



Emulex® Drivers for Linux

User Guide
Release 14.4

Copyright © 2003–2024 Broadcom. All Rights Reserved. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries. For more information, go to www.broadcom.com. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

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	5
1.1 Overview	5
1.2 Abbreviations	5
Chapter 2: Installing and Uninstalling	7
2.1 General Installation Requirements	7
2.2 Emulex FC Driver Kit	8
2.2.1 Installing the Emulex FC Driver Kit	8
2.2.2 Uninstalling the Emulex FC Driver Kit	9
2.3 Booting from a Nonzero LUN Attached to an Emulex FC Adapter	9
2.4 Emulex HBA Manager Application	9
2.5 Updating the Firmware	10
2.6 Installing the NVMe over FC Driver	11
2.6.1 Installing the NVMe over FC Driver on SLES Operating Systems	11
2.6.2 Installing the NVMe over FC Driver on RHEL Operating Systems	11
Chapter 3: Configuration	12
3.1 FC Driver Configuration	12
3.1.1 FC Driver Parameters	12
3.1.1.1 Static FC Driver Parameters	12
3.1.1.2 Dynamic FC Driver Parameters	15
3.1.2 Configuring FC Driver Parameters	18
3.1.2.1 Temporary Configuration with the modprobe Command	18
3.1.2.2 Persistent Configuration with the elx-lpfc.conf File	18
3.1.2.3 Configuring Parameters with the sysfs Command	19
3.1.3 Creating a New Ramdisk Image	20
3.1.4 Dynamically Recognizing LUNs and Targets (Using Scan)	20
3.1.5 Persistent Naming	20
3.1.5.1 Using udev to Discover Logical-to-Physical Mappings for sd Devices	20
3.1.5.2 Configuring the System to Boot Using Persistent Names	21
3.1.5.3 Using udev with st Devices	22
3.1.5.4 Persistent Naming References	24
3.1.6 ExpressLane Support	24
3.1.7 VMID Support	25
3.2 vPort Configuration	26
3.2.1 vPort Configuration Prerequisites	26
3.2.2 Creating, Deleting, and Displaying vPorts	26
3.2.2.1 Creating vPorts Using the mkvport.sh Script	26
3.2.2.2 Deleting vPorts Using the rmvport.sh Script	27

3.2.2.3 Displaying vPorts Using the <code>lsvport.sh</code> Script	27
3.2.3 vPort sysfs	27
3.2.3.1 vPort sysfs Tree	27
3.2.3.2 vPort sysfs Entries	28
3.2.4 Monitoring vPorts with <code>fc_vport</code>	30
3.2.5 vPort Configuration Limits	31
3.3 FC Driver Performance Tuning	31
3.3.1 Overview	31
3.3.1.1 <code>lpfc_hdw_queue</code>	32
3.3.1.2 <code>lpfc_irq_chann</code>	32
Chapter 4: Configuring NVMe over FC	33
4.1 Configuring an NVMe over FC Initiator Using Native CLI Commands	33
4.2 Emulex Support for NVMe CLI Commands	36
Chapter 5: Troubleshooting	37
5.1 Hardware Situations and Resolutions	37
5.2 FC Driver Situations and Resolutions	37
5.3 Firmware Update Error Messages	40
5.4 Dynamic Firmware Logging	40
5.5 NVMe Driver Situations and Resolutions	41
5.6 FC Driver Log Messages	42
5.6.1 Retrieving FC Driver Log Messages	42
5.6.2 FC Error Log Messages and Their Descriptions	43
Appendix A: FCP Driver BlockGuard Functionality	112
A.1 Overview	112
A.2 Enabling BlockGuard	113
A.3 SCSI Command Support – SCSI Operation Codes	113
A.4 FCP Driver Module Parameters	114
A.4.1 <code>lpfc_prot_mask</code>	115
A.4.2 <code>lpfc_prot_guard</code>	115
Appendix B: Native NVMe Multipathing	116
Appendix C: Trunking Support	123
Appendix D: FC-SP-2 Authentication (DH-CHAP) Support	125
Appendix E: Arm Support	126

Chapter 1: Introduction

1.1 Overview

This guide provides procedures for installing, updating, uninstalling, configuring, and troubleshooting several types of drivers for Linux that are supported by Emulex®.

