



Emulex[®] Drivers for Linux for LightPulse[®] Adapters

User Guide
Release 12.0

Broadcom, the pulse logo, Connecting everything, Avago Technologies, Avago, the A logo, Emulex, LightPulse, OneCommand, OneConnect, and ExpressLane are among the trademarks of Broadcom and/or its affiliates in the United States, certain other countries, and/or the EU.

Copyright © 2003–2018 Broadcom. All Rights Reserved.

The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries. For more information, please visit www.broadcom.com.

Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design. Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.



Table of Contents

Chapter 1: Introduction	6
1.1 Overview	6
1.2 Abbreviations	6
Chapter 1: Installing and Uninstalling	9
1.1 General Installation Requirements	9
1.2 Binary RPM FC and FCoE Driver Kit	9
1.2.1 Installing the Binary RPM FC and FCoE Driver Kit	9
1.2.2 Uninstalling the Binary RPM FC and FCoE Driver Kit	10
1.3 Ethernet Driver Kit for PPC CNAs	10
1.3.1 Installing the Ethernet Driver Kit	10
1.3.2 Uninstalling the Ethernet Driver Kit	11
1.4 Booting from a Nonzero LUN Attached to an Emulex FC Adapter	11
1.5 Emulex OneCommand Manager Application	11
Chapter 2: Configuration	12
2.1 FC and FCoE Driver Configuration	12
2.1.1 FC and FCoE Driver Parameters	12
2.1.1.1 Static FC and FCoE Driver Parameters	12
2.1.1.2 Dynamic FC and FCoE Driver Parameters	16
2.1.2 Configuring FC and FCoE Driver Parameters	18
2.1.2.1 Temporary Configuration with modprobe	19
2.1.2.2 Persistent Configuration with the lpfc.conf file	19
2.1.2.3 Configure Parameters with a Read/Write to sysfs	19
2.1.3 Creating a New Ramdisk Image	20
2.1.4 Dynamically Recognizing LUNs and Targets (Using scan)	21
2.1.5 Persistent Naming	21
2.1.5.1 Using udev to Discover Logical to Physical Mappings for sd Devices	21
2.1.5.2 Configuring the System to Boot Using Persistent Names	21
2.1.5.3 Using udev with st Devices	22
2.1.5.4 Persistent Naming References	24
2.1.6 ExpressLane Support	24
2.2 vPort Configuration	24
2.2.1 vPort Configuration Prerequisites	25
2.2.2 Creating, Deleting, and Displaying vPorts	25
2.2.2.1 Creating vPorts Using the mkvport.sh Script	25
2.2.2.2 Deleting vPorts Using the rmvport.sh Script	25
2.2.2.3 Displaying vPorts Using the lsvport.sh Script	26
2.2.3 vPort sysfs	26

2.2.3.1	vPort sysfs Tree	26
2.2.3.2	vPort sysfs Entries	27
2.2.4	Monitoring vPorts with fc_vport	29
2.2.5	vPort Configuration Limits	29
2.3	Ethernet Driver Configuration for PPC CNAs	30
2.3.1	Ethernet Driver Configuration Parameters	30
2.3.2	Support for Ethtool set-channels/get-channels Commands	30
2.3.3	Support for Ethtool set-dump Command	31
2.3.4	Transmit/Receive Queue Counts	31
2.3.5	Support for Ethtool set-rxfh/get-rxfh Commands	32
2.3.6	Support for Showing Onboard Die Temperature	32
2.3.7	SR-IOV Configuration	32
2.3.7.1	Introduction	32
2.3.7.2	Setting Up SR-IOV	33
2.3.7.3	Assigning VFs to a VM on the SLES Operating System	36
2.3.7.4	Link State Reporting with SR-IOV	37
2.3.7.5	Configuring VFs	37
2.3.7.6	Link State Configuration	39
2.3.7.7	Spoof Check Configuration	39
2.3.7.8	Viewing VF Properties	39
2.3.8	Bonding Considerations	39
2.4	FC and FCoE Driver Performance Tuning	40
2.4.1	Overview	40
2.4.1.1	lpfc_fcp_io_channel	40
2.4.1.2	lpfc_fcp_io_sched	40
2.4.1.3	lpfc_fcp_imax	41
2.4.1.4	lpfc_vector_map.sh	41
2.5	Network Performance Tuning	42
2.5.1	Memory Bandwidth Considerations	42
2.5.1.1	Enabling Optimal Bandwidth Options	42
2.5.1.2	Populating DIMM Slots	42
2.5.1.3	Disabling Memory Mirroring	42
2.5.1.4	Using a Fast Clock Speed for the Front Side Bus (FSB)	42
2.5.2	Network Memory Limits	42
2.5.3	TCP Segmentation Offload (TSO)	43
2.5.4	Flow Control	43
2.5.5	RX Frame Coalescing	44
2.5.6	Maximum Transmission Unit (MTU)	44
2.5.7	Interrupt Coalescing	45
2.5.7.1	Setting the Interrupt Delay Duration to a Range of Values (AIC)	45

2.5.7.2 Setting the Interrupt Delay Duration to a Constant Value	45
2.5.8 Receive-Side Scaling (RSS)	45
2.5.9 Analyzing Performance Issues	46
Chapter 3: Troubleshooting	47
3.1 Hardware Situations and Resolutions	47
3.2 FC and FCoE Driver Situations and Their Resolutions	47
3.3 Ethernet Driver Situations and Their Resolutions	50
3.4 Log Messages	50
3.4.1 FC and FCoE Driver Log Messages	50
3.4.1.1 Retrieving FC and FCoE Driver Log Messages	50
3.4.1.2 LPFC Error Log Messages and their Descriptions	51
3.4.2 Ethernet Driver Log Messages for NIC PPC CNAs	122
3.4.2.1 Retrieving Ethernet Driver Log Messages	122
3.4.2.2 Ethernet Driver Log Messages and Their Descriptions	122
Appendix A: Ethtool -S Option Statistics for PPC CNAs	127
Appendix B: Ipfc Driver BlockGuard Functionality	129
B.1 Overview	129
B.2 Enabling BlockGuard	130
B.3 SCSI Command Support – SCSI Operation Codes	130
B.4 Ipfc Driver Module Parameters	131
B.4.1 lpfc_prot_mask	131
B.4.2 lpfc_prot_guard	132

Chapter 1: Introduction

1.1 Overview

This guide provides installation, updating, uninstalling, configuring, and troubleshooting procedures for several types of Emulex®-supported drivers for Linux.

This guide is applicable to several versions of drivers, operating systems, firmware, and adapters.

- For supported firmware versions and their latest release, refer to the Documents and Downloads area of www.broadcom.com for the specific adapter.

This product supports the following adapters:

- Emulex LightPulse® Fibre Channel (FC) host bus adapters (HBAs):
 - LPe12000-series adapters
 - LPe16000-series adapters
 - LPe31000-series adapters
 - LPe32000-series adapters
- IBM PowerPC (PPC) converged network adapters (CNAs):
 - Network interface cards (NIC)
 - FC over Ethernet + NIC (FCoE+NIC)

For a list of Emulex adapters that are compatible with the stand-alone driver kits, refer to the Documents and Downloads area of www.broadcom.com for the specific driver.

1.2 Abbreviations

The following table lists the acronyms and abbreviations used in this document.

Table 1: Acronyms and Abbreviations

Acronym/Abbreviation	Description
AIC	adaptive interrupt coalescing
ANSI	American National Standards Institute
API	application programming interface
ASIC	application-specific integrated circuit
BIOS	basic input/output system
CLI	command line interface
CNA	converged network adapter
CPU	central processing unit
CQ	completion queue
CQE	completion queue entry
DCBX	Data Center Bridging Capabilities Exchange
DIF	Data Integrity Field
DIMM	dual in-line memory module
DMA	direct memory access
ELS	Extended Link Service

Table 1: Acronyms and Abbreviations (Continued)

Acronym/Abbreviation	Description
EQ	event queue
FC	Fibre Channel
FCF	Fibre Channel over Ethernet Forwarder
FCoE	Fibre Channel over Ethernet
FCFI	Fibre Channel Forwarder Indicator
FCP	Fibre Channel Protocol
FSB	front side bus
FW	firmware
Gb/s	gigabits per second
GnuPG	GNU Privacy Guard
GRO	Generic Receive Offload
GUI	Graphical User Interface
HBA	host bus adapter
I/O	input/output
IOCB	input/output control block
IOV	I/O virtualization
IP	Internet Protocol
IRQ	interrupt request
KVM	kernel-based virtual machine
LAN	local area network
LBA	logical block address
LLC	logical link control
LOM	LAN on motherboard
LRO	large receive offload
LUN	logical unit number
MAC	Media Access Control
Mb/s	megabits per second
MDS	Multilayer Director Switch
MSI	message signaled interrupts
MSI-X	message signaled interrupts – extended
MTU	maximum transmission unit
NAA	Network Address Authority
NIC	network interface card (or controller)
NPIV	N_Port ID Virtualization
NUMA	non-uniform memory access
PCI	Peripheral Component Interconnect
PCIe	Peripheral Component Interconnect Express
PF	physical function
POST	power-on self-test
PPC	PowerPC
QoS	quality of service
RHEL	Red Hat Enterprise Linux
RPI	remote port indicator

Table 1: Acronyms and Abbreviations (Continued)

Acronym/Abbreviation	Description
RPM	resource package manager
RQ	receive queue
RSCN	registered state change notification
RSS	receive-side scaling
Rx	receive
SAN	storage area network
SCSI	Small Computer System Interface
SGE	Oracle (formerly Sun) grid engine
SLES	SUSE Linux Enterprise Server
SLI	Service Level Interface
SNMP	Simple Network Management Protocol
SR-IOV	single-root I/O virtualization
tar	tape archive
TCP	Transmission Control Protocol
TSO	TCP Segmentation Offload
Tx	transmit
UDP	User Datagram Protocol
UEFI	Unified Extensible Firmware Interface
VEB	virtual Ethernet bridging
VEPA	virtual Ethernet port aggregator
VF	virtual function
VGT	virtual guest tagging
vPort	virtual port
VLAN	virtual local area network
VM	virtual machine
VPD	vital product data
WWPN	World Wide Port Name
XRI	extensible resource indicator

Chapter 1: Installing and Uninstalling

The Linux driver supports both FCP and NVMe over Fibre Channel. For NVMe support, refer to the *Emulex NVMe over Fibre Channel User Guide*.

Emulex releases Linux binary RPMs that are digitally signed using the GNU Privacy Guard (GnuPG) standard. This will allow certification of the contents of the RPMs and verification that the contents have not been modified since they were created by Emulex. The RPMs have been digitally signed by Emulex with a GnuPG private key that is only held by Emulex. Instructions for creating the GnuPG public key file are located at:

docs.broadcom.com/docs/1334566

1.1 General Installation Requirements

NOTE: Update the latest driver on the operating system before updating the firmware.

Prior to driver installation, follow these general requirements:

- Install a supported Emulex adapter in the system. Refer to the adapter's installation guide for specific hardware installation instructions.
- Use a supported operating system. Refer to the Documents and Downloads area of www.broadcom.com for supported operating systems.

1.2 Binary RPM FC and FCoE Driver Kit

The binary RPM FC and FCoE driver kit contains the following:

- A zipped tar file that includes the driver binary RPMs for a specific driver version and Linux distribution.

NOTE: Use only officially released Linux distribution kernels. The binary RPM packages support only officially released Linux distribution kernels and do not support pre-release distribution kernels.

- An installation script, `elx_lpfcd_install.sh`, that installs the FC and FCoE driver binary RPM that corresponds to the target system's architecture and kernel memory variant.
- A `README` file that provides a description of the kit structure, its contents, and distribution support scope.

1.2.1 Installing the Binary RPM FC and FCoE Driver Kit

NOTE: You must uninstall any FC driver kits that are not part of this distribution. For example, you must uninstall any previous FC and FCoE driver kits before installing this driver kit. This installation fails if a previous version of the FC and FCoE driver kit is detected. For more information, see [Section 1.2.2, Uninstalling the Binary RPM FC and FCoE Driver Kit](#).

To install the binary RPM FC and FCoE driver:

1. Download the appropriate driver kit from the Documents and Downloads area of www.broadcom.com.
2. Log in as `root` to a terminal, and unpack the tarball:

```
tar xzf elx-lpfc-dd-<Linux distribution version>-<driver version>.tar.gz
```
3. Change to the directory into which the tarball was extracted:

```
cd elx-lpfc-dd-<Linux distribution version>-<driver version>/
```

4. Run the `elx_lpfc_install.sh` script without options to install the driver kit:

```
./elx_lpfc_install.sh
```

After the `elx_lpfc_install.sh` script has finished running successfully, the Emulex FC and FCoE driver is loaded, and devices that are connected to the system are accessible.

5. Reboot the system to enable the newly added driver options in the ramdisk. You can also reboot the system later.
6. You can use the following command to verify if the version of the newly installed LPFC driver is correct:

```
#cat /sys/module/lpfc/version
```

1.2.2 Uninstalling the Binary RPM FC and FCoE Driver Kit

NOTE: You must run the uninstall script that shipped with the version of the driver kit you want to remove.

To uninstall the binary RPM FC and FCoE driver:

1. Log in as *root*.
2. If possible, exit all applications that use FC-attached drives, and then unmount the drives. If you cannot exit all applications that use FC-attached drives, the uninstall script works properly, but you must reboot after the uninstall is complete.
3. Run the `elx_lpfc_install.sh` script with the `--uninstall` option:

```
./elx_lpfc_install.sh --uninstall
```

1.3 Ethernet Driver Kit for PPC CNAs

The Ethernet driver kit includes the driver that supports the NIC protocol. The Ethernet driver kit contains the following:

- A zipped tar file that includes the driver binary RPMs for a specific driver version, and for all of the supported Linux distribution kernels.

NOTE: Use only officially released Linux distribution kernels. The binary RPM packages support only officially released Linux distribution kernels, and do not support pre-release distribution kernels.

- An installation script, `lpnic_install.sh`, which installs (by default) the Ethernet driver binary RPM that corresponds to the target system's architecture and kernel memory variant.
- A README file that provides a description of the kit structure, its contents, and distribution support scope.

1.3.1 Installing the Ethernet Driver Kit

NOTE: Remove any previously installed Ethernet driver kits (that is, those that were not part of a distribution's kernel), before proceeding. For more information on uninstalling the driver, see [Section 1.3.2, Uninstalling the Ethernet Driver Kit](#).

To install the Ethernet driver:

1. Download the appropriate driver kit from the Documents and Downloads area of www.broadcom.com.
2. Log in as *root* to a terminal, and unpack the tarball:

```
tar xzf elx-lpnic-dd-<driver version>.tar.gz
```
3. Change to the directory that is extracted:

```
cd elx-lpnic-dd-<driver version>/
```
4. Run the `lpnic_install.sh` script without options to install the driver kit:

```
./lpnic_install.sh
```

After the `lpnic_install.sh` script has finished running successfully, the Emulex Ethernet driver is loaded, and devices that are connected to the system are accessible.

5. Reboot the system to enable the newly added driver options in the ramdisk. You can also reboot the system later.
6. You can use the following command to verify if the version of the newly installed LPNIC driver is correct:

```
#cat /sys/modules/lpnic/version
```

1.3.2 Uninstalling the Ethernet Driver Kit

NOTE: You must run the uninstall script that shipped with the version of the driver kit you want to remove.

To uninstall the Ethernet driver:

1. Log in as *root*.
2. If possible, exit all applications that use Ethernet-attached drives, and then unmount the drives. If you cannot exit all applications that use Ethernet-attached drives, the uninstall works properly, but you must reboot after the uninstall is complete.
3. Run the `lpnic_install.sh` script with the `--uninstall` option:

```
./lpnic_install.sh --uninstall
```

1.4 Booting from a Nonzero LUN Attached to an Emulex FC Adapter

To configure SLES 11 SPx to boot from an FC-attached disk device other than a LUN 0 device, refer to the *Boot for the Fibre Channel Protocol User Guide*, which is available at the Documents and Downloads area of www.broadcom.com.

1.5 Emulex OneCommand Manager Application

The Emulex OneCommand[®] Manager application is a powerful and centralized adapter management suite. It provides discovery, reporting, and management of local and remote adapters from a single console anywhere in the SAN and across platforms. Both a GUI and CLI are provided. For instructions on installing and using the Emulex OneCommand Manager application, refer to the *OneCommand Manager Application for LightPulse Adapters User Guide*, which is available at the Documents and Downloads area of www.broadcom.com.

Chapter 2: Configuration

2.1 FC and FCoE Driver Configuration

The following section describes how to configure parameters for the FC and FCoE driver.

2.1.1 FC and FCoE Driver Parameters

The FC and FCoE driver parameters determine some aspects of the driver's behavior. There are two main types, static and dynamic. Changes to the static parameters require a driver reload for the change to take effect. Changes to most dynamic parameters take effect immediately; some do not take effect until there is a link-down/link-up sequence.

For more information on driver parameters, see [Section 2.1.1.1, Static FC and FCoE Driver Parameters](#), and [Section 2.1.1.2, Dynamic FC and FCoE Driver Parameters](#).

2.1.1.1 Static FC and FCoE Driver Parameters

Changes to static parameters require a driver reload for the change to take effect. The following table lists the static FC and FCoE driver parameters.

Table 1: Static FC and FCoE Driver Parameters

Parameter	Description	sysfs Visible	Applicable Adapters	Activation
lpfc_ack0	When enabled, ACK0 is used for Class 2. The enabled value is 1. The disabled value is 0 (default).	Yes	SLI-3 only	Driver reload
lpfc_discovery_threads	Specifies the maximum number of ELS commands that can be outstanding for a discovery. NOTE: The <code>lpfc_discovery_threads</code> parameter defaults to a value of 64 for private loop topologies regardless of the configured value. If multiple ports are configured on the host, the value of 64 is used only for those ports that are connected in a private loop topology. The configured value is used for all other ports. The minimum value is 1. The maximum value is 64. The default value is 32.	No	All	Driver reload
lpfc_enable_bg	Enables the BlockGuard (T10-DIF) feature. The minimum value is 0 (default). The maximum value is 1.	Yes	All	Driver reload
lpfc_enable_da_id	When enabled, the FC and FCoE driver issues a <code>DA_ID CT</code> command to the fabric when vPorts log out of the fabric. The enabled value is 1. The disabled value is 0 (default).	No	All	Driver reload

Table 1: Static FC and FCoE Driver Parameters (Continued)

Parameter	Description	sysfs Visible	Applicable Adapters	Activation
lpfc_enable_fc4_type	When enabled, this parameter defines the FC4 types that are supported. The possible values are: <ul style="list-style-type: none"> ■ 1 = Enable just FCP (default). ■ 3 = Enable both FCP and NVMe Supported values are 1 or 3. Default value is 1 NOTE: For additional information, refer to the <i>Emulex NVMe over Fibre Channel User Guide</i> .	Yes	LPe31000-series and LPe32000-series	Driver reload
lpfc_enable_hba_heartbeat	When enabled, the heartbeat logic in the FC and FCoE driver is able to detect whether the adapter is functional. If the heartbeat logic detects that the adapter is not functional, the driver will shut down the adapter. The disabled value is 0 (default). The enabled value is 1.	Yes	All	Driver reload
lpfc_enable_hba_reset	When enabled, the FC and FCoE driver can pass resets to the adapter. This action is typically used for debugging purposes. The enabled value is 1 (default). The disabled value is 0.	Yes	All	Driver reload
lpfc_enable_npiv	When enabled, the FC and FCoE driver can use NPIV to create vPorts (if supported by the fabric). The enabled value is 1 (default). The disabled value is 0.	Yes	All	Driver reload
lpfc_enable_SmartSAN	Sets up FDMI support for SmartSAN. The disabled value is 0 (default). The enabled value is 1.	Yes	All	Driver reload
lpfc_EnableXLane	For Emulex LPe16000-series, LPe31000-series, and LPe32000-series adapters, this parameter specifies whether the ExpressLane™ feature is enabled or disabled. The minimum value is 0 (default). The maximum value is 1.	Yes	SLI-4 only	Driver reload
lpfc_fcp_class	Specifies either FC Class 2 or 3 for FCP data transmission. For Class 2, the value is 2. For Class 3, the value is 3 (default).	Yes	All	Driver reload
lpfc_fcp_io_channel	For Emulex LPe16000-series, LPe31000-series, and LPe32000-series adapters, this parameter defines the number of I/O channels supported by the driver. For more information, see Section 2.4, FC and FCoE Driver Performance Tuning . The default value is 4 I/O channels. The minimum value is 1 I/O channel. The maximum value is 16 I/O channels. For Emulex LPe12000-series adapters, this parameter is not applicable and has no effect.	Yes	SLI-4 only	Driver reload

Table 1: Static FC and FCoE Driver Parameters (Continued)

Parameter	Description	sysfs Visible	Applicable Adapters	Activation
lpfc_fdmi_on	Specifies if FDMI support is enabled or disabled. The enabled value is 1 (default). The disabled value is 0.	Yes	All	Driver reload
lpfc_hba_queue_depth	Limits the number of outstanding commands per HBA. The minimum value is 32. The maximum value is 8192 (also default).	Yes	All	Driver reload
lpfc_lun_queue_depth	Specifies the default maximum number of commands sent to a single logical unit (disk drive). The minimum value is 1. The maximum value is 512. The default value is 30. The value of <code>lpfc_lun_queue_depth</code> is dependent on the <code>max_xri</code> value received. The <code>max_xri</code> value might vary per adapter type.	Yes	All	Driver reload
lpfc_max_luns	Specifies the highest available LUN ID that is valid, per target. For example, a value of 19 means that LUN IDs from 0 to 19 are valid for the target. The SCSI layer scans each target until it reaches this specified LUN ID. The minimum value is 0. The maximum value is 65535. The default value is 255.	Yes	All	Driver reload
lpfc_max_scsicmpl_time	Uses command completion time to control queue depth. The units are in milliseconds. The minimum value is 0 (default). The maximum value is 6000.	Yes	All	Driver reload
lpfc_restrict_login	When enabled, restricts vPorts login to remote initiators. The enabled value is 1 (default). The disabled value is 0.	No	All	Driver reload
lpfc_scan_down	When enabled, selects the <i>scan down</i> method (scanning the AL_PA from high to low) to assign a SCSI ID. The enabled value is 1 (default). The disabled value is 0.	Yes	FC adapters only	Driver reload
lpfc_sg_seg_cnt	Controls the scatter/gather maximum segment count passed to the FC and FCoE driver. This variable is applicable per SCSI command. For Emulex LPe12000-series adapters, the minimum value is 64 (default), and the maximum value is 4096. For all other adapters, the minimum value is 64 (default), and the maximum value is 510.	Yes	All	Driver reload

Table 1: Static FC and FCoE Driver Parameters (Continued)

Parameter	Description	sysfs Visible	Applicable Adapters	Activation
lpfc_sli_mode	<p>For Emulex LPe12000-series adapters, this parameter allows you to force the SLI mode requested by the adapter driver. The possible values are:</p> <ul style="list-style-type: none"> ■ 0 = Auto-select (default). ■ 2 = SLI-2. ■ 3 = SLI-3. <p>NOTE: For all other adapters, this parameter is not applicable and has no effect.</p>	No	SLI-3 only	Driver reload
lpfc_use_msi	<p>When enabled, determines whether the driver uses MSI or MSI-X.</p> <ul style="list-style-type: none"> ■ 0 = MSI disabled; INTx mode is used . ■ 1 = MSI; allows a maximum of 32 interrupts. ■ 2 = MSI-X; allows a maximum of 2048 interrupts (default value). <p>NOTE: The default is 2. This value reverts to 1 if the system does not support MSI-X. This value reverts to 0 if the system does not support MSI.</p>	Yes	All	Driver reload
lpfc_xri_split	<p>When enabled, this parameter defines the division of XRI resources between SCSI and NVMe.</p> <p>This parameter is only used if the value of the <code>lpfc_enable_fc4_type</code> parameter is 3 - register both FCP and NVMe and the port is not configured for NVMET.</p> <p>Supported values are in percentages. The <code>lpfc_xri_split</code> value is the percentage of XRI resources allocated to the SCSI port. The remaining percentage of XRI resources is allocated to NVMe.</p> <p>Supported range is 10% to 90%, Default value is 50%.</p> <p>For additional information, refer to the <i>Emulex NVMe over Fibre Channel User Guide</i>.</p>	Yes	LPe31000-series and LPe32000-series	Driver reload

2.1.1.2 Dynamic FC and FCoE Driver Parameters

Changes to the dynamic parameters take effect immediately. All `lpfc` dynamic parameters are read/write using `sysfs`. The following table lists the dynamic FC and FCoE driver parameters.

Table 2: Dynamic FC and FCoE Driver Parameters

Parameter	Description	<code>sysfs</code> Visible	Applicable Adapters	Activation
<code>lpfc_cr_count</code>	For Emulex LPe12000-series adapters, this parameter determines the value for I/O coalescing for <code>lpfc_cr_count</code> outstanding commands. The minimum value is 1 (default). The maximum value is 255. NOTE: For all other adapters, this parameter is not applicable and has no effect.	Yes	SLI-3 only	Port reset
<code>lpfc_cr_delay</code>	For Emulex LPe12000-series adapters, this parameter determines the value for I/O coalescing for <code>lpfc_cr_delay</code> (milliseconds) outstanding commands. The minimum value is 0 (default). The maximum value is 63. NOTE: For all other adapters, this parameter is not applicable and has no effect.	Yes	SLI-3 only	Port reset
<code>lpfc_devloss_tmo</code>	Specifies the number of seconds to hold an I/O error when a device disappears. The minimum value is 1. The maximum value is 255. The default value is 30.	Yes	All	Dynamic
<code>lpfc_enable_bbcr</code>	For Emulex LPe31000-series and LPe32000-series adapters, this parameter specifies if the Buffer to Buffer Credit Recovery feature is enabled or disabled. The minimum value is 0. The maximum value is 1 (default).	Yes	SLI-4 only	Port reset
<code>lpfc_enable_mds_diags</code>	For Emulex LPe16000-series, LPe31000-series, and LPe32000-series adapters, this parameter enables or disables the Multilayer Director Switch (MDS) Diagnostics feature. The minimum value is 0 (default). The maximum value is 1. NOTE: The parameter should be disabled (set to 0) after diagnostics are complete.	Yes	SLI-4 only	Port reset
<code>lpfc_enable_auth</code>	Enables or disables DH-CHAP authentication. The possible values are: 0 = Authentication is disabled (default). 1 = Authentication is enabled.	Yes		

Table 2: Dynamic FC and FCoE Driver Parameters (Continued)

Parameter	Description	sysfs Visible	Applicable Adapters	Activation
lpfc_fcp_imax	For Emulex LPe16000-series, LPe31000-series, and LPe32000-series adapters, this parameter specifies the maximum number of fast-path FCP interrupts per second. The minimum value is 0 or OFF and 5000. The maximum value is 5,000,000. The default value is 150,000. NOTE: For Emulex LPe12000 adapters, this parameter is not applicable and has no effect.	Yes	SLI-4 only	Dynamic
lpfc_fcp_io_sched	For Emulex LPe16000-series, LPe31000-series, and LPe32000-series adapters, this parameter determines which algorithm to use when scheduling an FC or FCoE I/O to an I/O channel. For more information, see Section 2.4, FC and FCoE Driver Performance Tuning . The default value is 0, configuration by round-robin scheduling. A value of 1 sets configuration to CPU scheduling. For Emulex LPe12000-series adapters, this parameter is not applicable and has no effect.	Yes	SLI-4 only	Dynamic
lpfc_first_burst_size	Configures first burst size on targets that support first burst. The minimum value is 0 (default). The maximum value is 65536.	Yes	All	Dynamic
lpfc_link_speed	Specifies the FC link speed. The possible values are: <ul style="list-style-type: none"> ■ 0 = Auto-select (default). ■ 2 = 2Gb/s. ■ 4 = 4Gb/s. ■ 8 = 8Gb/s. ■ 16 = 16Gb/s. ■ 32 = 32Gb/s. NOTE: For LPe32000-series adapters running 11.2 firmware or later, this parameter is ignored. NOTE: Setting this option incorrectly can cause the adapter to fail to initialize. NOTE: If you configure the link speed in a BIOS utility, the link speed can be overwritten by the operating system according to its own configuration settings. To avoid this issue, configure the link speed in both the operating system driver and the Boot BIOS or UEFI driver.	Yes	FC adapters only	Port reset
lpfc_log_verbose	Specifies the log verbosity level of the messages posted by the driver. Extra activity logging (bit mask). The minimum value is 0x0 (default). The maximum value is 0xFFFFFFFF.	Yes	All	Dynamic
lpfc_req_fw_upgrade	Set to 1 to update firmware.	Yes		Dynamic

Table 2: Dynamic FC and FCoE Driver Parameters (Continued)

Parameter	Description	sysfs Visible	Applicable Adapters	Activation
lpfc_tgt_queue_depth	Limits the number of outstanding commands per target port. The minimum value is 10. The maximum value is 65535 (also default).	Yes	All	Dynamic
lpfc_throttle_log_cnt	Specifies the number of messages logged within throttle_log_time. The minimum value is 1. The maximum value is 1000. The default value is 10.	Yes	All	Dynamic
lpfc_throttle_log_time	Specifies the time limit for throttle_log_cnt. The minimum value is 1. The maximum value is 60. The default value is 1.	Yes	All	Dynamic
lpfc_topology	For FC adapters, this parameter sets the link topology. The possible values are: <ul style="list-style-type: none"> ■ 0x0 = Loop first; if loop fails, then point-to-point (default). ■ 0x2 = Point-to-point only. ■ 0x4 = Loop only. ■ 0x6 = Point-to-point first; if point-to-point fails, then loop. NOTE: For PPC FCoE+NIC CNAs, this parameter is not applicable and has no effect.	Yes	FC adapters only	Port reset
lpfc_use_adisc	When enabled, an ADISC is sent instead of a PLOGI for device discovery or RSCN. The enabled value is 1. The default value is 0 (disabled).	Yes	All	Dynamic
lpfc_XLanePriority	For Emulex LPe16000-series, LPe31000-series, and LPe32000-series adapters, when enabled, this parameter defines the frame priority to be used for the ExpressLane feature. The minimum value is 0 (default). The maximum value is 0x7f.	Yes		Dynamic

2.1.2 Configuring FC and FCoE Driver Parameters

You can configure the FC and FCoE driver parameters using any of the following methods:

- The `modprobe` Linux program for temporary configuration.
- The `lpfc.conf` file for persistent configuration.
- The `sysfs` interface (to view and modify parameters after loading the FC and FCoE driver).

NOTE: Not all parameters visible in the `sysfs` directory can be modified; some are read-only.

- The Emulex OneCommand Manager application (refer to the *OneCommand Manager Application for LightPulse Adapters User Guide* for more information).

