provided with the distribution. All advertising materials mentioning features or use of this software must display the following acknowledgment: "This product includes cryptographic software written by Eric Young(eyay@cryptsoft.com)" The word 'cryptographic' can be left out if the routines from the library being used are not cryptographic related.
4. If you include any Windows specific code (or a derivative thereof) from the apps directory (application code) you must include an acknowledgment: "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 licence and distribution terms for any publically available version or derivative of this code cannot be changed. i.e. this code cannot simply be copied and put under another distribution licence.

Joint Interoperability Test Command

- [JITC overview](#)..... 355

JITC overview

The Joint Interoperability Test Command (JITC) mode on a FastIron device is compliant with the standards established by JITC, a United States military organization that tests technology pertaining to multiple branches of the armed services and the government.

The JITC mode implemented on a FastIron device enforces default behavior for some features to ensure strict JITC certification compliance.

AES-CTR encryption mode support for SSH

The Advanced Encryption Standard - Cipher Block Chaining (AES-CBC) encryption mode for Secure Shell (SSH) is vulnerable to certain plain-text attacks. The JITC mode uses AES-CTR (Counter) encryption mode for SSH instead of AES-CBC mode for enhanced security.

In the JITC mode, by default, the AES-CBC encryption mode for SSH is disabled and the AES-CTR (Counter) encryption mode is enabled. The **ip ssh encryption disable-aes-cbc** command that disables the AES-CBC mode can be seen in the running configuration. The encryption algorithms such as aes256-ctr, aes192-ctr, or aes128-ctr are enabled and the CBC mode ciphers are removed.

The AES-CBC mode can be re-enabled by issuing the **no ip ssh encryption disable-aes-cbc** command, which will bring back the pre-existing CBC ciphers (aes256-cbc, aes192-cbc, aes128-cbc, and 3des-cbc) along with the CTR ciphers.

NOTE

The AES-CTR mode must be configured both on the client and server sides to establish an SSH connection.

SHA1 authentication support for NTP

In the JITC mode, the symmetric key scheme supported for cryptographic authentication of messages uses the SHA1 keyed hash algorithm instead of the MD5 authentication scheme. The MD5 authentication for Network Time Protocol (NTP) is disabled by default in the JITC mode and the **disable authentication md5** command can be seen in the running configuration. Only the SHA1 authentication scheme is available to define the authentication key for NTP in the JITC mode. SHA1 authentication must be enabled manually using the **authentication-key key-id** command. In the JITC mode, only the SHA1 option is available.

The MD5 authentication scheme can be re-enabled by issuing the **no disable authentication md5** command. By doing so, the default JITC mode behavior is overridden.

IPv6 ACL for SNMPv3 group

As part of the JITC requirement, from 08.0.20a release onwards, the IPv6 access list is supported for the SNMPv3 group, and the incoming SNMP packets can be filtered based on the IPv6 ACL attached to the group.

For more information, refer to [Defining an SNMP group](#) on page 153 and [Defining an SNMP group and specifying which view is notified of traps](#) on page 157.