NOTE: The FC and FCoE driver parameter changes made using the Emulex OneCommand Manager application remain after the FC and FCoE driver is uninstalled. To return to the default settings, you must reset them using the OneCommand Manager application.

2.1.2.1 Temporary Configuration with modprobe

When you manually load the FC and FCoE driver as a module using the `modprobe` command, and you change one or more driver parameter values in the command line, the configuration is temporary. These changes are considered temporary because they are valid only for the current session or until the FC and FCoE driver is unloaded.

The `modprobe` program uses the `lpfc.conf` file, but parameters passed to it using the command line override the parameters in the `lpfc.conf` file. Values can be expressed in hexadecimal or decimal notation.

If you want to temporarily set `lun_queue_depth` to 20 (the default is 30) for all HBAs in your system, load the FC and FCoE driver with the following command:

```
modprobe lpfc lpfc_lun_queue_depth=20
```

2.1.2.2 Persistent Configuration with the lpfc.conf file

To make the FC and FCoE driver parameters persist across module loads and reboots, perform these steps:

1. In the `/etc/modprobe.d` directory, create a file with the driver name `lpfc.conf`.
2. In `/etc/modprobe.d/lpfc.conf`, use the `options` command to add the appropriate FC and FCoE driver parameters and their desired values. For example, adding the following command to the `lpfc.conf` file sets the verbose flag:

```
options lpfc lpfc_log_verbose=0x3ffff
```

If driver parameters are added to the `lpfc.conf` file, the FC and FCoE driver must be reloaded for the parameters to take effect. Also, a new ramdisk image is required if you want the changes to take effect in the next boot. For information on creating a new ramdisk, see [Section 2.1.3, Creating a New Ramdisk Image](#).

If the same parameter is specified on the `modprobe` command line and in the `lpfc.conf` file, the value specified in the `modprobe` command line takes precedence.

2.1.2.3 Configure Parameters with a Read/Write to sysfs

`sysfs` is a virtual file system that exposes the structure of the system. It also includes interfaces to driver parameters through which the FC and FCoE driver parameters can be viewed and modified. Because these interfaces are available only after driver load, only dynamic FC and FCoE driver parameters can be changed. However, both static and dynamic FC and FCoE driver parameters can be read through `sysfs`.

NOTE:

- `sysfs` changes exist only during driver load and are lost when the FC and FCoE driver is unloaded or the system is rebooted.
- Driver parameters that are set through module parameters are global; setting them through `sysfs` is on a SCSI host (adapter port) basis.

For example:

```
echo 0x7f >> /sys/class/scsi_host/host7/lpfc_log_verbose only affects host 7.
```

```
modprobe lpfc lpfc_log_verbose=0x7f applies to all SCSI host (ports) managed by the lpfc driver.
```

Viewing Parameters with sysfs:

The `sysfs` file system is mounted and available as `/sys`. You must first identify the `scsi_host` that represents the adapter for which you want to modify the FC and FCoE driver parameters. All `scsi_hosts` bound to the FC and FCoE driver can be viewed with the following command:

```
ls -d /sys/bus/pci/drivers/lpfc/*/host*
```

Assuming you are interested in adapter `scsi_host 7`, you can list the FC and FCoE driver parameters for this particular adapter as:

```
ls -l /sys/class/scsi_host/host7/lpfc*
```

An example output follows:

```
-r--r--r-- 1 root root 4096 Feb 28 17:03 /sys/class/scsi_host/host7/lpfc_ack0
-r--r--r-- 1 root root 4096 Feb 28 17:03 /sys/class/scsi_host/host7/lpfc_fcp_class
-rw-r--r-- 1 root root 4096 Feb 28 17:03 /sys/class/scsi_host/host7/
lpfc_fdmi_on
-r--r--r-- 1 root root 4096 Feb 28 17:03 /sys/class/scsi_host/host7/lpfc_link_speed
-rw-r--r-- 1 root root 4096 Feb 28 15:34 /sys/class/scsi_host/host7/lpfc_log_verbose
-r--r--r-- 1 root root 4096 Feb 28 17:03 /sys/class/scsi_host/host7/lpfc_lun_queue_depth
-rw-r--r-- 1 root root 4096 Feb 28 17:03 /sys/class/scsi_host/host7/lpfc_max_luns
-rw-r--r-- 1 root root 4096 Feb 28 17:03 /sys/class/scsi_host/host7/lpfc_devloss_tmo
-rw-r--r-- 1 root root 4096 Feb 28 17:03 /sys/class/scsi_host/host7/lpfc_scan_down
-r--r--r-- 1 root root 4096 Feb 28 17:03 /sys/class/scsi_host/host7/lpfc_topology
-rw-r--r-- 1 root root 4096 Feb 28 17:03 /sys/class/scsi_host/host7/lpfc_use_adisc
```

Temporary Configuration Parameters with sysfs:

In the previous example, notice that the FC and FCoE driver parameters are available as files. Reading a file displays the current value of a driver parameter. If the permissions allow it, you can write a value to the file and it will take effect immediately.

Reading the `lpfc_log_verbose` file might show that its value is 0:

```
cat /sys/class/scsi_host/host7/lpfc_log_verbose
0
```

To modify the `lpfc_log_verbose` value to `0xffffffff`:

```
echo 0xffffffff > /sys/class/scsi_host/host7/lpfc_log_verbose
```

Reading the `lpfc_log_verbose` file now shows a value of `0xffffffff`:

```
cat /sys/class/scsi_host/host7/lpfc_log_verbose
0xffffffff
```

NOTE: Setting the `lpfc_log_verbose` value to `0xffffffff` might degrade system performance.

2.1.3 Creating a New Ramdisk Image

The `lpfc-install` script creates a ramdisk image containing the FC and FCoE driver for the currently running kernel.

NOTE: You must create a new ramdisk image whenever the `lpfc` options in `/etc/modprobe.conf` are changed and you want the change to take effect on the next reboot.

- To create a new initial ramdisk image for inbox FC and FCoE drivers and installed binary RPM FC driver kits, type:


```
dracut -f /boot/initramfs-<kernel-version>.img <kernel-version>
```

2.1.4 Dynamically Recognizing LUNs and Targets (Using scan)

The FC and FCoE driver enables you to dynamically recognize LUNs and targets without unloading or reloading the FC/FCoE module and without resetting the adapter.

To rescan an adapter's targets with `sysfs`, given the adapter's host number (in this example, 3), type:

```
echo "--" > /sys/class/scsi_host/host3/scan
```

To limit the rescan to a particular target, given the adapter's host number (in this example, 3) and the target number (in this example, 2), type:

```
echo "- 2 -" > /sys/class/scsi_host/host3/scan
```

You can also use the `lun_scan` script in the `/usr/sbin/lpfc` directory.

2.1.5 Persistent Naming

The generic device manager for the Linux kernel is `udev`, which primarily manages device nodes in the `/dev` directory.

2.1.5.1 Using udev to Discover Logical to Physical Mappings for sd Devices

In Linux, the driver for SCSI disk drives is `sd`. A disk device name has an `sd` prefix. Persistent names for `sd` devices are provided in the `/dev/disk/by-id` directory. To find the persistent `udev` name for the disk, which is currently `sdc`, type:

```
cd /dev/disk/by-id
ls -l | grep sdc
```

The sample output is:

```
lrwxrwxrwx 1 root root 9 2006-08-01 19:08 scsi-32000000c5005d6e6 -> ../../sdc
```

In the previous example, the disk has no partitions. If the disk had two partitions, the output would look like the following:

```
lrwxrwxrwx 1 root root 9 2006-08-01 19:08 scsi-32000000c5005d6e6 -> ../../sdc
lrwxrwxrwx 1 root root 10 2006-08-01 19:08 scsi-32000000c5005d6e6-part1 -> ../../sdc1
lrwxrwxrwx 1 root root 10 2006-08-01 19:08 scsi-32000000c5005d6e6-part2 -> ../../sdc2
```

2.1.5.2 Configuring the System to Boot Using Persistent Names

For SLES 11 and SLES 12:

NOTE: SLES 11 and SLES 12 are configured by default with `udev` to provide persistent names for hard disks, including FC-attached disks.

To use a persistent name for a boot device:

1. In `/boot/grub/menu.lst`, find the kernel line for the default boot. For example:

```
kernel /boot/vmlinuz root=/dev/sda2 vga=0x314
```
2. Find the persistent name for the root partition (following `root=` on the kernel line) by using the instructions in [Section 2.1.5.1, Using udev to Discover Logical to Physical Mappings for sd Devices](#).
3. In the same file, `/boot/grub/menu.lst`, replace the text after `root=` with the partition's persistent name. For example:

```
kernel /boot/vmlinuz root=/dev/disk/by-id/scsi-32000000c5005d6e6-part2 vga=0x314
```
4. Change any mounts listed in `/etc/fstab` that refer to this root partition by either its `/dev/sd` name or a file system label to use the persistent name as well.

For RHEL 6.x and RHEL 7.x:

To use a persistent name for a boot device:

1. In `/boot/grub/grub.conf`, find the kernel line for the default boot. For example:
`kernel /boot/vmlinuz -<kernel version> ro root=/dev/sda2`
2. Find the persistent name for the root partition (following `root=` on the kernel line) by using the instructions in [Section 2.1.5.1, Using udev to Discover Logical to Physical Mappings for sd Devices](#).
3. In the same file, `/boot/grub/menu.lst`, replace the text after `root=` with the partition's persistent name. For example:
`kernel /boot/vmlinuz -<kernel version> ro root=/dev/disk/by-id/scsi-32000000c5005d6e6-part2`
4. Change any mounts listed in `/etc/fstab` that refer to this root partition by either its `/dev/sd` name or a file system label to use the persistent name as well.

2.1.5.3 Using udev with st Devices

In Linux, the driver for SCSI tape drives is `st`. A tape device name has an `st` prefix. The `udev` rules for tape devices are the same as for disk devices. A unique ID must exist that persists across initiator reboots and persists regardless of discovery order.

You must consider whether the tape device is an FC tape device or an FC-SCSI tape device (in which multiple SCSI tape devices reside behind an FC controller). If it is an FC tape device, then the WWPN is unique and can be used to create the persistent name. In this case, the `scsi_id` command should return this as the unique identifier with a single-digit prefix. If the FC controller has multiple SCSI tape devices behind it, the WWPN is not unique, and the persistent name must use multiple information elements to build the unique ID. [FC Tape Device Examples](#) and [FC-SCSI Tape Device Example](#) are examples of each scenario.

FC Tape Device Examples:

The following examples use the `scsi_id` command to retrieve and generate a unique SCSI identifier:

```
scsi_id [options]
```

For these examples, the following [options] are used:

- `-g` Treats the device as white listed. It is needed on the command line or in the `scsi_id.config` file for the `scsi_id` command to generate any output. In the examples, the `-g` option is needed on the command line because the vendor and model for this tape device were not in the `/etc/scsi_id.config` file.
- `-s` Generates an id for the `sysfs`-device. Note that `-s` is an invalid option for `scsi_id` version 147.

NOTE: Since the [options] can vary depending on the version of the `scsi_id` command, refer to the `scsi_id` man page on your system for the correct and complete list of the [options].

The following example is an FC tape device using the SCSI generic driver (`sg`) rather than the SCSI tape driver. The value returned has a leading prefix of 3, which is the NAA type. The remaining digits represent the FC controller's WWPN.

```
scsi_id -g -s /sys/class/scsi_generic/sg0
350060b000029b592
```

The following example is an FC tape device using the SCSI tape driver. The value returned is the same as the previous example.

```
scsi_id -g -s /sys/class/scsi_tape/nst0
350060b000029b592
```

The following example uses a different FC tape vendor. Notice that the value returned is similar to the previous examples, with respect to the leading digit and the WWPN.

```
/sbin/scsi_id -g -s sys/class/scsi_tape/nst0
35005076300015101
```

FC-SCSI Tape Device Example:

The following is an example of a FC controller with multiple SCSI tape devices behind it (FC-SCSI tape device). When the Emulex driver is loaded, the SCSI mid-level discovers the SCSI tape devices as follows:

```
scsi scan: INQUIRY to host 14 channel 0 id 0 lun 0
scsi: unknown device type 12
Vendor: ADIC      Model: SNC 4000      Rev: 42d4
Type: RAID      ANSI SCSI revision: 03
Attached scsi generic sg5 at scsi14, channel 0, id 0, lun 0, type 12
scsi scan: INQUIRY to host 14 channel 0 id 0 lun 1
Vendor: ADIC      Model: Scalar 24      Rev: 227A
Type: Medium Changer      ANSI SCSI revision: 02
Attached scsi generic sg6 at scsi14, channel 0, id 0, lun 1, type 8
scsi scan: INQUIRY to host 14 channel 0 id 0 lun 2
Vendor: IBM      Model: ULTRIUM-TD2      Rev: 38D0
Type: Sequential-Access      ANSI SCSI revision: 03
Attached scsi tape st0 at scsi14, channel 0, id 0, lun 2
st0: try direct i/o: yes (alignment 512 B), max page reachable by HBA 4503599627370495
Attached scsi generic sg7 at scsi14, channel 0, id 0, lun 2, type 1
scsi scan: INQUIRY to host 14 channel 0 id 0 lun 3
Vendor: IBM      Model: ULTRIUM-TD2      Rev: 38D0
Type: Sequential-Access      ANSI SCSI revision: 03
Attached scsi tape st1 at scsi14, channel 0, id 0, lun 3
st1: try direct i/o: yes (alignment 512 B), max page reachable by HBA 4503599627370495
Attached scsi generic sg8 at scsi14, channel 0, id 0, lun 3, type 1
```

This log output shows a controller at LUN 0, the medium changer at LUN 1, and two SCSI tape devices at LUNs 2 and 3.

The following example is the result of a `scsi_id` call:

```
scsi_id -g -s /sys/class/scsi_tape/nst0
1IBM      ULTRIUM-TD2      1110133831
scsi_id -g -s /sys/class/scsi_tape/nst1
1IBM      ULTRIUM-TD2      1110133994
```

Notice that the unique ID contains three values with space delimiters. A `udev` rule must have a unique ID for the device, so that all three parts of this returned string are required. To do this, use the following command:

```
scsi_id -u -g -s /sys/class/scsi_tape/nst0
1IBM_____ULTRIUM-TD2_____1110133831
scsi_id -u -g -s /sys/class/scsi_tape/nst1
1IBM_____ULTRIUM-TD2_____1110133994
```

Creating the udev Persistent Name for SCSI Tape Device:

After you know the SCSI ID call needed to extract a unique ID, use the same process to create a `udev` persistent name for a SCSI tape device as on a SCSI disk device.

The rule for the FC tape device is

```
BUS="scsi", SYSFS{vendor}="HP", SYSFS{model}="ULTRIUM 3-SCSI", PROGRAM="/sbin/scsi_id -p 0x83 -u -g -s /sys/class/scsi_tape/nst%n",RESULT="350060b000029b592", SYMLINK="fc_lun_st%n"
```

The rule for the FC-SCSI tape device is:

```
BUS="scsi", SYSFS{vendor}="IBM", SYSFS{model}="ULTRIUM-TD2", PROGRAM="/sbin/scsi_id -p 0x83 -u -g  
-s /sys/class/scsi_tape/nst%n", RESULT="1IBM_____ULTRIUM-TD2_____1110133831",  
SYMLINK="fc_lun_st%n"  
BUS="scsi", RESULT="1IBM_____ULTRIUM-TD2_____1110133994", SYMLINK="fc_lun_st%n"
```

Create a new file named `/etc/udev/rules.d/45-local.rules`, and put the appropriate rule in it. Then run `udevtrigger` to reload the `udev` rules, and the output of the rules will follow:

```
udevtrigger  
ls -al /dev/fc*  
lrwxrwxrwx 1 root root 3 Apr  7 15:03 fc_lun_st0 -> st0  
lrwxrwxrwx 1 root root 3 Apr  7 15:03 fc_lun_st1 -> st1
```

2.1.5.4 Persistent Naming References

See the following references for more information on persistent naming:

- www.reactivated.net/udevrules.php by Daniel Drake (dsd)
- kernel.org/pub/linux/utils/kernel/hotplug/udev_vs_devfs by Greg Kroah-Hartman

2.1.6 ExpressLane Support

The host servers do not distinguish between lower and higher priority workloads being sent to LUNs. For flash storage deployed in mixed storage environments, the combination of data from rotating media and flash devices can cause congestion on the network and reduced performance for the flash storage devices.

The ExpressLane feature enables special priority queuing for selected LUNs (ExpressLane LUNs). The ExpressLane LUN performance is superior to that of regular LUNs. Mixed storage environments can use the ExpressLane feature to alleviate congestion, reduce latency, and improve throughput, ensuring that key applications receive highest priority. Switches can provide up to three priority levels: Low, Medium, and High,

The following requirements must be met to use the ExpressLane feature:

- The ExpressLane feature is supported only on Emulex LPe16000-series, LPe31000-series, and LPe32000-series adapters.

An ExpressLane LUN must be enabled in the driver before it can be used by the Emulex OneCommand Manager application. Additionally, the priority value to mark each of the ExpressLane frames must be specified to the FC/FCoE driver through the appropriate driver parameters. For additional information, refer to the *OneCommand Manager for LightPulse Adapters User Guide*.

NOTE: Enabling the ExpressLane feature on a LUN, when maximum LUNs are already enabled for ExpressLane, it might result in an error. Use the `/usr/sbin/lpfc/lpfc_clean_xlane_conf.sh` script to clear any unwanted entries and retry enabling ExpressLane.

2.2 vPort Configuration

This section describes how to create, delete, and display vPorts.

2.2.1 vPort Configuration Prerequisites

Before configuring vPorts, consider the following points:

- Ensure that you are using the latest recommended firmware for vPort functionality. For the latest firmware, check the Documents and Downloads area of www.broadcom.com.

NOTE: Before performing a firmware update, driver installation is required. For more information on installing the driver, see [Chapter 1, Installing and Uninstalling](#).

- Loop devices and NPIV are not supported on the same port simultaneously. If you are running a loop topology and you create a vPort, the vPort link state is offline.
- vPorts do not persist across system reboots.

2.2.2 Creating, Deleting, and Displaying vPorts

vPorts are created through `sysfs` entries that are presented in the physical port's `sysfs` directory. There are three scripts to create, delete, and display vPorts. The scripts reside in the `/usr/sbin/lpfc` directory and are part of the OneCommand Manager application kit.

When NPIV is enabled and vPorts are configured, it can take longer for the adapter to finish discovery in some cases because each vPort must perform discovery independently. As more vPorts are configured, the amount of time that the driver and adapter take to finish discovery of remote ports on the SAN increases. To compensate for this extended amount of time taken in discovery, set the `lpfc_devloss_tmo` parameter to 60 seconds when NPIV is enabled.

2.2.2.1 Creating vPorts Using the `mkvport.sh` Script

You can use the `mkvport` script to create vPorts. To see the usage information, run the script with no parameters specified. The `mkvport.sh` script uses the following syntax:

```
./mkvport.sh <Physical Port's Host number> <Port Name> <Node Name>
```

You must supply the physical port's host number, WWPN, and WWNN when using the `mkvport.sh` script. For example, to create a vPort with port name of `1000000c94ac63a` and a node name of `20010000c94ac63a` on the physical port with `scsi_host` name `host7`, type:

```
./mkvport.sh host7 1000000c94ac63a 20010000c94ac63a
```

This script fails if the vPort is not created.

NOTE: It is possible for a vPort to be created successfully but to be in a failed state. For example, loop devices and NPIV are not supported on the same port simultaneously. If you are running a loop topology and you create a vPort, the vPort's link state will be offline.

2.2.2.2 Deleting vPorts Using the `rmvport.sh` Script

NOTE: You must unmap, unmount, and flush I/Os to vPort-connected devices before deleting the vPort.

You can use the `rmvport` script to delete vPorts. To see the usage information, run the script with no parameters specified. The `rmvport.sh` script uses the following syntax:

```
./rmvport.sh <VPort's Host number>
```

Or

```
./rmvport.sh <Port Name> <Node Name>
```

To delete the vPort with a port name of 10000000c94ac63a and a node name of 20010000c94ac63a, type:

```
./rmvport.sh 10000000c94ac63a 20010000c94ac63a
```

This script can take up to 30 seconds to finish. The script fails if the vPort is not deleted.

2.2.2.3 Displaying vPorts Using the lsvport.sh Script

You can use the `lsvport` script to display the vPorts and physical ports that are present on the system. Run the script with no parameters to display port information.

For example:

```
./lsvport.sh
lpfc0: host6 10000000c93a5b5e:20000000c93a5b5e LP10000 NPIV Not Supported
lpfc1: host7 10000000c93a5b5d:20000000c93a5b5d LP10000 NPIV Not Supported
lpfc2: host8 10000000c93cc8dd:20000000c93cc8dd LPe12000 NPIV Physical
lpfc3: host9 10000000c93cc8dc:20000000c93cc8dc LPe12000 NPIV Physical
lpfc4: host10 10000000c94ac63a:20010000c94ac63a NPIV Virtual (VPI 1)
```

In reference to the previous example:

- For LPFC0 and LPFC1, *NPIV Not Supported* indicates that this adapter/firmware combination does not support the creation of vPorts.
- For LPFC2 and LPFC3, *NPIV Physical* refers to a physical port of this adapter.
- For LPFC4, *NPIV Virtual* refers to a vPort of this adapter.

2.2.3 vPort sysfs

This section describes the vPort `sysfs` structure and vPort `sysfs` entries.

2.2.3.1 vPort sysfs Tree

When a vPort is created, three new directories are created in the class tree:

```
/sys/class/scsi_host/hostY/
/sys/class/fc_host/hostY/
/sys/class/fc_vports/vport-X:0-Z/-
```

Creating a new vPort also creates a new `sysfs` directory in the bus and devices tree:

```
/sys/bus/pci/drivers/lpfc/0000:A:B:C/hostX/vport-X:0-Z/hostY
/sys/devices/pci0000:A/0000:A:B:C/hostX/vport-X:0-Z/hostY
```

In both directories, there is a `hostY` directory that contains the remote ports that this new host can access:

```
/sys/bus/pci/drivers/lpfc/0000:A:B:C/hostX/vport-X:0-Z/hostY
/sys/bus/pci/drivers/lpfc/0000:A:B:C/hostX/vport-X:0-Z/hostY/rport-Y:0-0
/sys/bus/pci/drivers/lpfc/0000:A:B:C/hostX/vport-X:0-Z/hostY/rport-Y:0-1
/sys/bus/pci/drivers/lpfc/0000:A:B:C/hostX/vport-X:0-Z/hostY/rport-Y:0-2
```

In this example:

- X indicates the host value for the parent `fc_host` from which this vPort was created.
- Y indicates the new host value for the vPort that was created.
- Z indicates the instance of vPort created from the parent `fc_host`. A, B, and C indicate the PCI hierarchy for each physical FC or FCoE port.

`hostY` is the new host created for the new vPort. `vport-X:0-Z` uniquely identifies the vPort and indicates the parent host structure (X) that created this vPort.

For example, if a vPort is created from host5, a new `scsi_host`, a new `fc_host`, a new `fc_vport`, and a new entry under the bus tree are created as well.

```
ls /sys/class/scsi_host/
host0 host1 host4 host5 host6
ls /sys/class/fc_host/
host4 host5 host6
ls /sys/class/fc_vports/
vport-5:0-0
```

2.2.3.2 vPort sysfs Entries

NOTE: vPort `sysfs` entries shown in the following table are present only if the driver was loaded with `lpfc_enable_npiv` enabled.

Table 3: vPort sysfs Entries

vPort <code>sysfs</code> Entries	Type	Range/Input	Location and Description
<code>lpfc_peer_port_login</code>	Read/Write	0 = Off (default) 1 = On	<p><code>/sys/class/scsi_host/hostX/lpfc_peer_port_login</code></p> <p>This entry sets the port behavior when discovering targets in the SAN. The default behavior (value = 0) will log in only to <code>N_Ports</code> that are physically located on a different port. The port will still attempt to log in to targets on all other ports (including the other port in a dual-port adapter).</p> <p>If this parameter is turned on (value = 1), the port attempts to log in to all <code>N_Ports</code>, even if they are physically located on the same port.</p> <p>NOTE: This parameter was created to reduce the amount of hardware resources (for example, RPIs) that the driver requires. In a configuration where there are many vPorts on one physical port, this capability greatly reduces the number of RPIs that the driver uses.</p>
<code>lpfc_restrict_login</code>	Read/Write	0 = Off 1 = On (default)	<p><code>/sys/class/scsi_host/hostX/lpfc_restrict_login</code> (vPorts only)</p> <p>This entry sets the vPort behavior when discovering targets in the SAN. The default behavior (value = 1) prevents the vPort from logging into other initiator ports in the SAN. It also rejects logins from other ports in the SAN because it assumes that all ports that send a PLOGI are initiators.</p> <p>If this <code>sysfs</code> entry is turned off, the driver attempts to log in to every port that it can access in the SAN, and accepts logins from all ports.</p> <p>NOTE: This parameter was created to reduce the amount of hardware resources (for example, RPIs) that the driver requires. In a SAN where other initiators exist, this capability greatly reduces the number of RPIs that the driver uses.</p>

Table 3: vPort sysfs Entries (Continued)

vPort <i>sysfs</i> Entries	Type	Range/Input	Location and Description
max_npiv_vports	Read-only	integers	<p><code>/sys/class/fc_host/hostX/max_npiv_vports</code></p> <p>This entry displays the maximum number of vPorts that are supported by the underlying <code>fc_host</code> hardware.</p> <p>This <i>sysfs</i> entry exists only if the <code>vport_create</code> and <code>vport_delete</code> <i>sysfs</i> entries exist. If an <code>fc_host</code> does not support NPIV, this <i>sysfs</i> entry might not exist.</p> <p>Use this <i>sysfs</i> entry with the <code>npiv_vports_inuse</code> entry to determine whether the maximum number of vPorts have been created on this <code>fc_host</code>.</p>
node_name	Read-only	16-byte hexadecimal value	<p><code>/sys/class/fc_host/hostX/node_name</code> <code>/sys/class/fc_vports/vport-X:0-Z/node_name</code></p> <p>These entries display the physical or vPort node name.</p> <p>You assign this value when the vPort is created, and it is transmitted to the fabric upon fabric login.</p>
npiv_vports_inuse	Read-only	integers	<p><code>/sys/class/fc_host/hostX/npiv_vports_inuse</code></p> <p>This entry displays the number of vPorts that were created on this <code>fc_host</code>.</p> <p>This <i>sysfs</i> entry exists only if the <code>vport_create</code> and <code>vport_delete</code> <i>sysfs</i> entries exist. If an <code>fc_host</code> does not support NPIV, this <i>sysfs</i> entry might not exist.</p> <p>Use this <i>sysfs</i> entry with <code>max_npiv_vports</code> to determine whether the maximum number of vPorts have been created on this <code>fc_host</code>.</p>
port_name	Read-only	16-byte hexadecimal value	<p><code>/sys/class/fc_host/hostX/port_name</code> <code>/sys/class/fc_vports/vport-X:0-Z/port_name</code></p> <p>This entry displays the physical or vPort port name.</p> <p>You assign this value when the vPort is created, and it is transmitted to the fabric upon fabric login.</p>
vport_create	Write-only	WWPN; WWNN	<p><code>/sys/class/fc_host/hostX/vport_create</code></p> <p>This entry creates a vPort on the physical port on which <code>hostX</code> is located. The new vPort will have a WWPN and WWNN present on the fabric based on the WWPN and WWNN that are entered with this <i>sysfs</i> entry.</p> <p>This entry returns a 0 if the vPort creation was successful. A nonzero value indicates that the vPort was not created.</p> <p>If an <code>fc_host</code> does not support NPIV, then this <i>sysfs</i> entry might not exist.</p> <p>NOTE: It is possible for the vPort creation to succeed but for the vPort to be in a failed or inoperative state. Use the new <i>sysfs</i> tree created by the new vPort to check the state of the new vPort.</p>

Table 3: vPort sysfs Entries (Continued)

vPort <i>sysfs</i> Entries	Type	Range/Input	Location and Description
vport_delete	Write-only	WWPN; WWNN	<p><code>/sys/class/fc_host/hostX/vport_delete</code></p> <p>This entry deletes a vPort on the physical port on which <code>hostX</code> is located. The vPort matching the WWPN and WWNN is immediately deleted.</p> <p>This entry returns a 0 if the vPort deletion was successful. A nonzero value indicates that the vPort was not deleted.</p> <p>If an <code>fc_host</code> does not support NPIV, then this <i>sysfs</i> entry might not exist.</p> <p>NOTE: This entry deletes the vPort even if there are mounted file systems being accessed through this vPort, or if there are open files on it.</p>

2.2.4 Monitoring vPorts with `fc_vport`

The transport creates an `fc_vports` directory that you can use to monitor vPorts. This directory is populated entirely of vPorts and has links from each to the `fc_host` associated with that vPort.

```
ls /sys/class/fc_vports/
vport-5:0-0
ls -d /sys/bus/pci/drivers/lpfc/*/host*/*/host*
/sys/bus/pci/drivers/lpfc/0000:03:06.1/host5/vport-5:0-0/host6
ls /sys/devices/pci*/*/host5/vport-5*/host6
power rport-6:0-0 rport-6:0-1 rport-6:0-2 uevent
ls /sys/devices/pci*/*/host5/vport-5*/host6/rport-*
/sys/devices/pci00:03/00:03:06.1/host5/vport-5:0-0/host6/rport-6:0-0:
power uevent

/sys/devices/pci00:03/00:03:06.1/host5/vport-5:0-0/host6/rport-6:0-1:
power uevent

/sys/devices/pci00:03/00:03:06.1/host5/vport-5:0-0/host6/rport-6:0-2:
power target6:0:0 uevent
```

In this example:

- A new entry exists in the `fc_vports` directory for the vPort (`vport-5:0-0`). The `vport-5:0-0` entry indicates that the vPort was created from `host5`, and it is the first (0) vPort to be created on that `fc_host`.
- The new host for the vPort is `host6`, and it will appear in the usual directories.
- A new directory exists in the bus tree. This new directory indicates that `host6` was created under `vport-5:0-0` (which was created from `host5`).

2.2.5 vPort Configuration Limits

vPort configuration limits are designated as enforced or unenforced. Enforced limits are limits that the driver enforces and prevents you from exceeding. Unenforced limits are limits that the driver cannot enforce, but configurations that exceed them are unsupported.

The following vPort configuration limits have been tested with and are supported by the Emulex driver. Configurations that exceed one or more of these limits are unsupported.

- Before the vPort is deleted or the driver is unloaded, I/O devices accessed through a vPort must be stopped and file systems must be unmounted.

- The maximum number of vPorts configurable on a physical port is 255.
- The maximum number of LUNs supported on each driver port is 256.
- The maximum number of targets supported for each driver port is 255.
- The maximum number of driver ports in one zone is 64. This limit is based on the system's ability to recover from link events within the time constraints of the default timers.

The NPIV use-cases that involve virtual server environment include associating a vPort with a virtual machine, and placing the virtual machine in its own zone, which results in one vPort per zone. In the case of load-balanced environments, this can increase typically to two vPorts per virtual machine, to a practical limit of something far less than 50.

In the NPIV cases not related to virtual server environments, zoning is typically initiator-zoning, again resulting in one vPort, or a low number of vPorts in the case of load-balancing, within a given zone. If there are too many vPorts within a single zone, expected behavior includes devices being lost after link events.

- The minimum lifetime of a vPort is 60 seconds. An unenforced limit of 60 seconds exists between the creation of a vPort and the deletion of the same vPort. vPorts are designed to exist for a long time in the system, and the creation of vPorts is asynchronous, which means that a vPort might not be finished with FC or SCSI discovery when the command to create a vPort is finished.

2.3 Ethernet Driver Configuration for PPC CNAs

The following section describes how to configure parameters for the Ethernet driver.

2.3.1 Ethernet Driver Configuration Parameters

The Ethernet driver configurable parameters are described in the following table.

Table 4: Ethernet Driver Configuration Parameters

Parameter	Description
num_vfs	In systems supporting SR-IOV, when IOV is enabled, this parameter indicates the number of VFs to be enabled per PF. For configuring SR-IOV, see Section 2.3.7, SR-IOV Configuration . The default value is 0 (SR-IOV is not enabled). The possible values are 0 to 63. NOTE: This parameter is obsolete. Use the <code>sysfs</code> method to enable or disable VFs.
rss_on_mc	Enables receive-side scaling (RSS) on multichannel functions that have the capability. The default value is 0 (disabled). The enabled value is 1. NOTE: Currently, this parameter is ignored by the driver. The driver enables RSS on multichannel functions that have the capability by default.
rx_frag_size	The size of fragments used to DMA received data. The possible values are 2048 (default), 4096, and 8192.

You can configure Linux to automatically load the driver with any of these options after each reboot. To do so, add a line to `/etc/modprobe.conf` with the required options. For example, to load the driver with the fragment size of 4096, add the following line:

```
options lpnic rx_frag_size=4096
```

2.3.2 Support for Ethtool set-channels/get-channels Commands

The Ethtool `get-channels` command – `ethtool -l <ethx>` displays the following information:

- The current number of Rx/Tx/interrupt queue pairs (a *combined channel*) created by the NIC function.
- The maximum number of channels supported by the NIC function.

The Ethtool `set-channels` command `ethtool -L <ethx> combined N` configures the number of channels requested for the NIC function.

NOTE:

- You can increase (up to the maximum supported limit) or decrease the number of channels used by the NIC function.
- The NIC driver supports only *combined* channels.
- The maximum number of channels supported by an interface is also limited to the number of cpu-cores in the server.

NOTE: In older Linux distributions, use the `sysfs` interface as follows:

- The maximum number of channels supported by NIC function:
`cat /sys/class/net/<ethx>/max_qs`
- The current number of channels configured:
`cat /sys/class/net/<ethx>/num_qs`
- To configure the *N* number of channels requested for the NIC function:
`echo N > /sys/class/net/<ethx>/num_qs`

2.3.3 Support for Ethtool set-dump Command

The Ethtool `set-dump` command `ethtool -W <ethx> N` can be used as follows:

- To trigger a FW dump:
`ethtool -W <ethx> 1`
- To delete an existing dump:
`ethtool -W <ethx> 2`

NOTE:

- Delete an existing dump first before triggering a new dump.

NOTE: In older Linux distributions, use the `sysfs` interface as follows:

- To trigger a FW dump:
`# echo 1 > /sys/class/net/<ethx>/trigger_dump`
- To delete an existing dump:
`# echo 1 > /sys/class/net/<ethx>/delete_dump`

2.3.4 Transmit/Receive Queue Counts

Use the command `ethtool -l <ethx>` to view the Transmit and Receive queue counts.

For example:

```
ethtool -l <ethx>
```

Output

```
Channel parameters for enp9s0f0:
Pre-set maximums:
RX:                7
TX:                7
Other:             0
Combined:          8
Current hardware settings:
RX:                0
```

```
TX:                0
Other:             0
Combined:         8
```

2.3.5 Support for Ethtool set-rxfh/get-rxfh Commands

The Ethtool `get-rxfh` and `set-rxfh` commands are available in kernel versions of RHEL 7.1, 3.16.0 and higher.

The Ethool `get-rxfh` command – `ethtool -x <ethx>` shows the following information:

- The current indirection table entries
- The programmed RSS hash key

The Ethtool `set-rxfh` command – `ethtool -X <ethx> [hkey aa:bb:cc...] [equal N | weight W0 W1...]` sets the following parameters:

- Indirection table entries
- RSS hash key

NOTE: In older Linux distributions, use the `sysfs` interface as follows:

- To show the currently configured RSS hash key:

```
cat /sys/class/net/<ethx>/rss_hashkey
```
- To configure the new RSS hash key:

```
echo -e "aa:bb:cc...." > /sys/class/net/<ethx>/rss_hashkey
```

2.3.6 Support for Showing Onboard Die Temperature

Onboard die temperature is exported through the `sysfs` interface. On kernels that support the hardware monitoring framework, the temperature is exported using the `hwmonsysfs` interface. Temperatures exported using the `hwmon` interface can also be read using the `lm-sensors` utility with the `sensors` command.

To show the current board temperature using the `sysfs` node:

```
cat /sysclass/hwmon/<hwmonx>/device/templ_input
```

Otherwise, the current board temperature is exported with the regular `sysfs` interface:

```
cat /sys/class/net/<ethx>/device/templ_input
```

2.3.7 SR-IOV Configuration

2.3.7.1 Introduction

This section contains requirements and instructions to use SR-IOV with the following host operating systems:

- RHEL 6.8 – 64-bit KVM
- RHEL 6.9 – 64-bit KVM
- RHEL 6.10 – 64-bit KVM
- RHEL 7.2 – 64-bit KVM
- RHEL 7.3 – 64-bit KVM
- RHEL 7.4 – 64-bit KVM
- RHEL 7.5 – 64-bit KVM
- SLES 11 SP3 – 64-bit Xen and KVM
- SLES 11 SP4 – 64-bit Xen and KVM
- SLES 12 SP2 – 64-bit Xen and KVM

- SLES 12 SP3 – 64-bit Xen and KVM
- SLES 15 – 64-bit Xen and KVM
- SLES and RHEL supported with Xen

These environments support the enabling of multiple PCIe virtual functions (VFs) for a PCIe physical function (PF). Each of these VFs can be assigned to virtual machines (VMs). A VF enables the guest operating system to access to the Emulex OneConnect[®] adapter directly, such that guest performance is not limited by the overheads of the hypervisor.

With SR-IOV, VMs directly drive I/Os on the NIC. Therefore, SR-IOV has the following advantages over traditional virtualized I/O:

- Improved device performance for virtual guests
- Increased scalability
- Reduced CPU utilization
- Reduced latency

Known issues include:

- The `Kdump` (kernel dump) feature is not supported when SR-IOV is enabled.

2.3.7.2 Setting Up SR-IOV

Prerequisites:

To set up SR-IOV on your system, you need the following:

- A server or blade with an SR-IOV-capable motherboard BIOS.

NOTE: Configuration mechanisms for parameters, such as MAC address, VLAN, and QoS for VF are supported in RHEL 6 and later, and SLES 11 and later distributions.

- A PPC CNA that supports SR-IOV.
- If a KVM hypervisor is installed, it must contain the `qemu-kvm` packages.

Depending on your system, perform one or more of the following tasks to set up your BIOS. For more information, refer to the manufacturer's instructions for your system.

- Enable SR-IOV in the system BIOS, which supports SR-IOV functionality.
- Enable Intel Virtualization Technology support for Direct I/O VT-d.

You can use the Emulex OneCommand Manager CLI application to set up SR-IOV. Refer to the *Emulex OneCommand Manager Command Line Interface for LightPulse Adapters User Guide* for more information.

1. Install the required Linux operating system that serves as a hypervisor.
2. For legacy boot, perform the following step:
Update the `/boot/grub/menu.lst` file to include the following command line load parameter for the Linux kernel:
`intel_iommu=on`
For UEFI boot, perform the following step:
Update the `/etc/default/grub` file to include the following command line load parameter for the Linux kernel:
`intel_iommu=on`
3. Install the appropriate version of the Emulex Ethernet driver that supports SR-IOV for the operating system version that you are using.
4. Reboot the server for the new changes to become operational.

5. Use the `lspci -vvv` output command to check if SR-IOV is properly enabled.

The `lspci -vvv` command returns an SR-IOV capability section for each Ethernet PF. The Initial VFs and Total VFs should be nonzero. Make a note of the `lspci` command output in the hypervisor. This output is needed in step 10. Specifically, make note of the `pci-id` of the VFs that have been created.

The following is an example of the SR-IOV capabilities section output of the PF with SR-IOV enabled:

```
Capabilities: [180] Single Root I/O Virtualization (SR-IOV)
IOVCap: Migration-, Interrupt Message Number: 000
IOVctl: Enable- Migration- Interrupt- MSE- ARIHierarchy-
IOVSta: Migration-
Initial VFs: 32, Total VFs: 32, Number of VFs: 0, Function Dependency
Link: 00
VF offset: 0, stride: 1, Device ID: 0710
Supported Page Size: 00000557, System Page Size: 00000001
Region 0: Memory at 0000000000000000 (64-bit, non-prefetchable)
VF Migration: offset: 00000000, BIR: 0
```

The following is an example of the SR-IOV capabilities section output of the PF with SR-IOV disabled:

```
Capabilities: [180] Single Root I/O Virtualization (SR-IOV)
IOVCap: Migration-, Interrupt Message Number: 000
IOVctl: Enable- Migration- Interrupt- MSE- ARIHierarchy-
IOVSta: Migration-
Initial VFs: 0, Total VFs: 0, Number of VFs: 0, Function Dependency
Link: 00
VF offset: 0, stride: 1, Device ID: 0710
Supported Page Size: 00000557, System Page Size: 00000001
Region 0: Memory at 0000000000000000 (64-bit, non-prefetchable)
VF Migration: offset: 00000000, BIR: 0
```

6. Create a VM using the Virtual Machine Manager utility, and install any supported operating system on the VM.
7. Shut down the VM.
8. Echo the number of VFs/PF to the `sriov_numvfs` file, where *X* is the number of VFs per PF:

```
echo X > /sys/bus/pci/devices/.../sriov_numvfs
```

The possible values are 0 to 64 per physical port. The default value is 0 (SR-IOV is not enabled).

With the latest kernels, the `sysfs` path to enable VFs is

```
/sys/bus/pci/devices/<device-id>/sriov_numvfs
```

and for earlier kernels, the `sysfs` path to enable VFs is

```
/sys/bus/pci/devices/<device-id>/net/<ifaceName>/sriov_numvfs
```

The total number of VFs can be distributed among available ports as required, but each port has a maximum of 64 VFs.

[Table 5: Default Link Status of the PFs and VFs](#) on page 37 lists the total number of VF counts that are supported for various adapter configurations.

The maximum number of VFs supported per PF can be read from:

```
cat /sys/bus/pci/devices/.../sriov_totalvfs
```

NOTE: VFs are supported only for network functions; they are not supported for storage functions.

9. The total VF count that is supported for a two-port FCoE+NIC PPC CNAs (two NIC PFs and two FCoE PFs) is 128. The VF count is symmetrically distributed across the ports, which means that the VF count per NIC port is 64. Detach VFs from the host before adding them to the guest.

Example 1 (using only one VF per physical port):

RHEL 6.x KVM: View Emulex PCI devices with the `lspci` command:

```
lspci | grep Emulex 16:00.0 Ethernet controller: Emulex Corporation OneConnect NIC (Lancer) (rev 30)<-- LOM Port 0 (Function 0) 16:00.1 Ethernet controller: Emulex Corporation OneConnect NIC (Lancer) (rev 30)<-- LOM Port 1 (Function 1) 16:04.0 Ethernet controller: Emulex Corporation OneConnect NIC (Lancer) (rev 30)<-- VF LOM Port 0 (Function 0) 16:08.0 Ethernet controller: Emulex Corporation OneConnect NIC (Lancer) (rev 30)<-- VF LOM Port 1 (Function 1) # lspci -nn | grep Emulex 16:00.0 Ethernet controller: [0200]: Emulex Corporation OneConnect NIC (Lancer)[10df:0e220] (rev 03) <-- LOM Port 0 (Function 0) 16:00.1 Ethernet controller: [0200]: Emulex Corporation OneConnect NIC (Lancer)[10df:0e220] (rev 03) <-- LOM Port 1 (Function 1) 16:04.0 Ethernet controller: [0200]: Emulex Corporation OneConnect NIC (Lancer)[10df:0e220] (rev 03) <-- VF LOM Port 0 (Function 0) 16:08.0 Ethernet controller: [0200]: Emulex Corporation OneConnect NIC (Lancer)[10df:0e220] (rev 03) <-- VF LOM Port 1 (Function 1)
```

Detach desired VFs:

```
# virsh nodedev-dettach pci_0000_16_04_0 Device pci_0000_16_04_0 detached <-- VF LOM Port 0 (Function 0)
# virsh nodedev-dettach pci_0000_16_08_0 Device pci_0000_16_08_0 detached <-- VF LOM Port 1 (Function 1)
```

Example 2 (using only one VF per physical port):

SLES 11.x Xen: View Emulex PCI devices with the `lspci` command:

```
lspci | grep Emulex 16:00.0 Ethernet controller: Emulex Corporation OneConnect NIC (Lancer) (rev 30)<-- LOM Port 0 (Function 0) 16:00.1 Ethernet controller: Emulex Corporation OneConnect NIC (Lancer) (rev 30)<-- LOM Port 1 (Function 1) 16:04.0 Ethernet controller: Emulex Corporation OneConnect NIC (Lancer) (rev 30)<-- VF LOM Port 0 (Function 0) 16:08.0 Ethernet controller: Emulex Corporation OneConnect NIC (Lancer) (rev 30)<-- VF LOM Port 1 (Function 1) # lspci -nn | grep Emulex 16:00.0 Ethernet controller: [0200]: Emulex Corporation OneConnect NIC (Lancer)[10df:0e220] (rev 03) <-- LOM Port 0 (Function 0) 16:00.1 Ethernet controller: [0200]: Emulex Corporation OneConnect NIC (Lancer)[10df:0e220] (rev 03) <-- LOM Port 1 (Function 1) 16:04.0 Ethernet controller: [0200]: Emulex Corporation OneConnect NIC (Lancer)[10df:0e220] (rev 03) <-- VF LOM Port 0 (Function 0) 16:08.0 Ethernet controller: [0200]: Emulex Corporation OneConnect NIC (Lancer)[10df:0e220] (rev 03) <-- VF LOM Port 1 (Function 1)
```

List Emulex PCI device by device code:

```
# virsh nodedev-list | grep 10df pci_10df_e220 <-- LOM Port 0 (Function 0) pci_10df_e220_0 <-- LOM Port 1 (Function 1) pci_10df_e220_1 <-- VF LOM Port 0 (Function 0) pci_10df_e220_2 <-- VF LOM Port 1 (Function 1) # virsh nodedev-dettach pci_10df_e220_1 Device pci_10df_e220_1 detached <-- VF LOM Port 0 (Function 0) # virsh nodedev-dettach pci_10df_e220_2 Device pci_10df_e220_2 detached <-- VF LOM Port 1 (Function 1)
```

10. Use the Virtual Machine Manager GUI to attach the VF ([step 5 on page 34](#)) to the guest operating system by using the add physical device option.

NOTE: To reconfigure a system that is already set up, remove the attached VF from the guest operating system by selecting the VF and using the remove option. Refer to the documentation for the host operating system for information on using the Virtual Machine Manager to attach and remove virtual interfaces.

11. Start the RHEL 6.x guest operating system. After the guest operating system is booted, use the `lspci` command to confirm the visibility of the NIC function. The output shows a NIC function, for example:

```
06:00.0 Ethernet controller: Emulex Corporation OneConnect NIC (Lancer)
```

12. The Ethernet driver automatically loads with the out-of-box driver, and creates the network interfaces. Use the `ifconfig` command to verify that the interface is created.

13. After configuring the network interfaces with proper IP addresses, you can send and receive network traffic from the VM. Refer to the documentation for the host and guest operating systems for information on network configuration.

2.3.7.3 Assigning VFs to a VM on the SLES Operating System

To assign VFs to the VM in the SLES Xen kernel, the VF must be unbound from the NIC module and then bound to the pciback module.

NOTE: In the following steps, `0000:07:0b.5` is used as an example. To match those instances to the port that you want to use, you need to select the entry that matches the PCI bus, device, or the function that corresponds to the port that you want to assign. Use the `ethtool` utility to determine this information, such as `ethtool -i eth0` (where `eth0` is the interface you want to assign).

1. Load the `pciback` driver `modprobe pciback`.

2. Navigate to the `/sys/bus/pci/drivers/pciback` directory, and ensure that the following is displayed:

```
Panama-Sles11sp2:/sys/bus/pci/drivers/pciback # ls -lrt
total 0
--w----- 1 root root 4096 Sep 5 15:29 unbind
--w----- 1 root root 4096 Sep 5 15:29 uevent
-r----- 1 root root 4096 Sep 5 15:29 slots
--w----- 1 root root 4096 Sep 5 15:29 remove_slot
--w----- 1 root root 4096 Sep 5 15:29 remove_id
-rw----- 1 root root 4096 Sep 5 15:29 quirks
-rw----- 1 root root 4096 Sep 5 15:29 permissive
--w----- 1 root root 4096 Sep 5 15:29 new_id
lrwxrwxrwx 1 root root 0 Sep 5 15:29 module ->
../../../../module/pciback
--w----- 1 root root 4096 Sep 5 15:34 new_slot
--w----- 1 root root 4096 Sep 5 15:34 bind
```

3. Navigate to the `/sys/bus/pci/drivers/lpnic` directory, and ensure that the following is displayed:

```
--w----- 1 root root 4096 Sep 5 15:32 uevent
--w----- 1 root root 4096 Sep 5 15:32 remove_id
--w----- 1 root root 4096 Sep 5 15:32 new_id
lrwxrwxrwx 1 root root 0 Sep 5 15:32 module ->
../../../../module/lpnic
--w----- 1 root root 4096 Sep 5 15:32 bind
--w----- 1 root root 4096 Sep 5 15:33 unbind
lrwxrwxrwx 1 root root 0 Sep 5 15:32 0000:07:0b.5 ->
../../../../devices/pci0000:00/0000:00:09.0/0000:07:0b.5
lrwxrwxrwx 1 root root 0 Sep 5 15:32 0000:07:0b.4 ->
../../../../devices/pci0000:00/0000:00:09.0/0000:07:0b.4
lrwxrwxrwx 1 root root 0 Sep 5 15:32 0000:07:0b.3 ->
../../../../devices/pci0000:00/0000:00:09.0/0000:07:0b.3
lrwxrwxrwx 1 root root 0 Sep 5 15:32 0000:07:0b.2 ->
../../../../devices/pci0000:00/0000:00:09.0/0000:07:0b.2
lrwxrwxrwx 1 root root 0 Sep 5 15:32 0000:07:0b.1 ->
../../../../devices/pci0000:00/0000:00:09.0/0000:07:0b.1
lrwxrwxrwx 1 root root 0 Sep 5 15:32 0000:07:0b.0 ->
../../../../devices/pci0000:00/0000:00:09.0/0000:07:0b.0
```

4. Unbind the VF from the `lpnic` driver.

```
echo -n "0000:07:0b.5" > /sys/bus/pci/drivers/lpnic/unbind
```

5. Bind the driver to the `pciback` module.

```
echo -n "0000:07:0b.5" > /sys/bus/pci/drivers/pciback/new_slot
echo -n "0000:07:0b.5" > /sys/bus/pci/drivers/pciback/bind
```

- Navigate to the `/sys/bus/pci/drivers/pciback` directory, and ensure that the device `0000:07:0b.5` is listed under it. In addition, check that `dmesg` logs report the same device.

You can now launch `qemu-kvm` and attach VF `0000:07:0b.5` to any desired VM.

2.3.7.4 Link State Reporting with SR-IOV

When VEB is used to switch traffic between the functions of an SR-IOV-enabled port, the link status reported to the operating system stack when the physical port is not connected varies with the adapter type. The VFs indicate the link as DOWN when the physical port is not linked up. This behavior enables two or more VFs to be configured in a bond.

The following table summarizes the default link status of the PFs and VFs when SR-IOV is enabled in BIOS and in the driver.

Table 5: Default Link Status of the PFs and VFs

Function	Physical Link Is Not Connected
Physical function (PF)	Down
Virtual function (VF)	Down

The default behavior of VFs can be changed using the `ip link` command in distributions of Linux that support `iproute` version 3.11.0 or later. For details on configuring the link state for VFs, see [Section 2.3.7.6, Link State Configuration](#).

When VEPA is configured to switch traffic between the functions instead of VEB, the VF link state always reflects the physical state of the associated port and this behavior cannot be changed.

2.3.7.5 Configuring VFs

In operating system distributions with newer IP commands that support VF configuration options, the host administrator can perform the following by using the `ip link set` command:

- Change the default MAC address
- Configure VLAN
- Configure the transmission rate
- Configure the link state
- Set the QoS parameter on VFs

MAC Address Configuration:

PPC CNAs are shipped with factory-configured MAC addresses for the network interfaces corresponding to the PFs. The driver generates random MAC addresses for the network interfaces corresponding to the VFs based on the factory-configured MAC address. Other MAC addresses can be assigned for the interfaces corresponding to the VF using IP utility commands in the hypervisor.

To configure the MAC address for the virtual function, run the following command in the hypervisor:

```
# ip link set eth<X> vf <VFN> mac <MAC-ADDR>
```

where:

- `eth<X>` is the interface corresponding to the physical function.
- `<VFN>` is the VF number (0-based) corresponding to the interface for which you are configuring the MAC address.
- `<MAC-ADDR>` is the MAC address you are configuring.

For example, to configure the MAC address for VF number 0 on eth0(PF), run the following command in the hypervisor:

```
# ip link set eth0 vf 0 mac 00:16:88:AA:BB:AA
```

If the VM is already running and the VF driver is loaded, the VF MAC address can be updated using one of the following two methods in the VM:

- If the VF interface is configured using `ifcfg` script:
 - a. Update the `ifcfg` script for the VF, with the MAC address configured in the hypervisor in the earlier step. For example, update the `HWADDR` line in the VF `ifcfg-eth<X>` script.

```
HWADDR=00:16:88:AA:BB:AA
```

- b. Restart the network service.

```
# service network restart
```

Or

- If the VF interface is not configured using `ifcfg` script:
Update the VF MAC address using the `ifconfig` or `ip link` command.

```
# ifconfig eth0 hw ether 00:16:88:AA:BB:AA
```

VLAN Configuration:

This section includes examples for transparent tagging and virtual guest tagging (VGT).

When transparent tagging is configured for a VF, the NIC transparently tags all non-VLAN traffic from the VF with the configured transparent VLAN ID. The VM is not aware of the VLAN tag.

To assign a transparent VLAN ID to the VF, run the following IP command in the hypervisor:

```
# ip link set eth<X> vf <VFN> vlan <VLAN>
```

where:

- `eth<X>` is the interface corresponding to the physical function,
- `<VFN>` is the VF number corresponding to the interface for which you are configuring the VLAN,
- `<VLAN>` is the VLAN ID you are configuring,

Example:

```
# ip link set eth0 vf 0 vlan 5
```

NOTE: Guest VLAN tagging is not allowed on the VF in the guest operating system, when transparent VLAN tagging is enabled on the VF.

Transmission Rate Configuration:

Configure the transmission rate limit (TX-RATE) on a VF interface from the hypervisor using the following IP command syntax:

```
# ip link set eth<X> vf <VFN> rate <TX-RATE>
```

where:

- `eth<X>` is the interface corresponding to the physical function.
- `<VFN>` is the VF number corresponding to the interface that you are configuring the TX-RATE.
- `<TX-RATE>` is the transmission rate limit, in Mb/s.

For example, to set a TX-RATE of 5000 Mb/s for the VF 0, run the following command:

```
# ip link set eth0 vf 0 rate 5000
```

2.3.7.6 Link State Configuration

NOTE: The link state configuration feature is supported on RHEL 7.x and SLES 12 and later and is not supported in earlier kernels.

This section provides link state configuration for a VF.

Configure the link state on a VF interface from the hypervisor using the following IP command syntax:

```
# ip link set eth<X> vf <VFN> state < auto | enable | disable >
```

where:

- `eth<X>` is the interface corresponding to the physical function.
- `<VFN>` is the VF number corresponding to the interface for which the link state is being configured.
- `auto` – VF link state will reflect the PF link state.
- `enable` – VF link state will be always up.
- `disable` – VF link state will be always down.

2.3.7.7 Spoof Check Configuration

Configure the HW packet source MAC spoof check on a VF interface from the hypervisor using the following IP command syntax:

```
# ip link set eth<X> vf <VFN> spoofchk [on|off]
```

where:

- `eth<X>` is the interface corresponding to the physical function.
- `<VFN>` is the VF number corresponding to the interface for which you are configuring the spoof check.

For example, to turn on spoof checking for the VF 0, run the following command:

```
# ip link set eth0 vf 0 spoofchk on
```

2.3.7.8 Viewing VF Properties

To view the properties configured to VFs attached to a PF, use the following IP command syntax:

```
# ip link show eth<X>
```

where `eth<X>` is the interface corresponding to the physical function.

For example, to view the properties of a PF at `eth0` (which has VFs 0, 1 associated with it), run the following command in the hypervisor:

```
# ip link show eth0
```

Expected example output:

```
eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP qlen 1000
link/ether 00:00:c9:bb:16:ee brd ff:ff:ff:ff:ff:ff
vf 0 MAC 00:00:c9:9d:90:80, tx rate 10000 (Mbps)
vf 1 MAC 00:00:c9:9d:90:81, tx rate 10000 (Mbps)
```

2.3.8 Bonding Considerations

VFs can be added to a bonding configuration. But because VFs are typically assigned to a VM, it might not be possible to identify two VFs from the same PF or physical port, which could lead to incorrect configurations. To avoid this, the driver exports a physical port identifier corresponding to a given network interface. VFs from the same physical port will have the same identifier.

For example, to obtain the physical port identifier of a VF `eth0` in a VM, read the corresponding `sysfs` file `phys_port_id` in the VM as shown here. The number in the output shown is the physical port identifier of `eth0`.

```
# cat /sys/class/net/eth0/phys_port_id
0100000000000000000000000000000000383838314643
```

The NetworkManager reports incorrect bonding configuration with a message like this:

```
"NetworkManager: <warn> (bond0): slave eth1 shares a physical port with existing slave eth0"
```

2.4 FC and FCoE Driver Performance Tuning

This section describes how to tune the FC and FCoE driver for best performance.

2.4.1 Overview

The following configurable parameters can enhance performance:

- `lpfc_fcp_io_channel`
- `lpfc_fcp_io_sched`
- `lpfc_fcp_imax`

These features are available through module parameters that are defined in the FC and FCoE driver as well as `sysfs` entries defined by the Linux kernel.

In addition, you can use the `lpfc_vector_map.sh` script to map a specific I/O channel to a specific CPU. Determination of the mapping of the I/O channel to a specific CPU is also handled by the driver.

This section provides more information about how the tuning parameters and script can improve Emulex adapter performance.

The parameters in this section do not apply to LPe12000-series adapters.

2.4.1.1 `lpfc_fcp_io_channel`

The `lpfc_fcp_io_channel` module parameter can be configured at driver load time. It defines the number of I/O channels supported by the driver. The driver is capable of supporting parallel I/O paths, and each I/O path is capable of posting and completing FCP commands independent of the other.

Emulex adapters that are running in MSI-X interrupt mode can use more than one I/O path. Each I/O channel is composed of a unique MSI-X vector- EQ/CQ/WQ tuple. This parameter will override the value of the `lpfc_fcp_eq_channel` parameter.

NOTE: The Emulex LPe12000 adapters support only one I/O path, so this parameter has no effect on them.

By default, the driver is configured for four I/O channels per port. The driver will also limit the number of I/O channels to not exceed the number of online *logical* CPUs (as reported by `/proc/cpuinfo`).

2.4.1.2 `lpfc_fcp_io_sched`

The `lpfc_fcp_io_sched` module parameter can be configured at driver load time. It also can be set dynamically as an `sysfs` entry. The driver uses the parameter to determine which algorithm to use when scheduling an FCP I/O to an I/O channel.

When multiple I/O channels are in use, I/Os can be scheduled to an I/O channel in a round-robin fashion, or by determining which CPU is running when the I/O is submitted.

The default value (0) configures the driver for round-robin scheduling. A value of 1 configures scheduling by CPU.

NOTE: This parameter is only applicable for LPe16000, LPe31000, and LPe32000-series adapters and by default the value is configured to CPU scheduling.

2.4.1.3 lpfc_fcp_imax

The `lpfc_fcp_imax` can be configured at driver load time. It also can be set dynamically as an `sysfs` entry. This parameter defines the maximum number of interrupts per second that each adapter port will support.

Considerations:

- The lower the value set, the more completions are coalesced by the adapter, which causes the driver to handle multiple completions under the context of one interrupt. The higher the value, the faster an interrupt is generated for a completed command. Therefore, a balanced or *tuned* system must be found.
- A lower value equals higher interrupt latency; a higher value equals lower interrupt latency.
- Faster completions consume more system resources and CPU cycles, as the overhead of one interrupt completes fewer commands. The value is divided by the number of I/O channels, and each I/O channel is separately configured for its own interrupt latency.

By default, the module parameter is configured for 50,000 interrupts per second per adapter port. Older driver versions have a default value of 10,000.

2.4.1.4 lpfc_vector_map.sh

The `lpfc_vector_map.sh` script uses kernel `sysfs` entry points to map a specific I/O channel (MSI-X vector-EQ/CQ/WQ tuple) to a specific CPU. Run the script immediately after the driver is loaded.

This script resides in `/usr/sbin/lpfc`. Run this script by adding the following line to `/etc/modprobe.d/lpfc.conf` and through `/etc/rc.d` for the initial boot:

```
install lpfc /sbin/modprobe --ignore-install lpfc; /usr/sbin/lpfc/lpfc_vector_map.sh
```

The script maps each interrupt vector allocated by the driver to a CPU, thereby spreading the interrupt load of the ports across multiple CPUs. Each vector, with its associated I/O channel, is sequentially mapped to a CPU in a round-robin fashion. The number of vectors assigned to each adapter port is defined by the `lpfc_fcp_io_channel` module parameter.

I/O channels, which correspond to vectors, are typically mapped to unique CPUs to enhance the ability of the driver to run multiple FCP commands in parallel. In addition, running this script forces I/O scheduling to be by CPU (`lpfc_fcp_io_sched = 1`), which increases performance when a specific I/O channel is mapped to a specific CPU.

The script has two modes of operation: Driver mode and HBA mode. By default, the script runs in Driver mode.

Driver Mode:

Driver mode maps all vectors for all driver ports, starting with CPU0, sequentially assigning a new CPU for each vector belonging to the entire driver. If more interrupt vectors exist than CPUs, the vector assignment wraps back to CPU0 as needed.

HBA Mode:

HBA mode maps all vectors for each specific adapter port, starting with CPU0, sequentially assigning a new CPU for each vector belonging to that adapter port.

2.5 Network Performance Tuning

This section describes optimizing network performance.

2.5.1 Memory Bandwidth Considerations

The availability of higher memory bandwidth leads to better network performance. The following sections describe how memory bandwidth can be increased.

2.5.1.1 Enabling Optimal Bandwidth Options

Most computers offer multiple distinct memory channels, or memory interleaves, that might not be enabled by default. Check the manufacturer's documentation and BIOS parameters for details on enabling optimal memory bandwidth options.

2.5.1.2 Populating DIMM Slots

Typically, all the dual inline memory module (DIMM) slots must be populated to make use of all the memory channels. As a general rule, using more DIMMs provides better performance by allowing a higher degree of memory-access interleaving to occur.

2.5.1.3 Disabling Memory Mirroring

Some servers might allow memory mirroring, where the total memory is divided in half and each location is stored twice. This situation allows fault recovery if one memory location detects an error, but it greatly reduces the perceived memory bandwidth of the system. Consider disabling memory mirroring if it is not needed.

2.5.1.4 Using a Fast Clock Speed for the Front Side Bus (FSB)

Nearly any desktop or low-end server has enough memory bandwidth for NIC PPC CNAs to support DMA at 20Gb/s of data (10Gb/s read, 10Gb/s write). However, most of the memory demands come from the processor accessing the data for either packet copies in the non-offloaded networking stack or application accesses. All processor memory accesses use the FSB. The clock speed of this bus is critical for allowing efficient memory bandwidth. A system with a faster processor FSB clock speed performs better than a system with a slower FSB clock speed.

2.5.2 Network Memory Limits

The default values of tunable parameters in the Linux network stack are optimal for most network applications involving several TCP/UDP streams. The optimal size for the network queues and buffers depend on several factors such as protocol, number of streams (connections), request size, and application behavior.

The following network configuration settings are a good combination to get the best unidirectional transmit and receive performance with six or more TCP connections/UDP streams:

```
echo 4096 87380 4194304 > /proc/sys/net/ipv4/tcp_rmem
echo 4096 16384 4194304 > /proc/sys/net/ipv4/tcp_wmem
echo 64000000 > /proc/sys/net/core/rmem_default
```

```
echo 64000000 > /proc/sys/net/core/rmem_max
echo 32000000 > /proc/sys/net/core/wmem_default
echo 32000000 > /proc/sys/net/core/wmem_max
```

These settings assume ideal conditions such as low latency, zero (or close-to-zero) packet loss in the network, enough free memory, and 10Gb/s path-to-peer system.

These `tcp_rmem` and `tcp_wmem` values are also the default values in recent Linux distributions. If your application requires best throughput with very small number of connections (less than four), it can help to increase the `tcp_rmem` and `tcp_wmem` to much larger values:

```
echo 4096 87380 16777216 > /proc/sys/net/ipv4/tcp_rmem
echo 4096 65536 16777216 > /proc/sys/net/ipv4/tcp_wmem
```

2.5.3 TCP Segmentation Offload (TSO)

In low-loss networks, TSO considerably improves performance, and therefore must be enabled. TSO is enabled by default in the OneConnect network driver.

The `/proc/sys/net/ipv4/tcp_tso_win_divisor` process variable controls how aggressive the network stack can be in making TSO requests. For low-loss networks, the TSO divisor values must be in the range of 2 to 16. In most distributions, the default value of 3 might be the optimal choice for a no-loss network.

Smaller divisor values result in larger TSO chunks and better throughput, as well as better CPU utilization. However, if the receiver or the network is dropping frames (too many retransmits on the transmit side as indicated by `netstat -st`), it might help to make TSO chunks smaller (by increasing the divisor value) or to turn TSO off. For example, to set the divisor level to a value of 8, run:

```
echo 8 > /proc/sys/net/ipv4/tcp_tso_win_divisor
```

To turn TSO on or off, run one of the following `ethtool` commands:

```
ethtool -K eth<N> tso on
ethtool -K eth<N> tso off
```

where `eth<N>` is the name of the Ethernet device you are working on (for example, `eth0`).

2.5.4 Flow Control

PPC CNAs support IEEE 802.3x standard flow control, which uses control packets to temporarily pause the transmission of packets between two endpoints. These control messages are point-to-point and are not forwarded by switches or routers. The adapter can respond to flow control packets by temporarily pausing transmits. The adapter can send flow control pause packets when the transmitter is overwhelming the system's receive bandwidth.

Flow control can greatly improve performance, as described in the following examples:

- The adapter is installed in 4x PCIe slot or an underpowered server system.
If the PCIe bus does not provide 10Gb/s of throughput due to chipset limitations or the bus width, the adapter cannot maintain 10Gb/s of incoming receive data. It starts dropping packets quickly. In this situation, it may be beneficial to enable receive flow control in the adapter, and enable flow control in the attached switch for all devices. This action helps to slow down the transmitters.
- The adapter transmits to 1Gb/s devices, especially when using a non-TCP protocol.
If the adapter transmits to a 10Gb/s switch with attached 1Gb/s clients, the adapter can overwhelm the switch. The switch is then forced to start dropping packets because, although it can receive a 10Gb/s stream, the client can only sink a 1Gb/s stream. In this situation, it might be beneficial to enable transmit flow control in the adapter and enable flow control for the 10Gb/s switch port.

You can configure the adapter to respond to flow control pause frames from the other side (switch or router) using the following `ethtool` commands:

```
ethtool -A eth<N> pause rx on
ethtool -A eth<N> pause rx off
```

where `eth<N>` is the name of the Ethernet device you are working on (for example, `eth0`).

You can configure the adapter to send flow control pause frames using the following `ethtool` commands:

```
ethtool -A eth<N> pause tx on
ethtool -A eth<N> pause tx off
```

where `eth<N>` is the name of the Ethernet device you are working on (for example, `eth0`).

RX and TX flow control are enabled by default. When priority flow control is enabled in the adapter, normal flow control cannot be enabled.

Refer to the switch or router documentation to determine how link level flow control can be configured on the switch or router to which the adapter port is connected.

NOTE: In multichannel configurations where multiple PCI functions are exposed for a single 10GbE port, the flow control parameter for a port can be configured through any interface associated with the physical port, and the configured property will apply to all interfaces associated with the port.

2.5.5 RX Frame Coalescing

The Ethernet driver coalesces regular-sized TCP segments to a large frame before passing it to the network stack, which might improve TCP receive performance. RX frame coalescing is implemented using the GRO mechanism (in Linux driver versions that support GRO) or the LRO mechanism (in older Linux driver versions).

RX frame coalescing is enabled by default. In some configurations where the endpoint for the TCP connection to which the packets belong is not in the current server (for example, the endpoint is a router), RX coalescing should not be enabled.

GRO can be disabled using the `-K` option with the `ethtool` command:

```
ethtool -K eth<N> gro off
```

LRO can be disabled using the `-C` option with the `ethtool` command:

```
ethtool -C eth<N> rx-frames 1
```

where `eth<N>` is the name of the Ethernet device you are working on (for example, `eth0`).

2.5.6 Maximum Transmission Unit (MTU)

The Ethernet driver supports MTUs between 256 bytes and 9000 bytes. The default MTU is set to 1500. If other elements in the network path support a larger MTU, you can increase the MTU up to 9000 using the `ifconfig` command. To do this run:

```
ifconfig eth<N> mtu 9000
```

where `eth<N>` is the name of the Ethernet device you are working on (for example, `eth0`).

2.5.7 Interrupt Coalescing

The Ethernet driver tries to reduce the number of interrupts by delaying the interrupts from the adapter, which reduces CPU utilization during a high traffic rate. The interrupt delay duration can be set to change dynamically within a range of values, depending on the receive rate (known as Adaptive Interrupt Coalescing (AIC)), or can be set to a constant value.

2.5.7.1 Setting the Interrupt Delay Duration to a Range of Values (AIC)

For receive interrupts, AIC is enabled by default. When AIC is enabled, the default low limit is 0 microseconds and the default high limit is 96 microseconds. In low traffic, the interrupt delay is set to 0 for best latency. As the number of interrupts per second increases, the delay is increased to higher values proportional to the receive rate, up to the default high limit of 96 microseconds. You can change the low and high limits using `ethtool`. For example, to set a low limit of 8 and a high limit of 40, run:

```
ethtool -C eth<N> rx-usecs-low 8
ethtool -C eth<N> rx-usecs-high 40
```

where `eth<N>` is the name of the Ethernet device you are working on (for example, `eth0`).

For transmit interrupts, AIC is not supported.

To disable AIC and set the interrupt delay duration to a constant value, see the following section.

2.5.7.2 Setting the Interrupt Delay Duration to a Constant Value

The interrupt delay duration can be set to a constant value for both receive and transmit interrupts. The possible interrupt delay duration values are 0 to 96 microseconds, in 8 microsecond increments.

For receive interrupts, disable AIC (because it is enabled by default) and set the interrupt delay duration using `ethtool`. For example, to disable AIC and set the constant RX interrupt delay to 8-microseconds, run:

```
ethtool -C eth<N> adaptive-rx off rx-usec 8
```

where `eth<N>` is the number of the Ethernet interface you are working on.

If your application requires low or predictive latency, turn off AIC and set `rx-usecs` to 0.

For transmit interrupts, the default interrupt delay duration is 96 microseconds. You can change this value using `ethtool`. For example, to set the transmit interrupt delay to 64 microseconds, run:

```
ethtool -C eth<N> tx-usec 64
```

where `eth<N>` is the number of the Ethernet interface you are working on.

2.5.8 Receive-Side Scaling (RSS)

Distributing the incoming traffic across several receive rings with separate interrupt vectors helps to distribute the receive processing across several CPU cores. This distribution could reduce the packet drop and improve the packet rate in certain applications. RSS is enabled in non-SR-IOV and non-multichannel configurations. In multichannel configurations, RSS is enabled in the first section of each port.

2.5.9 Analyzing Performance Issues

MSI-x interrupts are required for RSS to work. If your motherboard and operating system version support MSI-X, the Ethernet driver automatically uses MSI-X interrupts. If not enough MSI-X vectors are available, the Ethernet driver uses INTx interrupts, which might decrease performance. The proc node `/proc/interrupts` shows the interrupts and their types.

The Linux performance `top` utility can monitor the CPU utilization while troubleshooting performance issues. A low idle CPU percentage in any CPU core is an indication of excessive processing load for that CPU. The proc node `/proc/interrupts` shows the distribution of the interrupts across the CPU cores. If you see too many interrupts per second directed to one CPU, check to see if the `irqbalance` program is running. The `irqbalance` program is normally started at system boot. In some cases, you can get better performance by disabling `irqbalance` and manually distributing interrupts. You can manually distribute the interrupt load across the available CPU cores by setting the CPU affinity for any interrupt vector by setting the mask in the proc node `/proc/irq/<int-vector>/smp_affinity`.

Use the `netstat` command to look for excessive TCP retransmits or packet drops in the network stack.

In systems having more than one NUMA node, you can get better performance by pinning interrupts to the NUMA node local to the PCIe device.

Use the `-s` option of `ethtool` to see all statistics counters maintained by the Ethernet and driver. Excessive drop or error counters are an indication of a bad link or defective hardware. See [Table 11, Ethtool -S Option Statistics](#), and [Table 12, Transmit/Receive Queue Statistics](#).

Turning off auditing and SELinux can improve CPU utilization and, in some cases, increase throughput. You can disable auditing by appending `audit=0` in the boot command line. You can turn off SELinux by specifying `selinux=0` in the boot command line. For example, the following command boots the Linux kernel with the SELinux and auditing options disabled:

```
kernel /boot/vmlinuz-2.6.18 ro root=/dev/md0 selinux=0 audit=0
```

You can get better CPU utilization and, in some cases, better throughput, by disabling kernel debug options, such as `CONFIG_DEBUG_SLAB`. This situation requires you to build the kernel image and modules. Turning off the firewall and disabling `Hyper-Threading` can also improve performance.

Chapter 3: Troubleshooting

This section explains some of the situations in which your system can operate in an unexpected manner and some possible resolutions.

3.1 Hardware Situations and Resolutions

[Table 6: Hardware Situations and Their Resolutions](#) lists the hardware situations and their resolutions.

Table 6: Hardware Situations and Their Resolutions

Situation	Resolution
An unapproved optical transceiver is used to connect the adapter.	<p>The system log can generate one or more of these events if an unapproved optical transceiver is detected:</p> <ul style="list-style-type: none"> ■ Unsupported module ■ Optics faulted/incorrectly installed/not installed ■ Incompatible optics ■ Unknown port SFB status <p>To resolve this issue, do the following:</p> <ol style="list-style-type: none"> 1. Power off the system. 2. Replace the unapproved optical transceiver with an approved one. 3. Power on the system. <p>NOTE: For more information on replacing an optical transceiver, refer to the specific adapter hardware guide available at the Documents and Downloads area of www.broadcom.com.</p>

3.2 FC and FCoE Driver Situations and Their Resolutions

[Table 7: FC and FCoE Driver Situations and Their Resolutions](#) lists the FC and FCoE driver situations and their resolutions.

Table 7: FC and FCoE Driver Situations and Their Resolutions

Situation	Resolution
FC link fails to come up.	If an FC link fails to come up, verify that an 8Gb/s or 16Gb/s adapter is not attempting to connect to a 1Gb/s device. Only 2Gb/s, 4Gb/s, and 8Gb/s devices are supported on 8Gb/s adapters. Only 2Gb/s, 4Gb/s, 8Gb/s, and 16Gb/s devices are supported on 16Gb/s adapters.
Authentication is enabled but authentication service is not running error message.	<p>If you see this message in <code>/var/log/messages</code> and the adapter is in an error state, the <code>fcauthd</code> daemon probably is not running. To determine whether <code>fcauthd</code> is running, run:</p> <pre>/etc/init.d/fcauthd status.</pre> <p>To start <code>fcauthd</code>, run</p> <pre>/etc/init.d/fcauthd start.</pre>
If a SAN configuration has 256 targets mapped by the FC and FCoE driver, any additional added targets do not get a target ID mapping by the driver and cause target discovery to fail.	<p>Removing targets or reinitializing the link does not solve this issue. Unload and reload the driver to reset available target IDs. Ensure that the SAN configuration is correct prior to rebooting the driver. This action clears the driver consistent binding table and frees target IDs for new target nodes.</p>

Table 7: FC and FCoE Driver Situations and Their Resolutions (Continued)

Situation	Resolution
<p><code>rmmmod</code> fails to unload FC and FCoE driver module due to <code>ERROR: Module lpfc is in use.</code></p>	<p>This message can appear when you attempt to remove the driver and a Logical Volume Group is dependent on the driver.</p> <p>Driver unload is not supported in SAN booted environments unless the inbox driver supports SAN boot. Otherwise, use these steps to resolve this situation:</p> <ol style="list-style-type: none"> 1. Make the Logical Volume Group unavailable. Type: <code>lvchange -a n xxxxxxxx</code> The <code>xxxxxxx</code> parameter is the Volume Group Name. 2. Stop the Emulex OneCommand Manager application. 3. Stop Device Mapper.
<p><code>rmmmod</code> hangs and module reference count is 0.</p>	<p>Due to a small race condition in the kernel, it is possible for an <code>rmmmod</code> command to hang. Issue the <code>rmmmod -w</code> command. If this does not help, reboot the computer.</p>
<p><code>rmmmod</code> fails to unload driver due to device or resource busy.</p>	<p>This message occurs when you attempt to remove the driver without first stopping the Emulex OneCommand Manager application or the <code>fcauthd</code> daemon when the Emulex OneCommand Manager application is installed and running, or when FC disks connected to an Emulex LightPulse adapter are mounted. To resolve this situation:</p> <ol style="list-style-type: none"> 1. Stop the Emulex OneCommand Manager application before attempting to unload the driver. The script is located in the <code>/usr/sbin/ocmanager</code> directory. Type: <code>./stop_ocmanager</code> 2. Unmount any disks connected to the adapter. 3. Unload the driver. Type: <code>rmmmod lpfc</code>
<p>An <code>lspci</code> shows recent Emulex adapters as unknown.</p>	<p>This situation occurs because of the delay of getting new product IDs into the Linux development cycle.</p> <p>There is no resolution at this time.</p>
<p>Slow targets or extended link faults on the storage side might result in storage being marked offline by the mid-level and remaining offline (not recovered) when the link faults are corrected.</p>	<p>This version of the driver should eliminate this issue. However, if you experience offline device issues, increase the SCSI command timeout to a value greater than or equal to 60 seconds. Emulex also provides a script that addresses this issue.</p> <p>To access the <code>lun_change_state.sh</code> script, go to the Documents and Downloads area of www.broadcom.com.</p>
<p>Under certain conditions of an I/O load, some targets cannot complete an I/O issued by a Linux initiator within the default timeout of 30 seconds given by the SCSI mid-level.</p>	<p>If the situation is not corrected, the initiator-to-target condition deteriorates into abort/recovery storms, leading to I/O failures in the block layer. These types of failures are preceded by a SCSI I/O error of hex 6000000.</p> <p>Emulex provides a script that addresses this issue.</p> <p>To access the <code>set_target_timeout.sh</code> script, go to the Documents and Downloads area of www.broadcom.com.</p>
<p>The FC and FCoE driver fails to recognize an adapter and logs <code>unknown IOCB</code> messages in the system log during driver load.</p>	<p>The adapter is running outdated firmware.</p> <p>Install the latest firmware on the adapter.</p> <p>NOTE: Before performing a firmware update, driver installation is required. For more information on installing the driver, see Section 1.2.1, Installing the Binary RPM FC and FCoE Driver Kit.</p>
<p>Loading the FC and FCoE driver on SLES 11 SPx reports <code>unsupported module, tainting kernel</code> in system log.</p>	<p>This message is logged by the kernel whenever a module that is not shipped with the kernel is loaded.</p> <p>Ignore this message.</p>
<p>The system panics when it is booted with a failed adapter installed.</p>	<p>Remove the failed adapter and reboot the system.</p>

Table 7: FC and FCoE Driver Situations and Their Resolutions (Continued)

Situation	Resolution
Unloading the FC and FCoE driver on SLES 11 SPx might cause a message to be logged in the system log, such as the following: <pre>umount: /dev/disk/bypath/pci-0000:02:04.0-scsi-0:0:1:0: not mounted</pre>	These messages are normal output from the SLES 11 SPx hotplug scripts. Ignore them.
Driver installation fails.	The <code>lpfc-install</code> script fails to install the driver. The install script can fail for the following reasons: <ul style="list-style-type: none"> ■ A previous version of the driver is installed. Run the <code>lpfc-install --uninstall</code> script and then try to install the driver. ■ The current driver is already installed. ■ Run a supported RHEL or SLES kernel.
No module <code>lpfc</code> found for kernel <code>KERNELVERSION</code> RPM error message when upgrading the kernel.	These three situations can be resolved by upgrading the kernel. There are two ways to install the driver into an upgraded kernel. The method you use depends on whether you are updating the driver.
A recently upgraded kernel cannot find the ramdisk. After upgrading the kernel, the kernel cannot find the ramdisk, which halts or panics the system.	<ul style="list-style-type: none"> ■ Upgrade the kernel using the same version of the driver. ■ Upgrade the kernel using a new version of the driver. For more information on upgrading the kernel, see Chapter 1, Installing and Uninstalling .
The driver is not loaded after a system reboot after upgrading the kernel.	
Driver uninstall fails.	The <code>lpfc-install --uninstall</code> script fails with an error. Try the following solutions: <ul style="list-style-type: none"> ■ Uninstall the Emulex OneCommand Manager application; refer to the <i>OneCommand Manager Application for LightPulse Adapters User Guide</i> for instructions. ■ Unmount all FC disk drives. ■ Unload the <code>lpfcdfc</code> and FC and FCoE driver. ■ Use <code>rpm -e lpfcdriver</code> and <code>-e ocmanager</code> and uninstall the new kits.
<code>lpfc-install</code> script exit code.	The <code>lpfc-install</code> script contains exit codes that can be useful in diagnosing installation issues. Refer to the <code>lpfc-install</code> script for a complete listing of codes and definitions.
The Emulex driver for Linux does not load in ramdisk for a custom-built kernel.	Custom built kernels are not supported by Emulex. However, the Emulex install script attempts to install the driver into a ramdisk that follows the naming scheme used by RHEL or SLES kernels. <ul style="list-style-type: none"> ■ The RHEL naming scheme for ramdisk images is: <code>/boot/initrd-KERNELVERSION.img</code>. ■ The SLES naming scheme for ramdisk images is: <code>/boot/initrd</code>. If a custom built kernel has a ramdisk image that does not follow the appropriate naming scheme, change the name of the image using the following procedure: <ol style="list-style-type: none"> 1. Change the name of the ramdisk image to match the SLES naming scheme. 2. Update any file links to the ramdisk image. 3. Edit the boot loader configuration file (for example, <code>/etc/lilo.conf</code>, <code>/etc/yaboot.conf</code>, <code>/boot/grub/grub.conf</code>, <code>/boot/grub/menu.lst</code>), find any references to the old ramdisk image name, and replace them with the new name. 4. Reboot the system to verify the changes. 5. Install the Emulex Linux driver kit.

Table 7: FC and FCoE Driver Situations and Their Resolutions (Continued)

Situation	Resolution
The Linux SCSI subsystem sees only eight LUNs when more are present.	Some SCSI drivers do not scan past eight LUNs when the target reports itself as a SCSI-2 device. To resolve this situation, force a SCSI bus scan with the following command: <code>/usr/sbin/lpfc/lun_scan</code> . SUSE supplies a <code>/bin/rescan-scsi-bus.sh</code> script, which can be changed to scan everything.

3.3 Ethernet Driver Situations and Their Resolutions

The following table lists the Ethernet driver situations and their resolutions. This section applies to PPC CNAs only.

Table 8: Ethernet Driver Situations and Their Resolutions

Situation	Resolution
The <code>ethtool</code> configuration settings are not restored after system reboot.	The <code>ethtool</code> settings are not designed to persist across reboot. For persistence, invoke configuration commands from a boot script that is executed at system start, such as <code>/etc/rc.local</code> .
The Ethernet driver works but the transmit and receive data rates are not near a 10Gb/s line rate.	There could be several reasons for poor performance. For best performance practices, see Section 2.5, Network Performance Tuning .
When <code>MILI</code> and <code>SNMP</code> daemons start, they trigger warning messages within <code>SELinux</code> for certain operations.	This issue is known and no solution is available. However, to avoid <code>SELinux</code> warning messages, you can disable <code>SELinux</code> . To disable <code>SELinux</code> , open a terminal and enter the following command at the prompt: <code>echo 0 > /selinux/enforce</code> To enable <code>SELinux</code> , use the following command: <code>echo 1 > /selinux/enforce</code> In addition, see Section 2.5.9, Analyzing Performance Issues .

3.4 Log Messages

3.4.1 FC and FCoE Driver Log Messages

The following section describes retrieving and interpreting FC and FCoE driver log messages.

3.4.1.1 Retrieving FC and FCoE Driver Log Messages

LPFC error log messages are logged in the `/var/log/messages` file.

An example of an LPFC message:

```
Jul 2 04:23:34 daffy kernel: lpfc 0000:03:06.0: 0:1305 Link Down Event x2f2 received Data: x2f2 x20
x110
```

In this example:

- `lpfc 0000:03:06.0` – Identifies the PCI location of the particular LPFC hardware port.
- `0:` – Indicates Emulex adapter 0.
- `1305` – Indicates a log message number of 1305.

NOTE:

- If `Data :` is present in a log message, any information following `Data :` is intended only for Broadcom® Technical Support or Engineering use.
- If an error message instructs you to perform a firmware update, ensure that the driver is installed first. For more information on installing the driver, see [Chapter 1, Installing and Uninstalling](#).

3.4.1.2 LPFC Error Log Messages and their Descriptions

The following table lists the LPFC error log messages and their descriptions.

Table 9: LPFC Error Log Messages and Their Descriptions

<p>0111: Dropping received ELS cmd The driver decided to drop an ELS Response ring entry. Data: (1) ulpStatus (2) ulpWord[4] (3) ulpTimeout Severity: Error Log: Always Action: This error could indicate a software driver or firmware issue. If this issue persists, report these errors to Broadcom Technical Support.</p>
<p>0113: An FLOGI ELS command <elsCmd> was received from DID <did> in Loop Mode While in Loop Mode an unknown or unsupported ELS command was received. Data: None Severity: Error Log: Always Action: Check device DID.</p>
<p>0115: Unknown ELS command <elsCmd> received from NPORT <did> Received an unsupported ELS command from a remote N_Port. Data: None Severity: Error Log: Always Action: Check the remote N_Port for a potential issue.</p>
<p>0125: FDISC Failed (x%x). Fabric out of resources The fabric rejected an FDISC because the switch cannot support additional virtual ports. Data: lsRjtError Severity: Error Log: Always Action: Reconfigure the switch to support more NPIV logins. If this issue persists, contact Broadcom Technical Support.</p>
<p>0126: FDISC failed ulpStatus ulpWord4 Data: lsRjtError Severity: Error Log: Always Action: Reconfigure the switch to support more NPIV logins. If this issue persists, contact Broadcom Technical Support.</p>
<p>0127: ELS timeout An ELS IOCB command was posted to a ring and did not complete within ULP timeout seconds. Data: (1) elscmd (2) remote_id (3) ulpcommand (4) ulpIoTag Severity: Error Log: Always Action: If no ELS command is going through the adapter, reboot the system. If the issue persists, contact Broadcom Technical Support.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0133: PLOGI: no memory for reg_login Memory allocation error. Data: (1) nlp_DID (2) nlp_state (3) nlp_flag (4) nlp_rpi Severity: Error Log: LOG_ELS Action: Memory allocation error. Check system resources. Unload unused modules.</p>
<p>0134: PLOGI cannot issue reg_login The ELS PLOGI mailbox command has failed. Data: (1) nlp_DID (2) nlp_state (3) nlp_flag (4) nlp_rpi Severity: Error Log: LOG_ELS Action: Check the port and switch configuration.</p>
<p>0135: cannot format reg_login Could not allocate an RPI or DMA buffer for the mailbox command. Data: (1) nlp_DID (2) nlp_state (3) nlp_flag (4) nlp_rpi Severity: Error Log: LOG_ELS Action: None required.</p>
<p>0136: PLOGI completes to NPort <DID> completion A PLOGI has completed for which there is no NDLP. Data: (1) ulpStatus (2) ulpWord[4] Severity: Error Log: LOG_ELS Action: None required.</p>
<p>0137: No retry ELS command <ELS_CMD> to remote Data: (1) ulpStatus (2) ulpWord[4] Severity: Error Log: LOG_ELS Action: None required.</p>
<p>0138: ELS rsp: Cannot issue reg_login for <DID> REG_LOGIN mailbox command failed. Data: (1) nlp_DID (2) nlp_state (3) nlp_flag (4) nlp_rpi Severity: Error Log: LOG_ELS Action: None required.</p>
<p>0139: Ignoring ELS cmd tag <ioTag> completion Data This ELS command was aborted. Data: (1) ulpStatus (2) ulpWord[4] (3) ulpTimeout Severity: Error Log: LOG_ELS Action: None required.</p>
<p>0140: PLOGI Reject: invalid name Invalid node WWN provided. Data: None Severity: Error Log: LOG_ELS Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0141: PLOGI Reject: invalid pname Invalid port WWN provided. Data: None Severity: Error Log: LOG_ELS Action: None required.</p>
<p>0142: PLOGI RSP: Invalid WWN The PLOGI sent to the port by a remote port had an invalid WWN. Data: None Severity: Error Log: LOG_ELS Action: None required.</p>
<p>0143: SLI4 Adapter Hardware Error Data: <status0>/<status1> The HBA has encountered an unrecoverable error. Data: None Severity: Error Log: LOG_INIT Action: Use hbacmd to retrieve a dump file.</p>
<p>0144: Not a valid WCQE code: <Completion Code> The completion queue handler detected an invalid type. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>0147: Failed to allocate memory for RSCN event Memory could not be allocated to send the RSCN event to the management application. Data: None Severity: Error Log: LOG_ELS Action: None required.</p>
<p>0148: Failed to allocate memory for LOGO event Memory could not be allocated to send the LOGO event to the FC transport. Data: None Severity: Error Log: LOG_ELS Action: None required.</p>
<p>0149: Failed to allocate memory for ELS event Memory could not be allocated to send the ELS event to the FC transport. Data: None Severity: Error Log: LOG_ELS Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0154: Authentication not complete Authentication was restarted because the previous authentication did not complete. Data: None Severity: Error Log: LOG_DISCOVERY Action: Check the switch configuration.</p>
<p>0200: CONFIG_LINK bad hba state <hba_state> A CONFIG_LINK mailbox command completed and the driver was not in the right state. Data: None Severity: Error Log: Always Action: Software driver error. If this issue persists, report these errors to Broadcom Technical Support.</p>
<p>0203: Devloss timeout on WWPN <address> NPort <nlp_DID> A remote N_Port that was discovered by the driver disappeared for more than lpfc_devloss_tmo seconds. Data: (1) nlp_flag (2) nlp_state (3) nlp_rpi Severity: Error Log: Always Action: If the device generating this message is not a target to which the HBA is connected, this error will not affect the data integrity of the I/O between the HBA and the attached storage and can be ignored.</p>
<p>0206: Device discovery completion error This indicates that an uncorrectable error was encountered during device discovery after a link up. FC devices will not be accessible if this message is displayed. Data: None Severity: Error Log: Always Action: Reboot the system. If this issue persists, report the error to Broadcom Technical Support. Run with verbose mode enabled for more information.</p>
<p>0207: Device <DID> (<WWN>) sent invalid service parameters. Ignoring device. Invalid service parameters were received from DID. This remote port is ignored. Data: DID, WWN Severity: Error Log: Always Action: Verify the remote port's configuration. If this issue persists, report the error to Broadcom Technical Support. Run with verbose mode on for more details.</p>
<p>0217: Block sgl registration required DMAsize <reqlen> great than a page The request to post SGL pages does not fit on a page. Data: None Severity: Warning Log: LOG_INIT Action: None required.</p>
<p>0221: FAN timeout A link up event was received without the login bit set, so the driver waits E_D_TOV for the fabric to send a FAN. If no FAN if received, a FLOGI will be sent after the timeout. Data: None Severity: Warning Log: LOG_DISCOVERY verbose Action: None required. The driver recovers from this condition by issuing a FLOGI to the fabric.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0222: Initial FLOG/FDISKI timeout The driver sent the initial FLOGI or FDISK to the fabric and never got a response back. Data: None Severity: Error Log: Always Action: Check fabric configuration. The driver recovers from this situation and continues with device discovery.</p>
<p>0223: Timeout while waiting for NameServer login Our login request to the NameServer was not acknowledged within R_A_TOV. Data: None Severity: Error Log: Always Action: Check the fabric configuration. The driver recovers from this and continues with device discovery.</p>
<p>0224: NameServer Query timeout Node authentication timeout, node Discovery timeout. A NameServer Query to the fabric or discovery of reported remote N_Ports is not acknowledged within R_A_TOV. Data: (1) fc_ns_retry (2) fc_max_ns_retry Severity: Error Log: Always Action: Check fabric configuration. The driver recovers from this situation and continues with device discovery.</p>
<p>0226: Device discovery completion error This indicates that an uncorrectable error was encountered during device discovery after a link up. FC devices will not be accessible if this message is displayed. Data: None Severity: Error Log: Always Action: Reboot the system. If this issue persists, report the error to Broadcom Technical Support. Run with verbose mode on for more details.</p>
<p>0227: Node Authentication timeout The driver has lost track of what N_Ports are being authenticated. Data: None Severity: Error Log: Always Action: None required. The driver should recover from this event.</p>
<p>0228: CLEAR LA timeout The driver issued a CLEAR_LA that never completed. Data: None Severity: Error Log: Always Action: None required. The driver should recover from this event.</p>
<p>0230: Unexpected timeout, hba linkstate <link_state> Discovery has timed out, and the HBA state is not ready. Data: None Severity: Error Log: LOG_DISCOVERY Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0231: RSCN timeout The driver has lost track of which N_Ports have RSCNs pending. Data: (1) fc_ns_retry (2) lpfc_max_ns_retry Severity: Error Log: Always Action: None required. The driver should recover from this event.</p>
<p>0233: Nodelist not empty Driver unloaded or hotplug detected a node still in use. Data: None Severity: Error Log: LOG_DISCOVERY Action: None required.</p>
<p>0237: Pending Link Event during Discovery: State <hba_state> Received link event during discovery. Causes discovery restart. Data: None Severity: Warning Log: LOG_DISCOVERY verbose Action: None required, unless this issue persists. If persistent, check cabling.</p>
<p>0241: NameServer rsp error The driver received a NameServer response containing a status error. Data: (1) CommandResponse.bits.CmdRsp (2) ReasonCode (3) Explanation (4) fc_flag Severity: Error Log: LOG_DISCOVERY verbose Action: Check the fabric configuration. The driver recovers from this situation and continues with device discovery.</p>
<p>0246: RegLogin failed The firmware returned a failure for the specified RegLogin. Data: (1) Did (2) mbxStatus (3) hbaState Severity: Error Log: Always Action: This message indicates that the firmware could not do RegLogin for the specified DID. There might be a limitation on how many nodes an HBA can see.</p>
<p>0249: Cannot issue Register Fabric login: Err %d\ Could not issue the fabric register login, the error value is unique for each possible failure. Data: None Severity: Error Log: LOG_ELS Action: None required.</p>
<p>0251: NameServer login: no memory Could not allocate memory for the NDLP structure. Data: None Severity: Error Log: LOG_ELS Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0252: Cannot issue NameServer login Could not issue an ELS PLOGI to the NameServer DID. Data: None Severity: Error Log: LOG_ELS Action: Check the port connection and the switch configuration.</p>
<p>0253: Register VPI: Can't send mbox\ Could not issue the REG_VPI mailbox command for this VPort. Data: None Severity: Error Log: LOG_MBOX Action: None required.</p>
<p>0254: Register VPI: no memory" goto mbox_err_exit Could not allocate memory for the REG_VPI mailbox command. Data: None Severity: Error Log: LOG_MBOX Action: None required.</p>
<p>0255: Issue FDISC: no IOCB All of the preallocated IOCBs are in use. Data: None Severity: Error Log: LOG_ELS Action: None required.</p>
<p>0256: Issue FDISC: Cannot send IOCB\ Unable to send the fabric IOCB. Data: None Severity: Error Log: LOG_ELS Action: None required.</p>
<p>0257: GID_FT Query error The GID_FT CT request for the NameServer has failed. Data: None Severity: Error Log: LOG_ELS Action: Check the switch configuration.</p>
<p>0258: Register Fabric login error: The REG_LOGIN for the fabric has failed. Data: None Severity: Error Log: LOG_MBOX Action: Check the port connection and the switch configuration.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0259: No NPIV Fabric support The switch to which the port is connected does not support NPIV. Data: None Severity: Error Log: LOG_ELS Action: Check the switch configuration.</p>
<p>0260: Register NameServer error: The REG_LOGIN mailbox command has failed for the NameServer. Data: None Severity: Error Log: LOG_ELS Action: Check the switch configuration</p>
<p>0261: Cannot register NameServer login: Either a memory allocation issue or an invalid parameter was sent to the REG_LOGIN. Data: None Severity: Error Log: LOG_ELS Action: At least one message (0142 0121 0133 0134 0135) should precede this message.</p>
<p>0262: No NPIV Fabric support The switch to which the port is connected does not support NPIV. Data: None Severity: Error Log: LOG_ELS Action: Check the switch configuration.</p>
<p>0263: Discovery Mailbox error: state: Either the driver could not allocate resources or it could not send sparam_mbox or cfglink_mbox. Data: (1) address of sparam_mbox command (2) address of cfglink_mbox command Severity: Error Log: LOG_MBOX Action: Attempt to unload and reload the driver when it is convenient.</p>
<p>0264: No NPIV Fabric support The switch to which the port is connected does not support NPIV. Data: None Severity: Error Log: LOG_ELS Action: Check the switch configuration.</p>
<p>0266: Issue NameServer Req <cmdcode> err <rc> Data: <fc_flag> <fc_rscn_id_cnt> The driver was unable to send the NameServer CT command. Data: (1) vports fc_flag (2) vports fc_rscn_id_cnt Severity: Error Log: LOG_DISCOVERY Action: Check the port and switch configurations.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0267: NameServer GFF Rsp <did> Error (<ulpStatus> <un.ulpWord[4]>) Data: <fc_flag> <fc_rscn_id_cnt> The NameServer GFF CT request failed. Data: (1) vports fc_flag (2) vports fc_rscn_id_cnt Severity: Error Log: LOG_DISCOVERY Action: Check the port and switch configurations.</p>
<p>0268: NS cmd <cmdcode> Error (<ulpStatus> <un.ulpWord[4]>) The NameServer CT request failed. Data: None Severity: Error Log: LOG_DISCOVERY Action: Check the port and switch configurations.</p>
<p>0271: Illegal State Transition: node <nlp_DID> event <evt>, state <nlp_state> Data: <nlp_rpi> <nlp_flag> The current node state does not have a handler for this event. Data: (1) nlp_rpi (2) nlp_flag Severity: Error Log: LOG_DISCOVERY Action: Verify that all targets are still visible to the SCSI mid-layer.</p>
<p>0272: Illegal State Transition: node <nlp_DID> event <evt>, state <nlp_state> Data: <nlp_rpi> <nlp_flag> The driver is completing a PLOGI but does not have the rcv_plogi flag set. Data: (1) nlp_rpi (2) nlp_flag Severity: Error Log: LOG_DISCOVERY Action: Verify that all targets are still visible to the SCSI mid-layer.</p>
<p>0273: Unexpected discovery timeout, vport State x%x The discovery process has timed out. Data: None Severity: Error Log: LOG_DISCOVERY Action: Verify that all targets are still visible.</p>
<p>0274: lpfc_nlp_put: ndlp:x%pusgmap:x%x refcnt:%d, void *)ndlp, ndlp->nlp_usg_map, atomic_read(&ndlp->kref.refcount) Data: one Severity: Warning Log: LOG_NODE Action: None required.</p>
<p>0275: lpfc_nlp_put: ndlp:x%pusgmap:x%x refcnt:%d, void *)ndlp, ndlp->nlp_usg_map, atomic_read(&ndlp->kref.refcount) A kref_put was called again after the node was already inactive. Data: None Severity: Warning Log: LOG_NODE Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0276: lpfc_nlp_get: ndlp:x%pugmap:x%x refcnt:%d, void *)ndlp, ndlp->nlp_usg_map, atomic_read(&ndlp->kref.refcount)</p> <p>A kref_get was attempted on a node that was being released.</p> <p>Data: None</p> <p>Severity: Warning</p> <p>Log: LOG_NODE</p> <p>Action: None required.</p>
<p>0277: lpfc_enable_node: ndlp:x%pugmap:x%x refcnt:%d, void *)ndlp, ndlp->nlp_usg_map, atomic_read(&ndlp->kref.refcount)</p> <p>Enable node was attempted on an inactive node.</p> <p>Data: None</p> <p>Severity: Warning</p> <p>Log: LOG_NODE</p> <p>Action: None required.</p>
<p>0278: lpfc_enable_node: ndlp:x%pugmap:x%x refcnt:%d, void *)ndlp, ndlp->nlp_usg_map, atomic_read(&ndlp->kref.refcount)</p> <p>Enable node was attempted on an inactive node.</p> <p>Data: None</p> <p>Severity: Warning</p> <p>Log: LOG_NODE</p> <p>Action: None required.</p>
<p>0280: lpfc_cleanup_node: ndlp:x%pugmap:x%x refcnt:%d, void *)ndlp, ndlp->nlp_usg_map, atomic_read(&ndlp->kref.refcount)</p> <p>Node clean-up was attempted on a node that has already been marked for memory free.</p> <p>Data: None</p> <p>Severity: Warning</p> <p>Log: LOG_NODE</p> <p>Action: None required.</p>
<p>0281: lpfc_cleanup_node: ndlp:x%pugmap:x%x refcnt:%d, void *)ndlp, ndlp->nlp_usg_map, atomic_read(&ndlp->kref.refcount)</p> <p>Node clean-up was called to prepare the node for release.</p> <p>Data: None</p> <p>Severity: Warning</p> <p>Log: LOG_NODE</p> <p>Action: None required.</p>
<p>0282: ldid:x%x ndlp:x%pugmap:x%x refcnt:%d, ndlp->nlp_DID, (void *)ndlp, lpfc_init.c-ndlp->nlp_usg_map,</p> <p>Driver clean-up has found a node that is still on the node list during driver unload or PCI hot plug removal.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_NODE</p> <p>Action: None required.</p>
<p>0283: Failed to allocate mbox cmd memory</p> <p>Mailbox allocation error.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_INIT</p> <p>Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0285: Allocated DMA memory size <alloclen> is less than the requested DMA memorysize<reqlen> Memory allocation was truncated.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_INIT</p> <p>Action: None required.</p>
<p>0286: lpfc_nlp_state_cleanup failed to allocate statistical data buffer <nlp_DID> Memory allocation failed for node's statistical data.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_INIT</p> <p>Action: None required.</p>
<p>0287: lpfc_alloc_bucket failed to allocate statistical data buffer DID <nlp_DID> Memory allocation failed for node's statistical data.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_NODE</p> <p>Action: None required.</p>
<p>0288: Unknown FCoE event type <event_type> event tag <event_tag> The firmware has detected an unknown FCoE event.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_SLI</p> <p>Action: Check the FCoE switch configuration and the HBA DCBX mode.</p>
<p>0289: Issue Register VFI failed: Err <rc> The driver could not register the Virtual Fabric Index for the FCFI.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_ELS</p> <p>Action: Check the switch and port configurations.</p>
<p>0290: The SLI4 DCBX asynchronous event is not handled yet The SLI-4 DCBX asynchronous event is not handled yet.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_SLI</p> <p>Action: None required.</p>
<p>0291: Allocated DMA memory size (x%x) is less than the requested DMA memory size (x%x) The asynchronous DCBX events are not handled in the driver.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_INIT</p> <p>Action: Check the switch configuration.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0293: PM resume failed to start worker thread: error=<error> The PCI resume (hot plug) could not start the worker thread for the driver. Data: None Severity: Error Log: LOG_INIT Action: Unload and reload the driver.</p>
<p>0294: PM resume Failed to enable interrupt The PCI resume (hot plug) could not get an interrupt vector. Data: None Severity: Error Log: LOG_INIT Action: Unload and reload the driver.</p>
<p>0297: Invalid device group<pci_dev_grp> While unloading the driver, the driver detected a PCI device that it should not have claimed. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>0299: Invalid SLI revision <sli_rev> While processing a host attention or unrecoverable error, the driver detected an invalid SLI revision. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>0300: LATT: Cannot issue READ_LA: Data:<rc> The link attention handler could not issue a READ_LA mailbox command. Data: None Severity: Error Log: LOG_MBOX Action: None required.</p>
<p>0301: READ_SPARAM: no buffers The driver attempted to issue a READ_SPARAM mailbox command to the adapter, but there were no buffers available. Data: None Severity: Warning Log: LOG_MBOX verbose Action: This message indicates one of these two issues: (1) Kernel virtual memory is depleted. Check that the system meets minimum RAM requirements for the adapter. Try closing other applications to free some memory. (2) A possible driver buffer management issue. If this issue persists, report the error to Broadcom Technical Support.</p>
<p>0302: REG_LOGIN: no buffers The driver attempted to issue a REG_LOGIN mailbox command to the adapter, but no buffers were available. Data: (1) Did, (2) flag Severity: Warning Log: LOG_MBOX verbose Action: This message indicates one of these two issues: (1) Kernel virtual memory is depleted. Check that the system meets minimum RAM requirements for the adapter. Try closing other applications to free some memory. (2) A possible driver buffer management issue. If this issue persists, report the error to Broadcom Technical Support.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0313: Ring <ringno> handler: unexpected Rctl <Rctl> Type <Type> received The RCTL/Type of a received frame did not match any for the configured masks for the specified ring. Data: None Severity: Warning Log: LOG_SLI verbose Action: This error could indicate a software driver, firmware, or hardware issue. Report these errors to Broadcom Technical Support.</p>
<p>0303: Ring <ringno> handler: portRspPut <portRspPut> is bigger then rsp ring <portRspMax> The port rsp ring put index is larger than the size of the rsp ring. Data: None Severity: Error Log: Always Action: This error could indicate a software driver, firmware, or hardware issue. Report these errors to Broadcom Technical Support.</p>
<p>0304: Stray mailbox interrupt, mbxCommand <mbxcommand> mbxStatus <mbxstatus> Received a mailbox completion interrupt and there are no outstanding mailbox commands. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.</p>
<p>0306: CONFIG_LINK mbxStatus error <mbxStatus> HBA state <hba_state> The driver issued a CONFIG_LINK mailbox command to the HBA that failed. Data: None Severity: Error Log: Always Action: This error could indicate a firmware or hardware issue. Report these errors to Broadcom Technical Support.</p>
<p>0310: Mailbox command <mbxcommand> timeout A mailbox command was posted to the adapter and did not complete within 30 seconds. Data: (1) hba_state (2) sli_flag (3) mbox_active Severity: Error Log: Always Action: This error could indicate a software driver or firmware issue. If no I/O is going through the adapter, reboot the system. If this issue persists, report the error to Broadcom Technical Support.</p>
<p>0312: Ring <ringno> handler: portRspPut <rspPutInx> is bigger then rsp ring <numRiobc> The IOCB command rings put pointer is ahead of the get pointer. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>0313: Ring <ringno> handler: unexpected Rctl <Rctl> Type <Type> received The RCTL/Type of a received frame did not match any for the configured masks for the specified ring. Data: None Severity: Warning Log: LOG_SLI verbose Action: This error could indicate a software driver, firmware or hardware issue. Report these errors to Broadcom Technical Support.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0315: Ring <ringno> issue: portCmdGet <local_getidx> is bigger than cmd ring <max_cmd_idx> The port command ring get index is greater than the size of the command ring.</p> <p>Data: None Severity: Error Log: Always Action: This error could indicate a software driver, firmware or hardware issue. Report these errors to Broadcom Technical Support.</p>
<p>0317: iotag <ulp_IoTag> is out of range: max iotag <max_iotag> wd0 <wd0> The IoTag in the completed IOCB is out of range.</p> <p>Data: None Severity: Error Log: Always Action: This error could indicate a software driver, firmware or hardware issue. Report these errors to Broadcom Technical Support.</p>
<p>0318: Failed to allocate IOTAG. last IOTAG is <last_allocated_iotag> The driver cannot allocate an IoTag. Display the last value used.</p> <p>Data: None Severity: Error Log: Always Action: This message indicates the adapter HBA I/O queue is full. Typically this happens when heavy I/O is running on a low-end (3-digit) adapter. Upgrade to a higher-end adapter.</p>
<p>0319: READ_SPARAM mbxStatus error <mbxStatus> hba state <hba_state> The driver issued a READ_SPARAM mailbox command to the HBA that failed.</p> <p>Data: None Severity: Error Log: Always Action: This error could indicate a firmware or hardware issue. Report these errors to Broadcom Technical Support.</p>
<p>0320: CLEAR_LA mbxStatus error <mbxStatus> hba state <hba_state> The driver issued a CLEAR_LA mailbox command to the HBA that failed.</p> <p>Data: None Severity: Error Log: Always Action: This error could indicate a firmware or hardware issue. Report these errors to Broadcom Technical Support.</p>
<p>0322: Ring <ringno> handler: unexpected completion IoTag <IoTag> The driver could not find a matching command for the completion received on the specified ring.</p> <p>Data: (1) ulpStatus, (2) ulpWord[4], (3) ulpCommand, (4) ulpContext Severity: Warning Log: LOG_SLI verbose Action: This error could indicate a software driver or firmware issue. If this issue persists, report these errors to Broadcom Technical Support.</p>
<p>0323: Unknown Mailbox command <mbxCommand> Cmpl A unknown mailbox command completed.</p> <p>Data: None Severity: Error Log: Always Action: This error could indicate a software driver, firmware, or hardware issue. Report these errors to Broadcom Technical Support.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0324: Config port initialization error, mbxCmd <mbxCommand> READ_NVPARM, mbxStatus <mbxStatus> READ_NVPARMS mailbox command failed during port configuration.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: Always</p> <p>Action: This error could indicate a software driver, firmware, or hardware issue. Report these errors to Broadcom Technical Support.</p>
<p>0328: Rsp Ring <ring number> error: IOCB Data:</p> <p>The firmware has returned an error for this IOCB.</p> <p>Data: (1) <iocb word[0]:iocb word[7]>, (2) <rsp word[0]:rsp[word[7]></p> <p>Severity: Warning</p> <p>Log: LOG_SLI</p> <p>Action: None required.</p>
<p>0330: IOCB wake NOT set</p> <p>The completion handler associated with the IOCB was never called.</p> <p>Data: (1) timeout (2) timeleft/jiffies</p> <p>Severity: Error</p> <p>Log: Always</p> <p>Action: This error could indicate a software driver, firmware, or hardware issue. If this issue persists, report the error to Broadcom Technical Support.</p>
<p>0334: Unknown IOCB command</p> <p>Received an unknown IOCB command completion.</p> <p>Data: (1) type (2) ulpCommand (3) ulpStatus (4) ulploTag (5) ulpContext)</p> <p>Severity: Error</p> <p>Log: Always</p> <p>Action: This error could indicate a software driver or firmware issue. If this issue persists, report these errors to Broadcom Technical Support.</p>
<p>0335: Unknown IOCB command</p> <p>Received an unknown IOCB command completion.</p> <p>Data: (1) ulpCommand (2) ulpStatus (3) ulploTag (4) ulpContext)</p> <p>Severity: Error</p> <p>Log: Always</p> <p>Action: This error could indicate a software driver or firmware issue. If this issue persists, report these errors to Broadcom Technical Support.</p>
<p>0336: Rsp Ring <ringno> error: IOCB</p> <p>An IOCB error has occurred on the specified ring.</p> <p>Data: (1) ulpWord[0], (2) ulpWord[1], (3) ulpWord[2], (4) ulpWord[3], (5) ulpWord[4], (6) ulpWord[5], (7) irsp+6, (8) irsp+7</p> <p>Severity: Warning</p> <p>Log: LOG_SLI verbose</p> <p>Action: If this issue persists, check the targets. If the targets are not the problem, report the error to Broadcom Technical Support.</p>
<p>0340: Adapter temperature is OK now</p> <p>Adapter temperature has reverted to normal range.</p> <p>Data: Temperature in Celsius</p> <p>Severity: Error</p> <p>Log: LOG_TEMP verbose</p> <p>Action: No action needed, informational.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0341: Ring <ringno> Cannot find buffer for an unsolicited iocb tag <un.ulpWord[3]> No more preallocated buffers are available to handle unsolicited buffers. Data: None Severity: Error Log: LOG_SLI Action: Ensure that this port is not being managed by multiple ports.</p>
<p>0342: Ring <ringno> Cannot find buffer for an unsolicited iocb tag <un.sli3.sli3Words> This is a multiple IOCB unsolicited command, and sufficient buffer space cannot be allocated for it. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>0343: Ring <ringno> Cannot find buffer for an unsolicited iocb tag <un.ulpWord[3]> No more preallocated buffers are available to handle unsolicited buffers. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>0344: Ring <ringno> Cannot find buffer for an unsolicited iocb tag <un.sli3.sli3Words[7]> No more preallocated buffers are available to handle unsolicited buffers. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>0345: Resetting board due to mailbox timeout iocb. tag 0x%x A mailbox command failed to complete. The driver is resetting the port. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: If the mailbox command fails again, set the lpfc_log_verbose to LOG_MBOX and retry.</p>
<p>0346: Ring <ring number> handler: unexpected ASYNC_STATUS evt_code <evt code> W0 <hex w0> W1 <hex w1> W2 <hex W2> W3 <hex W3> W4 <hex W4> W5 <hex W5> W6 <hex W6> W7 <hex W7> W8 <hex W8> W9 <hex w9> W10 <hex W10> W11<hex W11> The HBA received an asynchronous event that was not a temperature event. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>0347: Adapter is very hot, please take corrective action Adapter temperature is above normal range. Data: Temperature in Celsius Severity: Error Log: LOG_TEMP verbose Action: Shutdown and remove the HBA. Contact Broadcom Technical Support.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0348: NameServer login: node freed The enable mode failed to free up the NameServer login. Data: None Severity: Error Log: LOG_ELSI Action: None required.</p>
<p>0349: rc should be MBX_SUCCESS The next mailbox command on the mailbox queue has failed. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.</p>
<p>0350: rc should have been MBX_BUSY Attempting to unregister a default RPI from an interrupt context and the mailbox state is not busy. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.</p>
<p>0351: Config MSI mailbox command failed, mbxCmd <u.mb.mbxComm>, mbxStatus <u.mb.mbxStatus> The mailbox command sent to the firmware to configure the adapter to use MSI-X has failed. Data: None Severity: Warning Log: LOG_MBOX Action: Ensure that the hardware platform supports MSI-X.</p>
<p>0352: Config MSI mailbox command failed, mbxCmd <u.mb.mbxCommand>, mbxStatus <u.mb.mbxStatus> The mailbox command sent to the firmware to configure the HBA to use MSI-X has failed. Data: None Severity: Error Log: LOG_MBOX Action: Ensure that the hardware platform supports MSI-X.</p>
<p>0353: Active Mailbox cleared - mailbox timeout exiting The mailbox timeout handler has determined that the driver is in the process of completing this mailbox command. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.</p>
<p>0357: MSI-X interrupt with no EQE SLI-4 adapter interrupt on the slow path but there is no associated EQE. Data: None Severity: Warning Log: LOG_SLI Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0358: MSI-X interrupt with no EQE SLI-4 adapter interrupt on the fast path but there is no associated EQE. Data: None Severity: Warning Log: LOG_SLI Action: None required.</p>
<p>0359:Not a valid slow-path completion event: majorcode=x%x, minorcode=x%x\n, bf_get(lpfc_eqe_ma- jor_code, eqe), bf_get(lpfc_eqe_minor_code, eqe)); SLI-4: The EQE is invalid. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>0360:Unsupported EQ count. <entry_count> Cannot create an event queue of this size. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>0361:Unsupported CQ count. <entry_count> Cannot create an completion queue of this size. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>0362:Unsupported MQ count. <entry_count> Cannot create MQ of this size. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>0364:Invalid param: SLI-4: The post-SGL function was passed an invalid XRI. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>0365:Slow-path CQ identifier <CQID> does not exist: The Completion Queue ID passed in the event queue entry does not reference a valid completion queue. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0366: Not a valid fast-path completion event: majorcode=<major code hex>, minor-code=<minor code hex> The major or minor code in the Event Queue field is invalid.</p> <p>Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>0367: Fast-path completion queue does not exist The fast-path completion queue referenced by the CQID does not exist.</p> <p>Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>0368: Mis-matched fast-path completion queue identifier: eqcqid=%d, fcpcqid=%d The CQID in the event queue entry does not match the fcp_cqid that was passed into the routine.</p> <p>Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>0369: No entry from fast-path completion queue fcpcqid=<queue_id> No completions in the completion queue were referenced by fcp_cqid.</p> <p>Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>0370: Invalid completion queue type <type> The event queue entry is not for a mailbox or a work queue entry.</p> <p>Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>0371: No entry from the CQ: identifier <queue_id>, type <type> No completion queue event existed for this event queue entry.</p> <p>Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>0372: iotag <iotag> is out of range: max iotag (<sli.last_iotag>) The IOCB lookup cannot be performed because the iocb_tag is out of range.</p> <p>Data: None Severity: Error Log: LOG_SLI Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0373: FCP complete error: status=<status> hw_status=<hw status>, total_data_specified=<total data transferred>, parameter=<rsp word[4]>, word3=<wcqe word 3></p> <p>Logs the FCP failure. Status and parameter are equivalent to ulpStatus and ulpWord[4].</p> <p>Data: None</p> <p>Severity: Warning</p> <p>Log: LOG_SLI</p> <p>Action: None required.</p>
<p>0374: FCP complete with no corresponding cmdiocb: iotag <iocb iotag></p> <p>No IOCB was on the in-progress list that matched this iotag.</p> <p>Data: None</p> <p>Severity: Warning</p> <p>Log: LOG_SLI</p> <p>Action: None required.</p>
<p>0375: FCP cmdiocb not callback function iotag: <iocb iotag></p> <p>The IOCB found for this iotag does not have a completion handler set in it.</p> <p>Data: None</p> <p>Severity: Warning</p> <p>Log: LOG_SLI</p> <p>Action: None required.</p>
<p>0377: Error <rc> parsing vpd. Using defaults.</p> <p>Could not parse the VPD data, so the driver is using the default values.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: Always</p> <p>Action: None required.</p>
<p>0378: No support for fcpi mode.</p> <p>Could not configure the port to run in FCP initiator mode.</p> <p>Data: None</p> <p>Severity: Warning</p> <p>Log: LOG_MBOX, LOG_SLI</p> <p>Action: None required.</p>
<p>0379: Feature Mismatch Data: <req_ftr word2 hex> <req_ftr word3 hex> <cfg_enable_npiv> <max_vpi hex></p> <p>The features passed in to the driver as module parameters do not match what the firmware can do. Setting to default values.</p> <p>Data: None</p> <p>Severity: Warning</p> <p>Log: LOG_MBOX, LOG_SLI</p> <p>Action: None required.</p>
<p>0381: Error %d during queue setup.</p> <p>Could not set up all the queues that the driver requires to exchange I/Os with the HBA.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_MBOX, LOG_SLI</p> <p>Action: Reload the driver.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0382: READ_SPARAM command failed status <issue status>, mbxStatus <mailbox status> The READ_SPARAM mailbox command has failed during initialization. The HBA has been set to error state. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI: Action: Perform a dump with hbacmd and then try reloading the driver.</p>
<p>0383: Error <rc> during scsi sgl post operation The SGL entries could not be registered with the adapter. Data: None Severity: Warning Log: LOG_MBOX, LOG_SLI Action: Reset the adapter using hbacmd.</p>
<p>0384: There is pending active mailbox cmd The mailbox commands have overlapped. This command should have been added to the mailbox queue. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.</p>
<p>0385: rc should have been MBX_BUSY The completion handler for REG_LOGIN detected the IMMED_UNREG flag and tried to issue the UNREG_LOGIN command from an interrupt level. The mailbox status should still be busy. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.</p>
<p>0386: ELS complete with no corresponding cmdioch: iotag <iotag> The completion that the ISR is handling cannot find a tag associated with the IOTAG. Data: None Severity: Warning Log: LOG_SLI Action: None required.</p>
<p>0387:Failed to allocate an iocbq Failed to get an IOCBQ from the list of available IOCBQs. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>0388:Not a valid WCQE code: x<hex cqe_code> The event code is invalid. This event will be dropped. Data: None Severity: Error Log: LOG_SLI Action: Ensure that the adapter firmware is current.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0391:Error during rpi post operation The driver was trying to post pages to the firmware to keep target login information and encountered a failure. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: Unload and reload the driver.</p>
<p>0393:Error <rc> during rpi post operation The driver was trying to post pages to the firmware to keep target login information and encountered a failure. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: Unload and reload the driver.</p>
<p>0394: Failed to allocate CQ_EVENT entry The asynchronous event handler was unable to allocate an event queue entry to which to transfer the asynchronous event. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: This could be a V-LINK clear from the switch or a fatal error from the firmware. Perform a dump from the Emulex OneCommand Manager application.</p>
<p>0395: The mboxq allocation failed The asynchronous link event handler could not allocate a mailbox command to issue the READ_LA (read link attention) mailbox command. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>0396:The lpfc_dmabuf allocation failed The asynchronous link event handler could not allocate a mailbox command to issue the READ_LA mailbox command. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>0397:The mbuf allocation failed The asynchronous link event handler could not allocate DMA for the READ_LA mailbox command. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>0398 Invalid link fault code: <hex link_fault> The attempt to read the link attention register has returned an unknown value. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0399 Invalid link attention type: <hex link_type> The READ_LA mailbox command has returned an invalid link type. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>0401: Ignoring change to nodev_tmo because devloss_tmo is set Attempting to change the nodev timeout when the devloss has already been set. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>0402: Cannot find virtual addr for buffer tag on ring <ringno> A DMA buffer is unavailable for this unsolicited command. Data: (1) tag (2) next (3) prev (4) postbufq_cnt Severity: Error Log: LOG_INIT Action: None required.</p>
<p>0404: lpfc_devloss_tmo attribute cannot be set to <val>, allowed range is [<LPFC_MIN_DEVLOSS_TMO>, <LPFC_MAX_DEVLOSS_TMO>] Attempt to set the devloss timeout value is outside the allowed range. Data: None Severity: Error Log: LOG_INIT Action: Set the devloss timeout between the minimum and maximum devloss range.</p>
<p>0405: lpfc_link_speed attribute cannot be set to %d, allowed values are ["LPFC_LINK_SPEED_STRING"] Attempt to set the link speed value is outside the allowed range. Data: None Severity: Error Log: LOG_INIT Action: Set the link speed between 0 and the maximum.</p>
<p>0406: Adapter maximum temperature exceeded <temperature>, taking this port offline The driver has received an error for the HBA indicating that the maximum allowable temperature has been exceeded. Data: (1) work_hs (2) work_status[0] (3) work_status[1] Severity: Error Log: LOG_INIT Action: Ensure that the server fans are not blocked. Shut down the server if the airflow is restricted.</p>
<p>0407: Ignoring nodev_tmo module parameter because devloss_tmo is set. Both module parameters (nodev and devloss) were set so the driver is ignoring the nodev parameter. Data: None Severity: Error Log: LOG_INIT Action: Set only one of these parameters.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0410: Cannot find virtual addr for mapped buf on ring <ringno> The driver cannot find the specified buffer in its mapping table. Therefore, it cannot find the virtual address needed to access the data. Data: (1) phys (2) next (3) prev (4) postbufq_cnt Severity: Error Log: Always Action: This error could indicate a software driver or firmware issue. If this issue persists, report these errors to Broadcom Technical Support.</p>
<p>0421: MSI-X slow-path request_irq failed <rc> The kernel API to request an interrupt request (IRQ) has failed. Data: None Severity: Warning Log: LOG_INIT Action: Use module parameter lpfc_use_msi=0 (INTx).</p>
<p>0422: lpfc_restrict_login attribute cannot be set to <val>, allowed range is [0, 1] Attempt to set the restrict login parameter to something other than on or off. Data: None Severity: Error Log: LOG_INIT Action: Use 0 (Off) or 1 (On)</p>
<p>0423: lpfc_"#attr" attribute cannot be set to %d, allowed range is ["#minval", "#maxval"] This compile time macro is used by several module parameters during initialization. Each module parameter has its own minimum and maximum values that are displayed. Data: None Severity: Error Log: LOG_INIT Action: Set the module parameter between the minimum and maximum values.</p>
<p>0424:lpfc_"#attr" attribute cannot be set to %d, allowed range is ["#minval", "#maxval"] This is a compile time macro that is used by several module parameters to set the value. Data: None Severity: Error Log: LOG_INIT Action: Set the module parameter between the minimum and maximum values.</p>
<p>0425:lpfc_restrict_login attribute cannot be set to %d, allowed range is [0, 1] The module parameter lpfc_restrict_login can only be set to 0 (off) or 1 (on). Data: None Severity: Error Log: LOG_INIT Action: Set lpfc_restrict_login = [0,1].</p>
<p>0426: Failed to enable interrupt The driver failed to start the interrupt. Data: None Severity: Error Log: LOG_INIT Action: Unload and reload the driver.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0427: Cannot re-enable interrupt after slot reset The driver was unable to enable the interrupt after an HBA reset. Data: None Severity: Error Log: LOG_INIT Action: Unload and reload the driver.</p>
<p>0429: MSI-X fast-path request_irq failed (<rc>) The driver received an error for the request_irq_call. Data: None Severity: Warning Log: LOG_INIT Action: Unload and reload the driver.</p>
<p>0430: PM resume Failed to enable interrupt The driver power management resume function could not enable the interrupt. Data: None Severity: Error Log: LOG_INIT Action: Perform another PM suspend and resume or HBA reset.</p>
<p>0431: Failed to enable interrupt. The driver failed to start the interrupt. Data: None Severity: Error Log: LOG_INIT Action: Unload and reload the driver.</p>
<p>0433: Wakeup on signal: rc=<rc> A signal other than the LPFC_DATA_READY was received on the worker thread. Data: None Severity: Error Log: LOG_ELS Action: Unload and reload the driver.</p>
<p>0434: PM resume failed to start worker thread: error=<error> The driver's power management resume function could not start the worker thread. Data: None Severity: Error Log: LOG_INIT Action: Unload and reload the driver.</p>
<p>0435: Adapter failed to get Option ROM version status <rc> The driver could not read the HBA option ROM. Data: None Severity: Error Log: LOG_INIT Action: Reset the HBA. Ensure the adapter firmware is current.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0436: Adapter failed to init, timeout, status reg <status> The adapter failed during power-up diagnostics after it was reset. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.</p>
<p>0437: Adapter failed to init, chipset, status reg <status> The adapter failed during power-up diagnostics after it was reset. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.</p>
<p>0438: Adapter failed to init, chipset, status reg <status> The adapter failed during power-up diagnostics after it was reset. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.</p>
<p>0439: Adapter failed to init, mbxCmd <mbxCommand> READ_REV, mbxStatus <mbxStatus> Adapter initialization failed when issuing a READ_REV mailbox command. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.</p>
<p>0440: Adapter failed to init, READ_REV has missing revision information A firmware revision initialization error was detected. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. Install the latest firmware revision. If this issue persists, report the error to Broadcom Technical Support.</p>
<p>0442: Adapter failed to init, mbxCmd <mbxCommand> CONFIG_PORT, mbxStatus <mbxStatus> Adapter initialization failed when issuing a CONFIG_PORT mailbox command. Data: (1) hbainit Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.</p>
<p>0446: Adapter failed to init, mbxCmd <mbxCommand> CFG_RING, mbxStatus <mbxStatus>, ring <num> Adapter initialization failed when issuing a CFG_RING mailbox command. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0447: Adapter failed init, mbxCmd <mbxCommand> CONFIG_LINK mbxStatus <mbxStatus> Adapter initialization failed when issuing a CONFIG_LINK mailbox command. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.</p>
<p>0448: Adapter failed to init, mbxCmd <mbxCommand> READ_SPARM, mbxStatus <mbxStatus> Adapter initialization failed when issuing a READ_SPARM mailbox command. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.</p>
<p>0449: lpfc_%attr attribute cannot be initialized to %d, allowed range is [%min, %max] The sysfs attribute value written exceeds attribute range. Data: (1) attribute name (2) value written (3) minimum value (3) maximum value Severity: Error Log: Always Action: Write a value within the supported range.</p>
<p>0450: lpfc_%attr attribute cannot be set to %d, allowed range is [%min, %max] The sysfs attribute value written exceeds attribute range. Data: (1) attribute name (2) value written (3) minimum value (3) maximum value Severity: Error Log: Always Action: Write a value within the supported range.</p>
<p>0451: Enable interrupt handler failed The driver attempted to register the HBA interrupt service routine with the host operating system, but failed. Data: None Severity : Error Log: Always Action: This error could indicate a hardware or driver issue. If this issue persists, report the error to Broadcom Technical Support.</p>
<p>0453: Adapter failed to init, mbxCmd <mbxCommand> READ_CONFIG, mbxStatus <mbxStatus> Adapter initialization failed when issuing a READ_CONFIG mailbox command. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.</p>
<p>0454: Adapter failed to init, mbxCmd <mbxCommand> INIT_LINK, mbxStatus <mbxStatus> Adapter initialization failed when issuing an INIT_LINK mailbox command. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0456: Adapter failed to issue ASYNCEVT_ENABLE mbox status x%x The mailbox command to enable an asynchronous event notification failed. Data: None Severity: Error Log: LOG_INIT Action: Ensure the adapter firmware is current. Reload the driver.</p>
<p>0457: Adapter Hardware Error The driver received an interrupt indicating a possible hardware issue. Data: (1) status (2) status1 (3) status2 Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.</p>
<p>0458: Bring adapter online The FC driver has received a request to bring the adapter online. This issue might occur when running lputil. Data: None Severity: Warning Log: LOG_INIT verbose Action: None required.</p>
<p>0459: Adapter heartbeat failure, taking this port offline. The Heartbeat mailbox command failed. Data: None Severity: Error Log: LOG_INIT Action: Ensure that the adapter firmware is current. Reload the driver.</p>
<p>0460: Bring adapter offline The FC driver has received a request to bring the adapter offline. This issue might occur when running lputil. Data: None Severity: Warning Log: LOG_INIT verbose Action: None required.</p>
<p>0466: Outstanding I/O when bringing Adapter offline The I/O is still pending while attempting to stop the driver. Data: None Severity: Warning Log: LOG_INIT Action: None required.</p>
<p>0467: lpfc_topology attribute cannot be set to %d, allowed range is [0, 6], phba->brd_no, val. The lpfc_topology module parameter is invalid. Data: None Severity: Error Log: LOG_INIT Action: Use a value in the valid range.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0468: lpfc_restrict_login must be 0 for Physical ports. "vport->cfg_restrict_login = 0; Cannot restrict the login for the physical port. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>0469: lpfc_link_speed attribute cannot be set to %d, allowed range is [0, 8] The link speed module parameter is invalid. Data: None Severity: Error Log: LOG_INIT Action: Use a link speed parameter in the valid range.</p>
<p>0472: PCI channel I/O permanent failure The PCI bus has detected an error. Data: None Severity: Error Log: LOG_INIT Action: Issue an HBA reset.</p>
<p>0474: Unable to allocate memory for issuing MBOX_CONFIG_MSI command Mailbox memory pool allocation error. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>0475: Not configured for supporting MSI-X cfg_use_msi: 0x%x The lpfc_use_msi module parameter should have been set to 2. Data: None Severity: Error Log: LOG_INIT Action: Set module parameter lpfc_use_msi= 2.</p>
<p>0476: HBA not supporting SLI-3 or later SLI Revision: <sli_rev> The HBA does not support SLI-3 or SLI-4. Data: None Severity: Error Log: LOG_INIT Action: This HBA does not support msi. Set lpfc_use_msi = 0.</p>
<p>0478: MSI request_irq failed (<rc>). The request_irq kernel API has failed. Data: None Severity: Warning Log: LOG_INIT Action: Set lpfc_use_msi = 0.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0479: Deferred Adapter Hardware Error An adapter hardware error was sent to the driver. Data: (1) work_hs, (2) work_status[0], (3) work_status[1] Severity: Error Log: LOG_INIT Action: Perform a dump using hbacmd.</p>
<p>0483:Invalid link-attention link speed: x%x", bf_get(lpfc_acqe_link_speed, acqe_link). The link speed reported in the link attention interrupt is invalid. Data: None Severity: Error Log: LOG_INIT Action: Check the switch configuration.</p>
<p>0485: MSI-X slow-path request_irq failed (<rc>). The request_irq kernel API has failed. Data: None Severity: Warning Log: LOG_INIT Action: Set module parameter lpfc_use_msi = 0.</p>
<p>0486: MSI-X fast-path (<index>) request_irq failed (<rc>). The request_irq kernel API has failed. Data: None Severity: Warning Log: LOG_INIT Action: Set module parameter lpfc_use_msi = 0.</p>
<p>0490: MSI request_irq failed (<rc>). The request_irq kernel API has failed. Data: None Severity: Warning Log: LOG_INIT Action: Set module parameter lpfc_use_msi = 0.</p>
<p>0492: Unable to allocate memory for issuing SLI_CONFIG_SPECIAL mailbox command Mailbox memory pool allocation error. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>0493: SLI_CONFIG_SPECIAL mailbox failed with status<rc> Mailbox command failed. Data: None Severity: Error Log: LOG_INIT Action: Ensure that the adapter firmware is current. Unload and reload the driver.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0494: Unable to allocate memory for issuing "SLI_FUNCTION_RESET mailbox command" Mailbox memory pool allocation error. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>0495: SLI_FUNCTION_RESET mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>. Mailbox command failed. Data: None Severity: Error Log: LOG_INIT Action: Reset the HBA.</p>
<p>0496: Failed allocate slow-path EQ The event queue for the slow path was not allocated. Data: None Severity: Error Log: LOG_INIT Action: Unload and reload the driver.</p>
<p>0497: Failed allocate fast-path EQ The event queue for the fast path was not allocated. Data: None Severity: Error Log: LOG_INIT Action: Unload and reload the driver.</p>
<p>0499: Failed allocate fast-path FCP CQ (<fcp_cqidx>). The completion queue event for the fast path could not be allocated. Data: None Severity: Error Log: LOG_INIT Action: Unload and reload the driver.</p>
<p>0500: Failed allocate slow-path mailbox CQ Failed to allocate slow-path mailbox CQ. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>0501: Failed allocate slow-path ELS CQ Failed to allocate slow-path ELS CQ. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0503: Failed allocate fast-path FCP Failed to allocate fast-path FCP. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>0504: Failed allocate slow-path ELS WQ Failed to allocate slow-path ELS WQ. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>0505: Failed allocate slow-path ELS MQ Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>0506: Failed allocate receive HRQ\n Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>0507: Failed allocate receive DRQ Failed to allocate receive DRQ. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>0520: Slow-path EQ not allocated The slow-path EQ not allocated. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>0522: Fast-path EQ <fcp_eqidx> not allocated The fast-path EQ is not allocated. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>0523: Failed setup of fast-path EQ <fcp_eqidx>, rc = <rc> The fast-path EQ setup failed. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0526: Fast-path FCP CQ <fcp_cqid> not allocated The fast-path FCP is not allocated. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>0527: Failed setup of fast-path FCP CQ <fcp_cqid>, rc = <rc> The fast-path FCP CQ setup failed. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>0528: Mailbox CQ not allocated The mailbox CQ is not allocated. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>0530: ELS CQ not allocated The ELS CQ is not allocated. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>0534: Fast-path FCP WQ <fcp_wqid> not allocated The fast-path FCP WQ is not allocated. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>0535: Failed setup of fast-path FCP WQ <fcp_wqid>, rc = <rc> The fast-path FCP WQ setup failed. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>0536: Slow-path ELS WQ not allocated The slow-path ELS WQ is not allocated. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0538: Slow-path MQ not allocated The slow-path MQ is not allocated. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>0540: Receive Queue not allocated The Receive Queue is not allocated. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>0542: lpfc_create_static_vport failed to allocate mailbox memory Failed to allocate mailbox memory for VPort creation. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>0543: lpfc_create_static_vport failed to allocate vport_info\n")) Failed to allocate vport_info. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>0544: lpfc_create_static_vport failed to issue dump mailbox command ret <rc> status <mbxStatus> Failed to issue a dump mailbox command for static VPort creation. Data: None Severity: Warning Log: LOG_INIT Action: None required.</p>
<p>0545: lpfc_create_static_vport bad information header 0x%x 0x%x\n"\, le32_to_cpu(vport_info->signature), le32_to_cpu(vport_info->rev) & VPORT_INFO_REV_MASK); Invalid information header; the signature or revision is invalid. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>0546: lpfc_create_static_vport failed to create vport Failed to create a VPort. Data: None Severity: Warning Log: LOG_INIT Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0582: Error <rc> during sgl post operation The SGL post operation failed. Data: None Severity: Error Log: LOG_MBOX, LOG_IP verbose Action: None required.</p>
<p>0602: Failed to allocate CQ_EVENT entry Failed to allocate a CQ_EVENT entry. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>0603: Invalid work queue CQE subtype (x%x)\n", cq-<subtype> Invalid work queue CQE. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>0700: Bus Reset on target <i> failed The bus reset for the specified target failed. Data: None Severity: Error Log: LOG_FCP Action: None required.</p>
<p>0704: At limitation of <total> preallocated command buffers The maximum number of command buffers have already been allocated. Data: None Severity: Warning Log: LOG_FCP verbose Action: None required.</p>
<p>0705: Allocation request of <num> command buffers will exceed max of <hba_queue_depth>. Reducing allocation request to <size> The number of command buffers requested will exceed the maximum so a smaller quantity will be allocated. Data: None Severity: Warning Log: LOG_FCP verbose Action: None required.</p>
<p>0708: Allocation request of <num_to_alloc> command buffers did not succeed. Allocated <num_allocated> buffers. The allocation request for the specified command buffers did not succeed. However, the specified number of buffers has been allocated. Data: None Severity: Warning Log: LOG_FCP Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0711: detected queue full - lun queue depth adjusted to%d The driver detected a queue full status on a SCSI command response. New LUN queue depth is reported. Data: (1) New LUN queue depth Severity: Warning Log: LOG_FCP verbose Action: This might indicate an oversubscribed target array. Check your SAN configuration and I/O workload.</p>
<p>0713: SCSI layer issued Device Reset (%d, %d) A device reset was issued. Data: None Severity: Error Log: LOG_FCP Action: None required.</p>
<p>0714: SCSI layer issued bus reset The SCSI layer is requesting the driver to abort all I/Os to all targets on this HBA. Data: (1) ret Severity: Error Log: Always Action: Check the state of the targets in question.</p>
<p>0720: FCP command <cmnd[0]> residual overrun error A residual overrun error has occurred while processing the specified FCP command. Data: (1) request_bufflen (2) resid Severity: Warning Log: LOG_FCP verbose Action: If this issue persists, check the targets for errors.</p>
<p>0721: Device Reset rport failure: rdata <rdata> The reset of the R_Port failed. Data: None Severity: Error Log: LOG_FCP Action: None required.</p>
<p>0722: Target Reset rport failure: rdata <rdata> The reset of the target failed. Data: None Severity: Error Log: LOG_FCP Action: None required.</p>
<p>0723: SCSI layer issued Target Reset (%d, %d) The SCSI layer issued a target reset. Data: None Severity: Error Log: LOG_FCP Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0724: I/O flush failure for context <"LUN","TGT","HOST","Unknown">: cnt <cnt> The I/O flush to the LUN, target, or host has failed. Data: None Severity: Error Log: LOG_FCP Action: None required.</p>
<p>0727: TME <cmd> to TGT <TGT#> LUN <LUN#> failed (<ulpStatus>, <ulpWord[4]>) The task management command failed. Data: None Severity: Error Log: LOG_FCP Action: None required.</p>
<p>0729: FCP cmd <cmd> failed <target>/<lun> status: <status> result: <result> The specified device failed an FCP command. Data: (1) ulpContext (2) iotag Severity: Warning Log: LOG_FCP verbose Action: Check the state of the target in question.</p>
<p>0730: FCP command failed: RSP The FCP command failed with a response error. Data: (1) resp_info (2) scsi_status (3) ResId (4) SnsLen (5) RspLen (6)rsplInfo3 Severity: Warning Log: LOG_FCP verbose Action: Check the state of the target in question.</p>
<p>0734: FCP read check error The issued FCP command returned a read check error. Data: (1) fcpDI (2) rspResId (3) fcp_i_parm (4) cmd[0] Severity: Warning Log: LOG_FCP verbose Action: Check the state of the target in question.</p>
<p>0735: FCP Read Check Error and Underrun Data HBA reported under run from storage array. Data: (1) vpi (2) fcpDI (3) res_id (4) fcp_i_parm Severity: Warning Log: LOG_FCP_ERROR verbose Action: No action needed, informational.</p>
<p>0748: Abort handler timed out waiting for abort to complete:ret <status> D <target id> LUN <lun id> The abort handler timed out waiting for abort to complete. Data: None Severity: Error Log: Always Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>0749: SCSI layer issued abort device The SCSI layer aborted a device. Data: (1) ret, (2) id, (3) lun, (4) snum Severity: Warning Log: LOG_FCP verbose Action: None required.</p>
<p>0915 Register VPI failed:<mbxStatus> Could not register the VPI. Data: None Severity: Error Log: LOG_MBOX Action: None required.</p>
<p>1019: Request tranid <tran_id> timed out A transaction with storage array could not complete due to timeout. Data: (1) tran_id Severity: Warning Log: LOG_SECURITY verbose Action: Software driver warning. If this issue persists, report these errors to Broadcom Technical Support.</p>
<p>1021: ERROR: attempted to queue security work, when no workqueue created Driver encountered missing queue required for processing security information. Data: None Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report these errors to Broadcom Technical Support.</p>
<p>1022: Security request does not exist A security request operation failed because no match was found for such request. Data: None Severity: Warning Log: LOG_SECURITY Action: Software driver warning. If this issue persists, report these errors to Broadcom Technical Support.</p>
<p>1023: Warning - data may have been truncated. Data: <data> reqdl: <data_len> mesdl:<data_len> A security message exchange operation failed because the response was missing or unreliable. Data: None Severity: Warning Log: LOG_SECURITY Action: Software driver warning. If this issue persists, report these errors to Broadcom Technical Support.</p>
<p>1028: Start Authentication: No buffers The authentication failed because some memory resources were not allocated. Data: None Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>1029: Reauthentication Failure The driver encountered errors and there was a failure to re-authenticate. Data: None Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.</p>
<p>1031: Start Authentication: Get config failed The authentication failed due to some error during port configuration. Data: None Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.</p>
<p>1032: Start Authentication: get config timed out The node authentication was aborted because waiting for port configuration to complete, timed out. Data: None Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.</p>
<p>1033: Received auth_negotiate from Nport: < nlp_DID> Unsolicited authentication negotiation message received from a port. Data: (1) nlp_DID Severity: Warning Log: LOG_SECURITY Action: No action, this message is informational.</p>
<p>1034: Not Expecting Challenge - Rejecting Challenge Unsolicited authentication challenge received from a port was rejected. Data: None Severity: Warning Log: LOG_SECURITY Action: Software driver warning. If this issue persists, report errors to Broadcom Technical Support.</p>
<p>elx_mag1036: Authentication transaction reject - re-auth request reason <reason> exp <explanation> An authentication was rejected and requested again due to reason as displayed with explanation. Data: (1) reason (2) explanation. Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.</p>
<p>1037: Authentication transaction reject - restarting authentication, reason <reason> exp <explanation> An authentication process was rejected then restarted and an authentication requested was again due to reason as displayed with the explanation. Data: (1) reason (2) explanation. Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>1039: Not Expecting Reply - rejecting. State <state> An unanticipated reply was received during authentication and was subsequently rejected. Data: (1) auth_state. Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.</p>
<p>1040: Bad Reply trans_id- rejecting. Trans_id < trans_id > Expecting: < trans_id> Unexpected transaction id was received during authentication and was subsequently rejected. Data: (1) auth_state Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.</p>
<p>1043: Authentication LS_RJT The authentication request was rejected. Data: None Severity: Error Log: LOG_ELS Action: None required.</p>
<p>1045: Issue AUTH_NEG failed Status:%x The authentication negotiation failed. Data: None Severity: Error Log: LOG_ELS Action: None required.</p>
<p>1048: Issue AUTH_REJECT failed Could not issue the reject for the authentication request. Data: None Severity: Error Log: LOG_ELS Action: None required.</p>
<p>elx_msg1050: Authentication mode is disabled, but is required by the fabric Discovery failed because the switch fabric required authentication, but either authentication was not configured or the authentication mode for this port pair is disabled. Data: None Severity: Error Log: LOG_SECURITY Action: Configure the driver to authenticate with the switch or disable authentication on the switch to this port.</p>
<p>1053: Start Authentication: Security service offline The authentication failed because security service was unavailable. Data: None Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>1055: Authentication parameter is disabled, but is required by the fabric FLOGI failed because the fabric has indicated that authentication is required, but authentication has not yet been configured or enabled on the HBA. Data: None Severity: Error Log: LOG_SECURITY Action: Configure authentication on this HBA.</p>
<p>1057: Authentication transaction reject. reason <reason> exp <explanation> An authentication was rejected and requested again due to the reason as displayed with the explanation. Data: (1) reason (2) explanation. Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.</p>
<p>1058: Waiting for authentication service A delay occurred when the authentication service was not initially available as expected. Data: None Severity: Warning Log: LOG_SECURITY Action: Software driver warning. If this issue persists, report these errors to Broadcom Technical Support.</p>
<p>1059: Authentication became available The authentication service came online, but it was not initially available as expected. Data: None Severity: Warning Log: LOG_SECURITY Action: Software driver warning. If this issue persists, report these errors to Broadcom Technical Support.</p>
<p>1201: Failed to allocate dfc_host Failed to allocate memory for the dfc_host_struct. Data: None Severity: Error Log: LOG_ELS Action: None required.</p>
<p>1210: Invalid cmd size: cmd <cmd> cmdsz <cmdsized> rspsz <rspsized> The management command for LPFC 2100 has failed. Data: None Severity: Error Log: LOG_LIBDFC Action: None required.</p>
<p>1211: genreq alloc failed\n"); Resource allocation failure. Data: (1) return code Log: LOG_LIBDFC Severity: Error Action: Kernel memory resources are too low.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>1213: FCoE cmd overflow: off <#> + cnt <#> > cmdsz <#> Application has tried to read more data than originally requested. Data: (1) response offset (2) size (3) cmd size Log: LOG_LIBDFC Severity: Error Action: The application might have sent a invalid command.</p>
<p>1214: Cannot issue FCoE cmd SLI not active: <#> rc= -EACCESS The SLI layer has not been initialized. Data: (1) offset Log: LOG_LIBDFC Severity: Error Action: Restart the HBA.</p>
<p>1215: Cannot issue FCoE cmd: not ready or not in maint mode Either the external link is unplugged, link down, and the FCoE is not in maintenance mode. Data: (1) current offset (2) return code. Log: LOG_LIBDFC Severity: Error Action: Plug external cable in or set FCoE in maintenance mode.</p>
<p>1216: FCoE IOCB failed: off <#> rc <#> FCoE command generated by the application has failed. Data: (1) offset (2) return code. Log: LOG_LIBDFC Severity: Error Action: Application should retry the command.</p>
<p>1223: menlo_write: couldn't alloc genreq Resource allocation failure. Data: None Log: LOG_LIBDFC Severity: Error Action: Kernel memory resources too low.</p>
<p>1224: FCoE iocb failed off <#> rc=<#>", FCoE command failed in SLI. Data: (1) offset (2) return code Log: LOG_LIBDFC Severity: Informational. Action: Retry the command. If it fails again, reset HBA when convenient.</p>
<p>1227: FCoE IOCB TMO: handler set for <context3> The management command for the LPFC 2100 has timed out. Data: None Severity: Warning Log: LOG_LIBDFC Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>1228: FCoE IOCB TMO: handler set for <context3> A management IOCB for the LPFC 2100 has timed out. Data: None Severity: Warning Log: LOG_LIBDFC Action: None required.</p>
<p>1229: Waiting for menlo mnt Waiting for the LPFC 2100 to enter maintenance method. Data: None Severity: Warning Log: LOG_LIBDFC Action: None required.</p>
<p>1230: Could not find buffer for FCoE cmd:off <#> indmp <addr> off <#> Could not find resources associated with this FCoE command. Data: (1) current offset (2) buffer desc pointer (3) size Severity: Error Log: LOG_LIBDFC Action: Try reloading the driver when convenient.</p>
<p>1231: bad bpl: An invalid buffer list was detected upon completion. Data: None Severity: Error Log: LOG_LIBDFC Action: None required.</p>
<p>1235: Could not find buffer for FCoE cmd: off:<#> poff:<#> cnt:<#> mlascnt:<#> addl:<x> addh:<x> mdsz:<#> FCoE command failed because it could not find the resource. Data: (1) current offset (2)previous offset (3) count (4) last count (5) address low (6) address high Severity: Error Log: LOG_LIBDFC Action: No action needed, informational.</p>
<p>1238: FCoE IOCB failed: off <#> rc=<#> The command generated by the driver to check the FCoE has failed. Data: (1) offset (2) return code Log: LOG_LIBDFC Severity: Error Action: Make sure link is up or the adapter has set menlo in maintenance mode.</p>
<p>1240: Unable to allocate command buffer memory Could not allocate memory for the command buffer. Data: None Severity: Error Log: LOG_LINK_EVENT Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>1243: Menlo command error. code=%d.\n", mlorsp->code The Menlo maintenance command failed. Data: None Severity: Error Log: LOG_LINK_EVENT Action: None required.</p>
<p>1244: Unable to allocate response buffer memory. Could not allocate memory for the management command response. Data: None Severity: Error Log: LOG_LINK_EVENT Action: None required.</p>
<p>1246: FCoE chip is running golden firmware. Update FCoE chip firmware immediately <fw_type> The FCoE is running the golden firmware. Data: (1) firmware-type Severity: Error Log: LOG_LINK_EVENT Action: Try resetting the FCoE to operational mode and disable maintenance mode.</p>
<p>1247: FCoE chip is running diagnostic firmware. Operational use suspended. <fw_type> The FCoE is running a diagnostic. Data: (1) firmware-type Severity: Error Log: LOG_LINK_EVENT Action: Try resetting the FCoE to operational mode.</p>
<p>1248: FCoE chip is running unknown firmware. <fw_type> The FCoE is running an unknown firmware version. Data: (1) firmware-type Severity: Error Log: LOG_LINK_EVENT Action: Try resetting the FCoE to operational mode. Try loading the latest FCoE firmware.</p>
<p>1249: Invalid FRU data found on adapter. Return adapter to Emulex for repair. The FRU data on the FCoE chip is invalid. Data: (1) firmware-type Severity: Error Log: LOG_LINK_EVENT Action: Try resetting the FCoE to operational mode. Try loading the latest FCoE firmware or send the HBA back to Broadcom for repair.</p>
<p>1250: Menlo command error. code=<#> The IOCB driver sent to check FCoE state has bad header size. Data: (1) return code Log: LOG_LINK_EVENT Severity: Error Action: Try resetting the FCoE to operational mode.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>1251: Menlo command error. code=<#> The IOCB driver sent to check FCoE state has failed, no resources. Data: (1) return code Log: LOG_LINK_EVENT Severity: Error Action: Try resetting the FCoE to operational mode.</p>
<p>1252: Menlo command error. code=<#> The IOCB driver sent to check FCoE state has failed. Data: (1) return code Log: LOG_LINK_EVENT Severity: Error Action: Try resetting the FCoE to operational mode.</p>
<p>1257: lpfc_menlo_issue_iocb: handler set for <context3>. Data: None Log: LOG_LIBDFC Severity: Warning Action: None required.</p>
<p>1259: mbox: Issued mailbox cmd <u.mb.mbxCommand> while in stopped state. Only the dump mailbox command and reset adapter mailbox command are allowed when in the stopped state. Data: None Severity: Warning Log: LOG_MBOX Action: None required.</p>
<p>1262: Failed to allocate dfc_host Could not allocate memory the dfc_host_struct. Data: None Log: LOG_LIBDFC Severity: Error Action: None required.</p>
<p>1268: Find ndlp returned NULL for oxid:x%x SID:x%x, oxid, sid.(int)off, rc. Could not find the node for this DID. Data: None Severity: Warning Log: LOG_ELS Action: None required.</p>
<p>1302: Invalid speed for this board: Reset link speed to auto: <cfg_link_speed> The driver is reinitializing the link speed to auto-detect. Data: None Severity: Warning Log: LOG_LINK_EVENT verbose Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>1303: Link Up Event <eventTag> received A link up event was received. It is possible for multiple link events to be received together. Data: (1) fc_eventTag (2) granted_AL_PA (3) UlnkSpeed (4) alpa_map[0] Detail: If link events received, log (1) last event number received, (2) ALPA granted, (3) Link speed, (4) number of entries in the loop init LILP ALPA map. An ALPA map message is also recorded if LINK_EVENT verbose mode is set. Each ALPA map message contains 16 AL_PAs. Severity: Error Log: Always Action: If numerous link events are occurring, check the physical connections to the FC network.</p>
<p>1304: Link Up Event ALPA map A link up event was received. Data: (1) wd1, (2) wd2, (3) wd3, (4) wd4 Severity: Warning Log: LOG_LINK_EVENT verbose Action: If numerous link events are occurring, check the physical connections to the FC network.</p>
<p>1305: Link Down Event <eventTag> received A link down event was received. Data: (1) fc_eventTag (2) hba_state (3) fc_flag Severity: Error Log: Always Action: If numerous link events are occurring, check the physical connections to the FC network.</p>
<p>1306: Link Up Event in loop back mode x%x received Data: x%x x%x x%x x%x Link up notification; configured for loopback. Data: (1) fc_eventTag (2) granted_AL_PA (3) UlnkSpeed (4) alpa_map[0] Severity: Error Log: LOG_LINK_EVENT Action: None required.</p>
<p>1308: Menlo Maint Mode Link up Event x%x rcvd Data: x%x x%x x%x Link down notification; configured for loopback. Data: (1) fc_eventTag (2) port_state (3) vport fc_flag Severity: Error Log: LOG_LINK_EVENT Action: None required.</p>
<p>1309: Link Up Event npiv not supported in loop topology NPIV is not supported in loop topology. Data: None Severity: Error Log: LOG_LINK_EVENT Action: None required.</p>
<p>1310: Menlo Maint Mode Link up Event <eventTag> rcvd The link is up in maintenance mode; only management commands are allowed. Data: (1) fc_eventTag (2) port_state (3) vport fc_flag Severity: Error Log: LOG_LINK_EVENT Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>1312: Link Down Event <eventTag> received Maintenance mode link up notification received without entering link down. Data: (1) fc_eventTag (2) port_state (3) vport fc_flag Severity: Error Log: LOG_LINK_EVENT Action: None required.</p>
<p>1400: Failed to initialize sgl list. Failed to initialize SGL list during initialization. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>1401: Failed to enable pci device. Failed to enable PCI device during initialization. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>1402: Failed to set up pci memory space. PCI initialization failed. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>1403: Failed to set up driver resource. Driver resource initialization failed. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>1404: Failed to set up driver resource. Driver resource initialization failed. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>1405: Failed to initialize iocb list. Driver resource initialization failed. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>1406: Failed to set up driver resource. Initialization failed to set up driver resource. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>1407: Failed to create scsi host. Initialization failed to create SCSI host. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>1408: Failure HBA POST Status: sta_reg=0x%x, "perr=x%x, sfi=x%x, nip=x%x, ipc=x%x, xrom=x%x, "dl=x%x, pstatus=x%x\n", sta_reg.word0, bf_get(lpfc_hst_state_perr, &sta_reg), The HBA POST has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>1409: Failed to enable pci device. Failed to enable PCI device during initialization. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>1410: Failed to set up pci memory space. Initialization failed to set up PCI memory space. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>1411: Failed to set up driver resource. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>1412: Failed to set up driver resource. Initialization failed to set up driver resource. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>1413: Failed to initialize iocb list. Initialization failed to initialize the IOCB list. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>1414: Failed to set up driver resource. Initialization failed to set up driver resource. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>1415: Failed to create scsi host. Initialization failed to create SCSI host. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>1416: Failed to allocate sysfs attr Initialization failed to sysfs attribute. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>1418: Invalid HBA PCI device group: <dev_grp> Invalid HBA PCI device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>1419: Invalid HBA PCI device group: <dev_grp> Invalid HBA PCI device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>1420: Invalid HBA PCI device group:<dev_grp> Invalid HBA PCI device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>1421: Failed to set up hba Initialization failed to set up the HBA. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>1422: HBA Unrecoverable error: uerr_lo_reg=<ue lo>, uerr_hi_reg=<ue hi>, online0_reg=<Online0>, online1_reg=<Online1> The HBA has notified the driver that it has encountered an unrecoverable error. Data: None Severity: Error Log: LOG_INIT Action: Perform a dump from the Emulex OneCommand Manager application. Then, unload and reload the driver.</p>
<p>1423: HBA Unrecoverable error: uerr_lo_reg=<ue lo>, uerr_hi_reg=<ue hi>, online0_reg=<Online0>, online1_reg=<Online1> The HBA has notified the driver that it has encountered an unrecoverable error. Data: None Severity: Error Log: LOG_INIT Action: Perform a dump from the Emulex OneCommand Manager application. Then, unload and reload the driver.</p>
<p>1424: Invalid PCI device group: <pci_dev_grp> Invalid HBA PCI device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>1425: Invalid PCI device group: <pci_dev_grp> Invalid HBA PCI device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>1426: Invalid PCI device group: <pci_dev_grp> Invalid HBA PCI device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>1427: Invalid PCI device group: <pci_dev_grp> Invalid HBA PCI device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>1428: Invalid PCI device group: <pci_dev_grp> Invalid HBA PCI device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>1429: Invalid PCI device group: <pci_dev_grp> Invalid HBA PCI device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>1430: Failed to initialize sql list Failed to initialize SQL list. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>1431: Invalid HBA PCI device group: <pci_dev_grp> Invalid HBA PCI device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>1432: Failed to initialize rpi headers. Failed to initialize RPI headers. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>1476: Failed to allocate sysfs attr Failed to allocate sysfs attributes. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>1477: Failed to set up hba Failed to set up the HBA. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>1603: Loopback test did not receive expected data length. actual length <len>expected length <full_size>.</p> <p>The loopback test did not receive the same amount of data that it transmitted.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_LIBDFC</p> <p>Action: None required.</p>
<p>1800: Could not issue unreg_vpi</p> <p>Driver attempt to unregister VPI failed.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_VPORT verbose</p> <p>Action: Software driver error. If this issue persists, report these errors to Broadcom Technical Support.</p>
<p>1801: Create vport work array FAILED: cannot do scsi_host_get</p> <p>The driver was unable to get a reference to a SCSI host.</p> <p>Data: None</p> <p>Severity: Warning</p> <p>Log: LOG_VPORT verbose</p> <p>Action: Software driver warning. If this issue persists, report these errors to Broadcom Technical Support.</p>
<p>1816: FLOGI NPIV supported, response data <port></p> <p>The fabric reports support for NPIV upon FLOGI.</p> <p>Data: (1) response_multiple_NPort</p> <p>Severity: Warning</p> <p>Log: LOG_VPORT verbose</p> <p>Action: No action needed, informational.</p>
<p>1817: Fabric does not support NPIV - configuring single port mode</p> <p>The fabric reports no support for NPIV upon FLOGI.</p> <p>Data: None</p> <p>Severity: Warning</p> <p>Log: LOG_VPORT verbose</p> <p>Action: No action needed, informational.</p>
<p>1818: VPort failed init, mbxCmd <mailbox command> READ_SPARM mbxStatus <mailbox status> , rc = <status></p> <p>A pending mailbox command that was issued to initialize the port failed.</p> <p>Data: (1) mbxCommand (2) mbxStatus (3) rc</p> <p>Severity: Error</p> <p>Log: LOG_VPORT verbose</p> <p>Action: Software driver error. If this issue persists, report these errors to Broadcom Technical Support.</p>
<p>1819: Unrecognized lpfc_sli_mode parameter: <mode></p> <p>An attempt was made to set the SLI mode to an invalid value. The valid values for the SLI mode are 0, 2, and 3.</p> <p>Data: (1) lpfc_sli_mode</p> <p>Severity: Error</p> <p>Log: LOG_VPORT verbose</p> <p>Action: Correct the lpfc_sli_mode driver parameter setting. Valid values are 0, 2, and 3.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>1820: Unable to select SLI-3. Not supported by adapter. The HBA is incapable of operating in a given mode. Data: None Severity: Error Log: LOG_VPORT verbose Action: SLI-3 mode is only available on some HBAs. Do not attempt to force the SLI mode to 3 on HBAs that do not support SLI-3 mode. This is an informational message. HBAs that do not support SLI-3 will be configured to run in SLI-2 mode, but it is recommended to use the auto setting (0).</p>
<p>1821: Create VPORT failed. Invalid WWN format The port could not be created due to an invalid WWNN or WWPN format. Data: None Severity: Error Log: LOG_VPORT verbose Action: Provide a valid WWN when creating Vports.</p>
<p>1822: Invalid <name>: <xx: xx: xx: xx: xx: xx: xx> An invalid WWN was used when creating a VPort. Data: (1) type_name (2) wwn[1] (3) wwn[3] (3) wwn[5] (4) wwn[7] Severity: Error Log: LOG_VPORT verbose Action: When creating a VPort, you must furnish a valid WWN.</p>
<p>1823: Create VPORT failed. Duplicate WWN on HBA. The port could not be created because it would duplicate an existing WWNN HBA address. The resources for the port were discarded. Data: None Severity: Error Log: LOG_VPORT verbose Action: Provide a unique WWN.</p>
<p>1824: NPIV enabled: Override lpfc_sli_mode parameter (<mode>) to auto(0) The lpfc_enable_npiv and lpfc_sli_mode driver parameter settings conflict. The HBA must be configured for SLI-3 mode to support NPIV. Data: (1) lpfc_sli_mode Severity: Error Log: LOG_VPORT verbose Action: This is an informational message that indicates that the lpfc_enable_npiv and lpfc_sli_mode parameter settings are not compatible. Resolve the parameter conflict by setting the SLI mode to 0 or 3 or, if SLI-2 mode is required, disable NPIV.</p>
<p>1825: Vport Created. This message indicates that a port was created in the system. It is displayed at this level to ensure that it always appears at all log levels. Data: None Severity: Error Log: LOG_VPORT verbose Action: No action, informational.</p>
<p>1826: Vport Disabled. The port had to be disabled in the system. Data: None Severity: Error Log: LOG_VPORT verbose Action: No action, informational.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>1827: Vport Enabled The port was enabled after possible recovery from some errors. Data: None Severity: Error Log: LOG_VPORT verbose Action: No action, informational.</p>
<p>1828: Vport Deleted A Vport was deleted. Data: None Severity: Error Log: LOG_VPORT verbose Action: No action, informational.</p>
<p>1829: CT command failed to delete objects on fabric. A command issued to the fabric to delete an associated resource for an object, such as for a port, failed. Data: None Severity: Error Log: LOG_VPORT verbose Action: Software driver error. If this issue persists, report these errors to Broadcom Technical Support.</p>
<p>1830: Signal aborted mbxCmd <command> A pending mailbox command was aborted because the thread received a signal. Data: None Severity: Error Log: LOG_VPORT verbose Action: Retry the attempted command.</p>
<p>1831: Create VPORT Interrupted The port creation process was unexpectedly interrupted at a critical time, and the operation was unsuccessful. Data: None Severity: Error Log: LOG_VPORT verbose Action: The process was interrupted while creating a VPort. Retry the command.</p>
<p>1832: No pending MBOX command to handle Data: None Severity: Error Log: LOG_MBOX Action: None required.</p>
<p>1835: Vport discovery quiesce failed: state <port_state> fc_flags <fc_flag> wait msec <jiffies_to_msecs(jiffies - start_time)> Could not pause discovery on this VPort. Data: None Severity: Error Log: LOG_VPORT Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>1836: Could not issue unreg_login(all_rpis) status <rc> The unreg_login cannot be issued. Data: None Severity: Error Log: LOG_MBOX, LOG_VPORT Action: None required.</p>
<p>1837: Vport_delete failed: Cannot delete static vport Static VPorts cannot be deleted. Data: None Severity: Error Log: LOG_VPORT Action: None required.</p>
<p>1838: Failed to INIT_VPI on vpi <vpi> status <rc> Failed to INIT_VPI. Data: None Severity: Error Log: LOG_VPORT Action: None required.</p>
<p>2000: Failed to allocate mbox for read_FCF cmd Failed to allocate mailbox for READ_FCF command. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>2001: Unable to allocate memory for issuing SLI_CONFIG_SPECIAL mailbox command Unable to allocate memory for issuing the SLI_CONFIG_SPECIAL mailbox command. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>2002: Error Could not grow rpi count An error occurred because the RPI count could not be increased. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>2004: Failed to allocate XRI.last XRITAG is <XRI> Max XRI is <MAX_XRI>, Used XRI is <USED_XRI>. All XRIs are in use. Data: None Severity: Warning Log: LOG_SLI Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>2005: Unable to deregister pages from HBA: <rc> The SGL pages could not be unregistered from the firmware. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>2007: Only Limited Edition cmd Format supported <iocb.ulpCommand> The SGL pages could not be unregistered from the firmware. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>2008: Error <rc> posting all rpi headers The RPI headers could not be posted to the firmware. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>2009: Failed to allocate mbox for ADD_FCF cmd Failed to allocate mailbox for ADD_FCF command. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>2010: Resume RPI Mailbox failed status <status>, mbxStatus <mbx status> Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>2011: Unable to allocate memory for issuing SLI_CONFIG_SPECIAL mailbox command Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>2012: Mailbox failed , mbxCmd <mbx_cmd> READ_CONFIG, mbxStatus <mbx status> The READ_CONFIG mailbox command failed. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>2013: Could not manually add FCF record 0, status <rc> Could not add an FCF record to the FCF list. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>2014: Invalid command <iocb.ulpCommand> The IOCB command is invalid. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>2015: Invalid CT %x command <iocb.ulpCommand> Invalid command type in the IOCB is not supported. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>2017: REG_FCFI mbxStatus error <mbx status> HBA state <port_state> The REG_FCFI mailbox command has failed. Data: None Severity: Error Log: LOG_MBOX Action: None required.</p>
<p>2018: REG_VFI mbxStatus error <mbx status> HBA state <port_state> The REG_VFI mailbox command has failed. Data: None Severity: Error Log: LOG_MBOX Action: None required.</p>
<p>2020: Failed to allocate mbox for ADD_FCF cmd Failed to allocate mailbox for ADD_FCF command. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>2022: VPI Mailbox failed status <status>, mbxStatus <mbxStatus> The INIT VPI mailbox command has failed. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>2401: Failed to allocate memory for ELS XRI management array of size <els_xri_cnt>. Initialization failed to allocate memory for the ELS XRI management array. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>2500: EQ_CREATE mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc></p> <p>The mailbox command sent to create the event queue has failed.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_INIT</p> <p>Action: None required.</p>
<p>2501: CQ_CREATE mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc></p> <p>The mailbox command sent to create the completion queue has failed.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_INIT</p> <p>Action: None required.</p>
<p>2502: MQ_CREATE mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc></p> <p>The mailbox command sent to create the mailbox queue has failed.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_INIT</p> <p>Action: None required.</p>
<p>2503: WWQ_CREATE mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc></p> <p>The mailbox command sent to create the work queue has failed.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_INIT</p> <p>Action: None required.</p>
<p>2504: RQ_CREATE mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc></p> <p>The mailbox command sent to create the receive queue has failed.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_INIT</p> <p>Action: None required.</p>
<p>2505: EQ_DESTROY mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc></p> <p>The mailbox command sent to delete the event queue has failed.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_INIT</p> <p>Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>2506: CQ_DESTROY mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc></p> <p>The mailbox command sent to delete the completion queue has failed.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_INIT</p> <p>Action: None required.</p>
<p>2507: MQ_DESTROY mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc></p> <p>The mailbox command sent to delete the mailbox queue has failed.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_INIT</p> <p>Action: None required.</p>
<p>2508: WQ_DESTROY mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc></p> <p>The mailbox command sent to delete the work queue has failed.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_INIT</p> <p>Action: None required.</p>
<p>2509: RQ_DESTROY mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc></p> <p>The mailbox command sent to delete the work queue has failed.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_INIT</p> <p>Action: None required.</p>
<p>2510: RQ_DESTROY mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc></p> <p>The mailbox command sent to delete the work queue has failed.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_INIT</p> <p>Action: None required.</p>
<p>2511: POST_SGL mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>n</p> <p>The mailbox command sent to post the SGL pages to the firmware has failed.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_INIT</p> <p>Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>2512: REMOVE_ALL_SGL_PAGES mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc></p> <p>The mailbox command sent to delete the SGL pages from the firmware has failed.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_INIT</p> <p>Action: None required.</p>
<p>2513: POST_SGL_BLOCK mailbox command failed status <shdr_status> add_status <shdr_add_status> mbx status <rc></p> <p>The mailbox command sent to post the SGL pages to the firmware has failed.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_INIT</p> <p>Action: None required.</p>
<p>2514: POST_RPI_HDR mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc></p> <p>The mailbox command sent to post the RPUI header pages to the firmware has failed.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_INIT</p> <p>Action: None required.</p>
<p>2515: ADD_FCF_RECORD mailbox failed with status <rc></p> <p>The mailbox command to add the FCF record has failed.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_INIT</p> <p>Action: None required.</p>
<p>2516: DEL_FCF of default FCF Index failed mbx status <rc>, status <shdr_status> add_status<shdr_add_status></p> <p>The mailbox command to delete the FCF record has failed.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_SLI</p> <p>Action: None required.</p>
<p>2517: Unregister FCFI command failed status %d, mbxStatus x%x", rc, bf_get(lpfc_mqe_status, &mbx->u.mqe)</p> <p>The driver was unable to unregister the FCFI from the firmware.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_SLI</p> <p>Action: None required.</p>
<p>2518: Requested to send 0 NOP mailbox cmd</p> <p>Data: None</p> <p>Severity: Warning</p> <p>Log: LOG_INIT</p> <p>Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>2519: Unable to allocate memory for NOP mailbox command Memory allocation for this mailbox command has failed.</p> <p>Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>2520: NOP mailbox command failed status x%x add_status x%x mbx status x%x, shdr_status, shdr_add_status, rc.</p> <p>The NOP mailbox command has failed.</p> <p>Data: None Severity: Warning Log: LOG_INIT Action: None required.</p>
<p>2521: READ_FCF_RECORD mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx</p> <p>The READ_FCF_RECORD mailbox command failed.</p> <p>Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>2523: Allocated DMA memory size (<alloc_len>) is less than the requested DMA memory size (<req_len>)</p> <p>The ADD_FCF_RECORD mailbox command failed to retrieve the length required from the firmware.</p> <p>Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>2524: Failed to get the non-embedded SGE virtual address</p> <p>The READ_FCF_RECORD mailbox command could not retrieve the SGE that was requested.</p> <p>Data: None Severity: Error Log: LOG_MBOX Action: None required.</p>
<p>2527: Failed to allocate non-embedded SGE array.</p> <p>Failed to allocate the non-embedded SGE array.</p> <p>Data: None Severity: Error Log: LOG_MBOX Action: None required.</p>
<p>2528: Mailbox command <vpi> cannot issue</p> <p>The mailbox command could not be issued because the mailbox interrupt is disabled.</p> <p>Data: (1) mbxCommand (2) sli_flag (3) flag Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>2529: Mailbox command <vpi> cannot issue Data: (1) mbxCommand (2) sli_flag (3) flag Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.</p>
<p>2530: Mailbox command <vpi> cannot issue The SLI layer in the driver is inactive. Data: (1) mb.mbxCommand (2) sli_flag (3) flag Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.</p>
<p>2531: Mailbox command <vpi> cannot issue Data: (1) mb.mbxCommand (2) sli_flag (3) flag Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.</p>
<p>2532: Mailbox command <vpi> (<mbxCommand>) cannot issue The mailbox bootstrap code detected that the SLI layer is active. Data: (1) sli4_mbox_opcode (2) sli_flag,(3) MBX_POLL Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.</p>
<p>2533: Mailbox command <vpi> (<mbxCommand>) cannot issue Data: (1) sli4_mbox_opcode (2) sli_flag (3) MBX_NOWAIT Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.</p>
<p>2535: Unsupported RQ count. (<entry_count>) The receive queue ring can only be 512, 1024, 2048, or 4096. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>2536: Unsupported RQ count. (<entry_count>) The receive queue ring can only be 512, 1024, 2048, or 4096. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>2537: Receive Frame Truncated! The receive unsolicited handler detected a truncated frame. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>2539: Dropped frame rctl:%s type:%s\n An unsupported frame was received by the port and dropped. Data: (1) rctl_names[fc_hdr->fh_r_ctl], (2) type_names[fc_hdr->fh_type] Severity: Error Log: Always Action: No action needed, informational.</p>
<p>2540: Ring <ring #> handler: unexpected Rctl <fh_rctl> Type <fh_type> The received frame has an unsupported RCTL or FH_TYPE. Data: None Severity: Warning Log: LOG_SLI Action: None required.</p>
<p>2541: Mailbox command <vpi> (<mbxCommand>) cannot issue Data: (1) sli_mbox_opcode (2) sli_flag (3) flag Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.</p>
<p>2542: Try to issue mailbox command <vpi> (<mbxCommand>) synchronously ahead of async mailbox command queue Attempting to send a synchronous mailbox command ahead of the asynchronous mailbox commands. Data: (1) sli4_mbx_opcode or sli_mbox_opcode, (2) sli_flag, (3) flag Severity: Warning Log: LOG_MBOX, LOG_SLI Action: None required.</p>
<p>2543: Mailbox command <vpi> (<mbxCommand>) cannot issue The mailbox command does not have all of the fields set correctly. Data: (1) sli_mbox_opcode (2) sli_flag (3) flag Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.</p>
<p>2544: Mailbox command <vpi> (<mbxCommand>) cannot issue The HBA cannot be accessed on the PCI bus. Data: (1) sli_mbox_opcode (2) sli_flag (3) flag Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.</p>
<p>2546: New FCF found index <index> tag <event_tag> A new FCF has been found. Data: None Severity: Error Log: LOG_DISCOVERY Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>2547: Read FCF record failed Could not read the FCF record from the firmware. Data: None Severity: Error Log: LOG_DISCOVERY Action: None required.</p>
<p>2548: FCF Table full count <count> tag <event_tag> The FCF table is full. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>2549: FCF disconnected from network index <index> tag <event_tag> The FCF has disconnected from the network. Data: None Severity: Error Log: LOG_DISCOVERY Action: None required.</p>
<p>2550: UNREG_FCFI mbxStatus error <u.mb.mbxStatus> HBA state <port_state>. The UNREG_FCFI mailbox command has failed. Data: None Severity: Error Log: LOG_DISCOVERY, LOG_MBOX Action: None required.</p>
<p>2551: UNREG_FCFI mbox allocation failed HBA state <port_state> The allocation for the UNREG_FCFI mailbox command has failed. Data: None Severity: Error Log: LOG_DISCOVERY, LOG_MBOX Action: None required.</p>
<p>2552: UNREG_FCFI issue mbox failed rc <rc> HBA state <port_state>. The UNREG_FCFI mailbox command has failed. Data: None Severity: Error Log: LOG_DISCOVERY, LOG_MBOX Action: None required.</p>
<p>2553: lpfc_unregister_unused_fcf failed to read FCF record HBA state. Data: None Severity: Error Log: LOG_DISCOVERY, LOG_MBOX Action: None required.</p>
<p>2554: Could not allocate memory for fcf record Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>2555: UNREG_VFI mbxStatus error <u.mb.mbxStatus> HBA state <port_state> The UNREG_VFI mailbox command has failed. Data: None Severity: Error Log: LOG_DISCOVERY, LOG_MBOX Action: None required.</p>
<p>2556: UNREG_VFI mbox allocation failed HBA state <port_state> Could not allocate memory for UNREG_VFI mailbox command. Data: None Severity: Error Log: LOG_DISCOVERY, LOG_MBOX Action: None required.</p>
<p>2557: UNREG_VFI issue mbox failed rc <rc> HBA state <port_state> Could not issue the UNREG_VFI mailbox command. Data: None Severity: Error Log: LOG_DISCOVERY, LOG_MBOX Action: None required.</p>
<p>2558: ADD_FCF_RECORD mailbox failed with status<shdr_status> add_status <shdr_add_status> The ADD_FCF_RECORD mailbox command has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>2559: Block sgl registration required DMA size <reqlen> great than a page. Attempting to register more SGEs with the firmware than can fit in a page. Data: None Severity: Warning Log: LOG_INIT Action: None required.</p>
<p>2560: Failed to allocate mbox cmd memory\n Failed to allocate mailbox command memory. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>2561: Allocated DMA memory size (<alloclen>) is less than the requested DMA memory size (<reqlen>) Could not get the memory required for the number of XRIs that are attempting to be posted. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>2562: No room left for SCSI XRI allocation max_xri=<sli4_hba.max_cfg_param.max_xri>, els_xri=<els_x-ri_cnt>n</p> <p>The number of allocated XRIs has reached the max_xri value.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_SLI</p> <p>Action: None required.</p>
<p>2563: Failed to allocate memory for SCSI XRI management array of size <sli4_hba.scsi_xri_max>. Initialization could not allocate memory to hold the XRIs.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_SLI</p> <p>Action: None required.</p>
<p>2564: POST_SGL_BLOCK mailbox command failed status <shdr_status> add_status <shdr_add_status> mbx status <rc></p> <p>The list of XRI SGEs failed to be registered with the firmware.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_SLI</p> <p>Action: None required.</p>
<p>2566: Failed to allocate table entry</p> <p>Failed to allocate connection table entry.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_INIT</p> <p>Action: None required.</p>
<p>2567: Config region 23 has bad signature</p> <p>The driver was unable to read Config Region 23 because it has an invalid signature.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_INIT</p> <p>Action: None required.</p>
<p>2568: Config region 23 has bad version</p> <p>The driver was unable to read Config Region 23 because it is an invalid version.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_INIT</p> <p>Action: None required.</p>
<p>2569: lpfc_dump_fcoe_param: memory allocation failed</p> <p>Memory allocation has failed.</p> <p>Data: None</p> <p>Severity: Warning</p> <p>Log: LOG_MBOX</p> <p>Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>2570: Failed to read FCoE parameters The driver failed to read FCoE parameters. Data: None Severity: Error Log: LOG_MBOX, LOG_INIT Action: None required.</p>
<p>2572: Failed allocate memory for fast-path per-EQ handle array Failed to allocate memory for the fast-path per-EQ handle array. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>2573: Failed allocate memory for msi-x interrupt vector entries The driver was unable to allocate memory during initialization of the MSI-X interrupt array. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>2574: Not enough EQs (<sli4_hba.max_cfg_param.max_eq>) from the pci function for supporting FCP EQs (<cfg_fcp_eq_count>) Failed to create the minimum fast-path event queues. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>2575: Not enough EQs (<max_eq>) from the pci function for supporting the requested FCP EQs (<cfg_fcp_eq_count>), the actual FCP EQs can be supported: <eq_count> The driver was not configured with enough fast-path event queues. Data: None Severity: Warning Log: LOG_INIT Action: None required.</p>
<p>2576: Failed allocate memory for fast-path EQ record array Failed to allocate memory for the fast-path EQ record array. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>2577: Failed allocate memory for fast-path CQ record array Failed to allocate memory for the fast-path EQ record array. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>2578: Failed allocate memory for fast-path WQ record array Failed to allocate memory for the fast-path EQ record array. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>2579: Slow-path wqe consume event carries miss-matched qid: wcqe-qid=<wcqe_qid>, sp-qid=<sp_qid> The consumed entry does not have the queueID of the slow-path. Data: None Severity: Warning Log: LOG_SLI Action: None required.</p>
<p>2580: Fast-path wqe consume event carries miss-matched qid: wcqe-qid=<fcq_wqid>. The consumed entry does not have the queueID of the fast path. Data: None Severity: Warning Log: LOG_SLI Action: None required.</p>
<p>2581: Not enough WQs (<sli4_hba.max_cfg_param.max_wq>) from the pci function for supporting FCP WQs (<cfg_fcq_wq_count>) The driver was not configured with the minimum number of fast-path work queues. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>2582: Not enough WQs (<max_wq>) from the pci function for supporting the requested FCP WQs (<cfg_fcq_wq_count>), the actual FCP WQs can be supported: <wq_count> The driver was not configured with enough fast-path work queues. Data: None Severity: Warning Log: LOG_INIT Action: None required.</p>
<p>2593: The FCP EQ count(<cfg_fcq_eq_count>) cannot be greater than the FCP WQ count(<cfg_fcq_wq_count>), limiting the FCP EQ count to <cfg_fcq_wq_count> The fast-path event queue cannot be greater than the fast-path work queue count. Data: None Severity: Warning Log: LOG_INIT Action: None required.</p>
<p>2597: Mailbox command <vpi> (<mbxCommand>) cannot issue Synchronou(2) sli_flag (3) flag Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>2598: Adapter Link is disabled. The adapter link has been disabled. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>2599: Adapter failed to issue DOWN_LINK mbox command rc <rc> The driver was unable to issue the DOWN_LINK mailbox command. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>2600: lpfc_sli_read_serdes_param failed to allocate mailbox memory Failed to allocate mailbox memory. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>2605: lpfc_dump_static_vport: memory allocation failed Failed to allocate mailbox memory. Data: None Severity: Error Log: LOG_MBOX Action: None required.</p>
<p>2606: No NPIV Fabric support No NPIV Fabric support. Data: None Severity: Error Log: LOG_ELS Action: None required.</p>
<p>2607: Failed to allocate init_vpi mailbox Failed to allocate the INIT_VPI mailbox command. Data: None Severity: Error Log: LOG_MBOX Action: None required.</p>
<p>2608: Failed to issue Init VPI mailbox The driver was unable to send an INIT_VPI mailbox command. Data: None Severity: Error Log: LOG_MBOX Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>2609: Init VPI mailbox failed <u.mb.mbxStatus> The INIT_VPI mailbox command failed. Data: None Severity: Error Log: LOG_MBOX Action: None required.</p>
<p>2610: UNREG_FCFI mbox allocation failed Failed to allocate mailbox memory. Data: None Severity: Error Log: LOG_DISCOVERY, LOG_MBOX Action: None required.</p>
<p>2611: UNREG_FCFI issue mbox failed Could not issue the UNREG_FCFI mailbox command. Data: None Severity: Error Log: LOG_DISCOVERY, LOG_MBOX Action: None required.</p>
<p>2619: Config region 23 has bad signature Configuration region 23 has an invalid signature. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>2620: Config region 23 has bad version Configuration region 23 has an invalid version. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>2621: Failed to allocate mbox for query firmware config cmd Failed to allocate mailbox memory. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>
<p>2622: Query Firmware Config failed mbx status <rc>, status <shdr_status> add_status <shdr_add_status> Could not read the firmware configuration. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>

Table 9: LPFC Error Log Messages and Their Descriptions (Continued)

<p>2623: FCoE Function not supported by firmware. Function mode = <function_mode> FCoE is not supported by this firmware. Data: None Severity: Error Log: LOG_SLI Action: Use the Emulex OneCommand Manager application to update to the latest firmware.</p>
<p>2707: Ring <Ring#> handler: Failed to allocate iocb Rctl <fh_rctl> Type <fh_type> received The driver was unable to allocate memory to send a query config mailbox command. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>2717: CT context array entry [<index>] over-run: oxid:<fh_ox_id>, sid:<fh_SID> All of the array slots to hold buffers that are passed to the application are in use. Data: None Severity: Warning Log: LOG_ELS Action: None required.</p>
<p>2718: Clear Virtual Link Received for VPI <index> tag <event_tag> A Clear virtual link was received from the fabric for this VPI. Data: None Severity: Error Log: LOG_DISCOVERY Action: None required.</p>
<p>2719: Invalid response length: tgt <TGT_ID> lun <LUN> cmd <CMD> rsplen <RSPLEN> The response length for this FCP command is not supported. Data: None Severity: Error Log: LOG_FCP Action: None required.</p>
<p>2721: ndlp null for oxid %x SID %x\n, icmd->ulpContext, dfchba->ct_ctx[tag].SID); The Node value for this SID is not in the node list. Data: None Severity: Warning Log: LOG_ELS Action: None required.</p>
<p>2726: READ_FCF_RECORD Indicates empty FCF table The driver requested the firmware provide a list of FCF entries to connect to and the firmware responded that the FCF table is empty. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>

3.4.2 Ethernet Driver Log Messages for NIC PPC CNAs

The following section describes how to retrieve and interpret Ethernet log messages.

3.4.2.1 Retrieving Ethernet Driver Log Messages

Ethernet driver messages are logged in the `/var/log/messages` file. This log file is an ASCII text file and can be viewed and searched with a text editor such as `vim`. A log file automatically rotates as it gets larger. Rotated log files are named `messages.x`, where `x` is an integer.

To search the log file for error messages, at the command prompt type:

```
# cd /var/log
# vim messages
```

A message is displayed similar to the following:

```
Aug 15 09:57:48 S74 kernel: Invalid MTU requested. Must be between 256 and 8974 bytes
```

3.4.2.2 Ethernet Driver Log Messages and Their Descriptions

When reporting an issue with the adapter, check the kernel message log using the `dmesg` command or the `/var/log/messages` file, and report any of these entries that are present.

There are three types of Ethernet log messages: error, information, and warning.

Ethernet driver warning messages logged by an adapter start with `lpnic <BID>`, where `<BID>` is the PCI bus identifier string. For example:

```
lpnic 0000:0d:00.1: MTU must be between 256 and 9000 bytes.
```

NOTE: In the following table, `<D>`, `<DD>`, or `<DDD>` refers to decimal values that appear in the log messages, and `<S>` refers to strings.

Table 10: Ethernet Driver Log Messages of Type

Ethernet Log Message	Description	Type
<code><S> initialization failed</code>	Initialization of the adapter or allocation of a resource for initializing the driver failed. In most cases, this message is accompanied by a more specific error message. Try rebooting the system after power cycling. If the issue persists, this could be a symptom of a hardware issue or corrupted firmware.	Error
<code><S>: Physical link <S></code>	This is an informational message about a change in physical link status.	Information
<code>Adapter busy, could not reset FW. Reboot server to activate new FW</code>	After flashing firmware on the adapter, the adapter is busy activating the new firmware. Reboot the machine for the new firmware to be active.	Error
<code>Adapter in non recoverable error</code>	Resetting the adapter failed, as adapter has gone into non-recoverable state.	Error
<code>Adapter recovery failed</code>	The adapter error recovery failed.	Error
<code>Adapter recovery succeeded</code>	The adapter has successfully recovered after an error.	Information
<code>Cannot set rx filter flags <XX></code> <code>Interface is capable of <XX> flags only</code>	Failed to change RX filter settings on an interface as requested, as interface does not have the capability.	Warning

Table 10: Ethernet Driver Log Messages of Type (Continued)

Ethernet Log Message	Description	Type
Could not create sysfs group	The creation of the <code>flash_fw</code> entry under the <code>/sys/class/net/eth<x></code> failed. The driver is fully functional, but you cannot install later firmware versions on the adapter.	Error
Failed to enable VLAN promiscuous mode	Firmware command failed to enable VLAN promiscuous mode.	Error
Could not get CRC for <DD> region	The driver could not get enough information from the adapter to decide whether a region from a firmware image should be flashed. The driver skips updating this section. This is a very unlikely error.	Error
Could not get <DD> MSI-x vector(s) Using <DD> vector(s)	Enabled only a subset of MSI-x vectors requested.	Information
Could not set PCI DMA Mask	The operating system call to set the DMA mask failed.	Error
created <DD> RX queue(s)	An informational message logging number of receive rings was created.	Information
created <DD> TX queue(s)	An informational message logging number of transmit rings was created.	Information
<DD> VFs are already enabled	When unloading the driver while a VF is still assigned to VM, the VFs remain enabled. When the driver is reloaded, this message is logged with the number of VFs that are enabled and active.	Information
Debug data event - <DDDD>, <DDDD>, <DDDD>, <DDDD>	These log messages are generated by firmware to identify certain infrequent hardware and firmware events and aid further analysis if required. These messages by themselves are not indicative of any problem. Data from these logs might help Broadcom troubleshoot other problems, if any, that are observed along with these messages.	Information
enabled <D> (out of <DD>) MSI-x vector(s) for NIC	The number of MSI-x vectors enabled for the NIC function.	Information
Enable VLAN promiscuous mode Disabling VLAN promiscuous mode	This is an informational log message when the driver enables or disables VLAN promiscuous mode.	Information
Error detected in the card ERR: sliport status <DDDDDD> ERR: sliport error1 <DDDDDD> ERR: sliport error2 <DDDDDD>	An SLIPOINT error was detected in the adapter, followed by status and error messages.	Error
Error in mapping PCI BARs	Initialization of the adapter failed due to an error while mapping PCI BARs.	Error
Failed to optimize SR-IOV resources	When the adapter is in an SRIOV-capable configuration, driver failed to redistribute the PF pool resources among the PF and requested number of VFs.	Error
FW config: function_mode=<DD>, function_caps=<DD>	A message indicating the function mode and function capabilities set on the adapter during initialization.	Information
FW dump deleted successfully	Previously generated FW dump is successfully deleted.	Information
FW dump initiated successfully	FW dump is successfully initiated on the adapter.	Information
FW dump not generated	Failed to initiate FW dump on the adapter.	Error
Firmware flashed successfully	This is an informational message that the firmware on the adapter has been updated.	Information
Firmware load error	Updating the adapter with new firmware failed. Usually this message is accompanied by a detailed message on the failure.	Error
Firmware load not allowed (interface is down)	Flashing firmware failed because the status of the interface is down.	Error

Table 10: Ethernet Driver Log Messages of Type (Continued)

Ethernet Log Message	Description	Type
Firmware update in progress	Firmware update is in progress on the adapter.	Information
Flashing firmware file <filename>	This is an informational message that the firmware in the adapter is being updated with the firmware image in the file indicated.	Information
Flashing section type <DD> failed	The driver could not get enough information from the FW UFI image to flash a section in the UFI image, after which the FW download will be aborted.	Error
FW image size should be multiple of 4	The FW UFI image size is invalid. Download the correct FW image and try again.	Error
FW version is <DD.DD.DDD.DDD>	This is an informational message to log the FW version loaded on the adapter.	Information
INTx request IRQ failed - err <DDD>	The request for INTx interrupt registration failed. The driver is nonfunctional if the INTx interrupt cannot be registered.	Error
Invalid Cookie. FW image may be corrupted	The firmware image under /lib/firmware/<filename> does not have the expected cookie. The firmware in this file will not be flashed. Copy the proper file and try flashing again.	Error
Invalid digital signature	Updating the adapter with new firmware failed. The firmware image being updated does not have a valid digital signature.	Error
invalid mailbox completion	The driver received an unexpected completion status for a firmware command.	Error
Lancer FW download timed out	FW command timed out, while downloading FW image.	Error
Link down event type: <DD>	The reason that the ASIC signaled the link status as down. Possible values are: 0 = Link down due to reasons other than those listed here. 1 = Link down caused by Dynamic Control channel protocol. 3 = Link down triggered by Virtual NIC configuration (for example, a zero bandwidth is assigned to a VNIC). 4 = Link down caused by Ethernet Pause frame flooding. 5 = Link down due to physical thermal temperature going up.	Information
Link down/Link up	This is an informational message about a change in link status.	Information
LPVID <DD>	VLAN filter configured in BIOS in multichannel configuration.	Information
Mac address assignment failed for VF <DD> MAC <S>M set on VF <DD> Failed	Firmware command fails to add a MAC address to a virtual function during initialization or when requested.	Error
MAC address change to <S> failed	Failed to change to new MAC address as requested.	Warning
Mac hash table alloc failed	The hash table to remember learned MAC addresses could not be allocated.	Error
Max: txqs <DD>, rxqs <DD>, rss <DD>, eqs <DD>, vfs <DD> Max: uc-macs <DD>, mc-macs <DD>, vlans <DD>	An informational message logging maximum resources available to the function. Resources include TX rings, RX rings, RX rings with RSS capability, number of VFs supported, unicast MACs, multicast MACs and VLAN filters, respectively.	Information
Memory allocation failure	The driver could not allocate the memory required for the requested operation.	Error
Memory allocation failure during GET_MAC_LIST	Failed to allocate memory to issue GET_MAC_LIST FW command.	Error
Missing digital signature	Updating the adapter with new firmware failed. The firmware image being updated does not have a digital signature.	Error

Table 10: Ethernet Driver Log Messages of Type (Continued)

Ethernet Log Message	Description	Type
Module param rx_frag_size must be 2048/4096/8192. Using 2048	An unsupported receive buffer size was passed for the rx_frag_size module parameter. The driver ignores the specified value and uses the default RX buffer size of 2048.	Warning
MSIx enable failed	A request for enabling MSIx interrupts registration failed. The driver will then use INTx interrupts.	Warning
MAC address changed to <S>	MAC address is changed successfully as requested.	Information
MSIX request IRQ failed -err <DDD>	The request for MSI-X interrupt registration failed. The driver will use INTx interrupts.	Warning
MTU changed from <DDD> to <DDD> bytes	This is an informational message that the MTU value changed as requested.	Information
MTU must be between 256 and 9000 bytes	A request to change the MTU was issued with an invalid MTU value. The request failed and MTU will not be changed.	Warning
opcode <DDD>-<D> failed: status <DD>-<DD>	A firmware command with opcode failed with the indicated status code and extended status code.	Error
PCIe error reporting enabled	PCIe error reporting is successfully enabled on this function.	Information
POST timeout; stage=<DD>	The POST of the adapter failed. This is an indication of a hardware or firmware issue. Try rebooting the system after a reset.	Error
Previous dump not cleared, not forcing dump	Initiating dump on the adapter failed because the previous dump is present. Clear the previous dump (using the <code>ethtool -W eth<x> 2</code> command) before initiating new dump.	Error
queue_setup failed	Firmware command failed to create the requested number of queues.	Error
Reboot server to activate new FW	After upgrading the firmware when virtual functions are loaded, the driver cannot reset the adapter to activate the new firmware. Reset the server to activate new firmware after upgrading the firmware on the adapter successfully.	Information
Resetting adapter to activate new FW	After upgrading the FW image on the adapter, driver is resetting the adapter to activate new FW without resetting the server.	Information
Setting HW VLAN filtering failed	Adding a VLAN filter to HW failed.	Error
SRIOV enable failed	Could not enable SRIOV because the call to enable SRIOV failed.	Error
TX-rate must be between 100 and <DDDD> Mbps	Request to change transmission rate was issued with an invalid TX rate value. The request failed, and speed will not be changed.	Error
TX-rate must be a multiple of <DDD> Mbps	Transmission rate on a virtual function should be given as multiples of %1Mb/s of link speed.	Error
TX-rate setting not allowed when link is down	Transmission rate on a virtual function cannot be modified when the physical link is down.	Error
Tx-rate setting of <DDD>Mbps on VF<DD> failed : <DDD>	The firmware command failed to change transmission rate as requested.	Error
txq<D>: cleaning <D> pending tx-wrbs	The driver did not get completions for some transmit requests from the adapter while unloading the driver. This usually indicates an issue with the adapter.	Error
Flash image is not compatible with adapter	The requested UFI image is not compatible with the chip on which the flash was requested.	Error
Unknown debug event <DD>	Other than QNQ type debug event, the driver logs as unknown for other asynchronous debug events received.	Warning
Unqualified SFP+ detected on <D> from <S> part no: <S>	The SFP module indicated in the message is not qualified or supported by Broadcom.	Information

Table 10: Ethernet Driver Log Messages of Type (Continued)

Ethernet Log Message	Description	Type
Unrecoverable Error detected in the adapter Please reboot server to recover UE LOW: <S> bit set UE HIGH: <S> bit set	An unrecoverable error is detected in the adapter that requires a reboot to recover. Low and high bits are set in the data path in which error occurred.	Error
User has aborted FW download	The user requested abort when FW download is in progress.	Error
Using profile <DD>	An informational message of the profile type currently enabled in the adapter.	Information
VF<DD> has FILTMGMT privilege	VF is provided with FILTMGMT privilege to program MAC/VLAN filters.	Information
VF is not privileged to issue opcode <DD>-<DD>	VF does not have enough privileges to issue opcode mentioned in the log message.	Warning
VFs are assigned to VMs: not disabling VFs	Do not disable virtual functions on the port during driver unload, on which VFs are assigned to guests that are powered ON.	Warning
VF setup failed	Failed to create VFs as FW commands failed to provide required resources.	Error
VLAN <DDDD> config on VF <DD> failed	Firmware command failed to set VLAN filter as requested.	Error
Waiting for FW to be ready after EEH reset	After a PCI EEH reset, wait until firmware becomes ready.	Information
Waiting for POST aborted	Waiting for POST of the adapter is aborted.	Error
Waiting for POST, <D>s elapsed	This is an informational log that logs the seconds elapsed while waiting for POST of the adapter.	Information
Disable/re-enable i/f in VM to clear Transparent VLAN tag	After clearing transparent VLAN tagging for a VF, disable and reenable the VF interface in guest operating system to clear VLAN-tagging for the traffic from guest.	Warning
Cannot disable VFs while they are assigned	Cannot disable VFs on the PF, when any of the VF is assigned to guest. Detach any VFs from guest to disable VFs on that PF.	Error
Invalid FW UFI file	The firmware UFI file is corrupted. Try flashing after copying correct UFI file.	Error
RSS hash key is longer than <DD> bytes	The requested RSS hash key is longer than 40 bytes. Request for a 40-byte RSS hash key.	Error
Invalid RSS hash key format	The hash key format must be in xx:yy:zz:aa:bb:cc format, meaning both the nibbles of a byte should be mentioned even if a nibble is zero.	Error
RSS hash key is too short (<DD> <DD>)	The requested RSS hash key is less than 40 bytes. Request only a 40-byte RSS hash key.	Error
<lpnic> version is <DD.DD.DDD.DDD>	This is an informational message of the loaded lpnic driver version.	Information
MGMT_STATUS_FLASHROM_SAVE_FAILED (0x17) / MGMT_ADDI_STATUS_DIGITAL_SIGNATURE_MISSING(0x57)	The firmware being used for flash is not digitally signed.	Error
MGMT_STATUS_FLASHROM_SAVE_FAILED (0x17) / MGMT_ADDI_STATUS_DIGITAL_SIGNATURE_INVALID (0x56)	The digital signature on the firmware being used for flash cannot be verified.	Error

Appendix A: Ethtool -S Option Statistics for PPC CNAs

The following table contains a list of `ethtool -S` option statistics and their descriptions.

Table 11: Ethtool -S Option Statistics

Name	Description
<code>dma_map_errors</code>	The number of packets dropped due to DMA mapping errors.
<code>eth_red_drops</code>	Received packets dropped due to the ASIC's Random Early Drop policy.
<code>forwarded_packets</code>	The number of packets generated by the ASIC internally. These packets are not handed to the host. This counter is shared across ports and all functions (NIC/FCoE).
<code>jabber_events</code>	The number of jabber packets received. Jabber packets are packets that are longer than the maximum size Ethernet frames and that have bad CRC.
<code>link_down_reason</code>	The reason that the ASIC signaled the link status as down. The various values are: <ul style="list-style-type: none"> ■ 0 – Link down due to reasons other than those listed here. ■ 1 – Link down caused by Dynamic Control channel protocol. ■ 3 - Link down triggered by Virtual NIC configuration (for example: zero bandwidth assigned to a VNIC). ■ 4 – Link down caused by Ethernet Pause frame flooding. ■ 5 – Link down due to physical thermal temperature going up.
<code>pmem_fifo_overflow_drop</code>	Received packets dropped when an internal FIFO going into main packet buffer tank (PMEM) overflows.
<code>rx_address_filtered</code>	Received packets dropped when they do not pass the unicast or multicast address filtering.
<code>rx_alignment_symbol_errors</code>	The number of packets dropped due to L1 alignment errors. This counter is on a per-port basis.
<code>rx_compl_err</code>	The number of RX completion errors received.
<code>rx_control_frames</code>	The number of control frames received.
<code>rx_crc_errors</code>	The number of packets dropped due to CRC errors.
<code>rx_dropped_header_too_small</code>	Received packets dropped when the IP header length field is less than 5.
<code>rx_dropped_runt</code>	Dropped receive packets due to runt packets (for example, packets shorter than the Ethernet standard).
<code>rx_dropped_tcp_length</code>	Received packets dropped when the TCP header length field is less than 5 or the TCP header length + IP header length is more than IP packet length.
<code>rx_dropped_too_short</code>	Received packets dropped when IP length field is greater than the actual packet length.
<code>rx_dropped_too_small</code>	Received packets dropped when IP packet length field is less than the IP header length field.
<code>rx_drops_mtu</code>	Received packets dropped when the frame length is more than 9018 bytes.
<code>rx_drops_no_erx_descr</code>	Received packets dropped due to the input receive buffer descriptor FIFO overflowing.
<code>rx_drops_no_pbuf</code>	Packets dropped due to lack of available HW packet buffers used to temporarily hold the received packets.
<code>rx_drops_no_tpre_descr</code>	Packets dropped because the internal FIFO to the offloaded TCP receive processing block is full. This could happen only for offloaded traffic.
<code>rx_drops_too_many_frags</code>	Received packets dropped when they need more than 8 receive buffers. This counter is always 0.
<code>rx_frame_too_long</code>	Received packets dropped when they are longer than 9216 bytes.
<code>rx_in_range_errors</code>	Received packets dropped when the Ethernet length field is not equal to the actual Ethernet data length.

Table 11: Ethtool -S Option Statistics (Continued)

Name	Description
rx_ip_checksum_errs, rx_tcp_checksum_errs, rx_udp_checksum_errs	Packets dropped due to TCP, IP, or UDP checksum errors.
rx_out_range_errors	Received packets dropped when their length field is ≥ 1501 bytes and ≤ 1535 bytes.
rx_pause_frames	The number of Ethernet pause frames (flow control) received.
rx_priority_pause_frames	The number of Ethernet priority pause frames (priority flow control) received per port.
rx_switched_unicast_packets, rx_switched_multicast_packets, rx_switched_broadcast_packets	The number of unicast, multicast, and broadcast packets switched internally.
rxpp_fifo_overflow_drop, rx_input_fifo_overflow_drop	The number of received packets dropped when a FIFO for descriptors going into the packet demux block overflows. In normal operation, this FIFO must never overflow.
tx_controlframes	The number of Ethernet control frames transmitted per port.
tx_dma_err	The number of errors occurred in the DMA operation associated with the transmit request from the host to the device.
tx_hdr_parse_err	The number of errors while parsing the packet header of a transmit request.
tx_internal_parity_err	The number of parity errors in the transmit request.
tx_pauseframes	The number of Ethernet pause frames (flow control) transmitted per port.
tx_priority_pauseframes	The number of Ethernet priority pause frames transmitted per port.
tx_qinq_err	The number of transmit requests with Q-in-Q style VLAN tagging, when such tagging is not expected on the outgoing interface.
tx_spoof_check_err	The number of spoof TX request failures, when MAC or VLAN spoof checking is enabled on the interface.
tx_tso_err	The number of transmit request errors, while performing TSO offload.

The following table contains a list of transmit/receive statistics per receive queue basis.

Table 12: Transmit/Receive Queue Statistics

Statistic	Description
rxq<x>:rx_bytes	The number of bytes received by the driver.
rxq<x>:rx_pkts	The number of packets received by the driver.
rxq<x>:rx_compl	The number of receive completions signaled to the driver by the ASIC.
rxq<x>:rx_mcast_pkts	The number of multicast packets received by the driver.
rxq<x>:rx_post_fail	The number of times the driver could not post received buffers to the ASIC.
rxq<x>:rx_drops_no_skbs	The number of times the driver could not allocate socket buffers.
rxq<x>:rx_drops_no_frags	Packets dropped due to insufficient buffers posted by the driver.
txq<x>:tx_compl	The number of transmit completions signaled by the ASIC.
txq<x>:tx_bytes	The number of bytes transmitted by the driver.
txq<x>:tx_pkts	The number of packets transmitted by the driver.
txq<x>:tx_reqs	The number of transmit request generated by the driver.
txq<x>:tx_stops	The number of times the driver requests the host to stop giving further transmit requests because the hardware transmit queue is filled up.
txq<x>:tx_drv_drops	The number of transmit packets dropped by the driver.

Appendix B: lpfc Driver BlockGuard Functionality

This appendix describes how to enable BlockGuard and set `lpfc` driver module parameters.


B.1 Overview

The BlockGuard feature checks the integrity of data read from and written to the host to the disk and back through the SAN. This check is implemented through the Data Integrity Field (DIF) defined in the ANSI T10 standard.

The Emulex lpfc driver supports T10 DIF Type 1. In the Type 1 implementation, the 8-byte DIF consists of a Ref Tag [or logical block address (LBA)], an App Tag, and a Guard Tag (or CRC). A Type 1 DIF is defined as having a 2-byte Guard Tag, a 2-byte App tag, and a 4-byte Ref tag, which consists of the lower 32 bits of the logical block address.

The following figure shows a data block (with a 512 byte sector) with the 8-byte footer attached to the end. The contents of the 8-byte footer are shown with the fields that make up the Type 1 DIF: the Guard Tag, the App Tag, and the Ref Tag. The App Tag is not used by the `lpfc` driver.

Figure 1: Data Block Showing Type 1 DIF



When data is written, the DIF is generated by the Host, or by the adapter, based on the block data and the logical block address. The DIF field is added to the end of each data block, and the data is sent through the SAN to the storage target. The storage target validates the CRC and Ref tag and, if correct, stores both the data block and DIF on the physical media. If the CRC does not match the data, then the data was corrupted during the write. A Check Condition is returned back to the host with the appropriate error code. The host records the error and retransmits the data to the target. In this way, data corruption is detected immediately on a write and never committed to the physical media. On a read, the DIF is returned along with the data block to the host, which validates the CRC and Ref tags. Because this validation is done by the hardware, it adds a very small amount of latency to the I/O.

The format of the Guard Tag can optionally be an IP Checksum instead of the CRC mandated by T10 DIF. This can be beneficial because the Initiator Host uses less CPU overhead to generate an IP Checksum than it does with a CRC. The IP Checksum is typically passed as the Guard Tag between the Initiator Host and the adapter. The adapter hardware will translate the IP Checksum into a CRC, or visa versa, on data being sent/received on the wire. The CRC is called a DIF protection type, and the IP Checksum is referred to as DIX protection type.

B.2 Enabling BlockGuard

BlockGuard is disabled by default. To enable it, the parameter `lpfc_enable_bg` must be passed to the driver as follows:

```
insmod lpfc.ko lpfc_enable_bg=1
```

For a permanent configuration that will persist across system reboots, create the `/etc/modprobe.d/lpfc.conf` file, and place the following line into it:

```
options lpfc lpfc_enable_bg=1
```

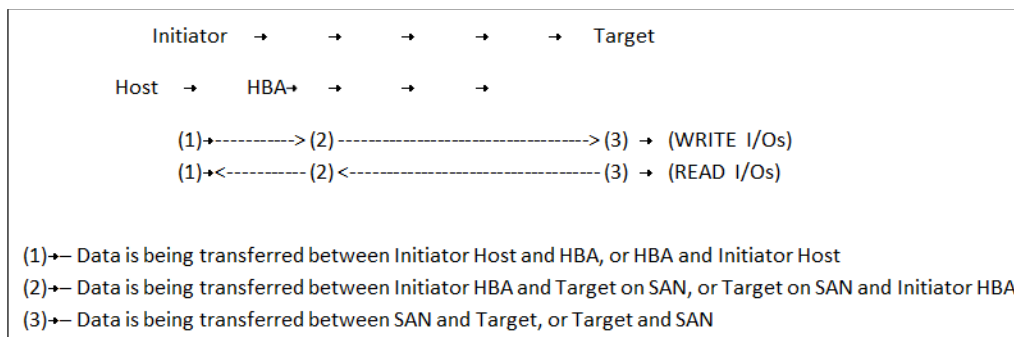
Additional module parameters can be added to this line, separated by spaces.

B.3 SCSI Command Support – SCSI Operation Codes

When there are both `READ` and `WRITE` requests, the command descriptor block (CDB) passed to the adapter from the Initiator Host has a read protect/write protect (`RDPROTECT`/`WRPROTECT`) field that indicates to the target whether to perform data integrity verification. It also indicates whether to transfer protection data between initiator and target. The adapter does not know if a target supports protection information or with which type of protection it is formatted. The Initiator Host, which has this knowledge, will always prepare a CDB with the appropriate `RDPROTECT`/`WRPROTECT` information, depending on target format and capabilities. The request will also include information about the protection type with which the target has been formatted.

In addition, the Initiator Host will also provide the adapter with an operation code that tells the controller how to place the protection data for the type of I/O to perform. Each I/O is logically a two-step process. The data is transferred between the Initiator Host and the adapter (over the PCI bus) and between the adapter and the target (over the SAN) as shown in the following figure. The type of operation defines whether the data transfer has protection data.

Figure 2: Data Transfer between Initiator Host and the Adapter



The initiator operations are listed in the following table.

Table 13: Initiator Operations

Initiator Operation	Initiator Host <-> Adapter	Adapter <-> Target	Comment
NORMAL	Unprotected	Unprotected	Used for unprotected I/O.
READ_INSERT	Protected	Unprotected	Reads the data from the target. The adapter then generates the protection data and transfers both data and protection data to the Initiator Host. No protection data is sent on the SAN. The adapter can insert the protection data guard tag as CRC or IP CSUM.
READ_PASS	Protected	Protected	Reads the data and protection data from the target on the SAN. The adapter will verify data integrity and transfer both data and protection data to the Initiator Host. The adapter can convert the protection data guard tag from CRC to IP CSUM.
READ_STRIP	Unprotected	Protected	Reads data and protection data from the target. The adapter will verify data integrity, discard protection data, and only transfers the data to the Initiator Host. It does not send the protection data to the Initiator Host. Protection data is only sent on the SAN.
WRITE_INSERT	Unprotected	Protected	Transfers the data from the Initiator Host. The adapter will then generate protection data and write both the data and protection data to the target. Protection data is only sent on the SAN.
WRITE_PASS	Protected	Protected	Transfers the data and protection data from the Initiator Host to the adapter. The adapter will verify protection data and write both data and protection data to the target on the SAN. The adapter can convert the protection data guard tag from IP CSUM to CRC.
WRITE_STRIP	Protected	Unprotected	Transfers data and protection data from the Initiator Host. The adapter will verify data integrity, discard protection data, and writes only the data to the target. No protection data is sent on the SAN.

B.4 lpfc Driver Module Parameters

The `lpfc` driver has two module parameters: `lpfc_prot_mask` and `lpfc_prot_guard`. Using these parameters, you can control which DIF capabilities the `lpfc` driver registers with the Linux SCSI subsystem. This, in turn, controls which initiator operations (BlockGuard profiles) are used during I/O operations. These parameters are set up when the driver loads and cannot be changed while the driver is running.

B.4.1 lpfc_prot_mask

This parameter controls the DIF operations that the driver registers with the operating system. The operating system selects an operation to use for each I/O command that matches the adapter DIF capability. The driver indicates its capabilities by the operations it registers with the operating system.

If the parameter is not passed to the driver, the default DIX_TYPE0 is used.

The SCSI layer will typically use the bit masks listed in the following table to determine how to place the protection data associated with I/Os to the SCSI Host. The default value for `lpfc_prot_mask` is to allow all of the options.

Table 14: `lpfc_prot_mask` Protection Types

Flag	Value	Indicates	Description
SHOST_DIF_TYPE1_PROTECTION	1	Adapter supports T10 DIF Type 1	Adapter to target Type 1 protection
SHOST_DIX_TYPE0_PROTECTION	8	Adapter supports DIX Type 0	Host to adapter protection only
SHOST_DIX_TYPE1_PROTECTION	16	Adapter supports DIX Type 1	Host to adapter Type 1 protection

The following table shows how protection data gets placed for each supported profile

Table 15: Protection Data Placement for Supported Profiles

Flag	Value	BlockGuard Profile	Operation
SHOST_DIF_TYPE1_PROTECTION	1	A1	READ_STRIP / WRITE_INSERT
SHOST_DIX_TYPE0_PROTECTION	8	AST2	READ_INSERT / WRITE_STRIP
SHOST_DIX_TYPE1_PROTECTION SHOST_DIF_TYPE1_PROTECTION	17	AST1 / C1	READ_PASS / WRITE_PASS

B.4.2 `lpfc_prot_guard`

This parameter correlates to the `SCSI_host_guard_type` of the Linux kernel. This parameter specifies the type of CRC the Linux operating system will pass to the `lpfc` driver. The following table shows the two guard types: CRC and IP-CSUM, with values of 0x1 and 0x2, respectively.

Table 16: `lpfc_prot_guard` Guard Types

Flag	Value	Indicates
SHOST_DIX_GUARD_CRC	1	Adapter supports T10 DIF CRC
SHOST_DIX_GUARD_IP	2	Adapter supports both T10 DIF CRC and IP-CSUM

The default value for `lpfc_prot_guard` is `SHOST_DIX_GUARD_IP`. This value defines the format for the guard tag when the data is transferred between the Host and the adapter. When data is transferred on the wire, the protection data guard tag is always translated into a T10 DIF CRC. To override the default, you can pass a module parameter value with either `insmod` or `modprobe`.

The SCSI layer will typically use an IP-CSUM as the method for computing the protection data guard tag because it uses less CPU overhead.