This guide is applicable to several versions of drivers, operating systems, firmware, and adapters. For supported firmware versions and their latest release, go to www.broadcom.com.

This product supports the following Emulex FC HBAs:

- LPe38000-series adapters
- LPe37000-series adapters
- LPe36000-series adapters
- LPe35000-series adapters
- LPe32000-series adapters
- LPe31000-series adapters

For a list of Emulex adapters that are compatible with the stand-alone driver kits, go to www.broadcom.com.

1.2 Abbreviations

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

Table 1: Acronyms and Abbreviations

Acronym/Abbreviation	Description
ABTS	abort sequence
ADISC	address discovery
ANA	asymmetric namespace access
CDB	Command Descriptor Block
CQ	completion queue
CQE	control queue entry
CQID	completion queue identifier
CT	common transport
DIF	data integrity field
DM	Device Mapper
ELS	Extended Link Service
EQ	Event Queue
EQE	event queue entry
FAN	File Area Network
FCFI	Fibre Channel Forwarder Indicator
FCP	Fibre Channel Protocol
FDMI	Fabric-Device Management Interface
FLOGI	fabric login
FPIN	Fabric Performance Impact Notification

Table 1: Acronyms and Abbreviations (Continued)

Acronym/Abbreviation	Description
GnuPG	GNU Privacy Guard
I/O	input/output
IOCB	input/output control block
IOCBQ	input/output control block queue
IOPS	input/output operations per second
IOTAG	input/output tag
ISR	interrupt service routine
LBA	logical block address
LPFC	LightPulse® Fibre Channel
MQ	mailbox queue
MSI	message signaled interrupt
MSI-X	extended message signaled interrupt
NAA	Network Address Authority
NPIV	N_Port ID Virtualization
OOB	out-of-box
PLOGI	port login
RPI	remote port indicator
RPM	Red Hat package manager
RQ	Receive Queue
RSCN	registered state change notification
SGE	scatter-gather element
SGL	scatter-gather list
ULP	upper layer protocol
VPI	virtual path identifier
WQ	work queue
WQE	work queue entry

Chapter 2: Installing and Uninstalling

The Linux driver supports both FCP and NVMe over FC.

Emulex releases Linux binary RPMs that are digitally signed using the GnuPG standard. This allows 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 held only by Emulex. Instructions for creating the GnuPG public key file are located at the following URL:

<https://docs.broadcom.com/docs/1334566>

The GnuPG key includes a link to the SHA2 key.

2.1 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 www.broadcom.com for supported operating systems.
- Install the latest driver.
- Update to the latest firmware on the adapter.

NOTE: In some cases, a firmware update requires either a firmware reset or a basic PCIe reset, depending on the features available in the new firmware. A firmware reset is performed automatically if it is needed, regardless of the operating system. A basic PCIe reset is also performed automatically, but only on the following Linux operating systems:

- SLES 12 SP4 and later
- SLES 15 SP4 and later
- RHEL 8.8 and later
- RHEL 9.2 and later

In some cases, a full reboot is required to activate new firmware or to enable a new feature. In that case, a message similar to one of the following appears after the firmware download is complete:

Download successfully completed. Please reboot the system to activate new firmware.

Download completed. Some features require an optional reboot. Refer to the Adapter's Firmware and Boot Code Release Notes for details.

For a list of features that require a reboot to be enabled, refer to the adapter's firmware and boot code release notes.

NOTE: Out-of-box Linux driver builds are built on and are intended to be run on general availability (GA) kernels for their respective distributions. If you need an out-of-box Linux driver build for a subsequent distribution's errata kernel, you might need to recompile the driver from the Linux driver source file on the errata kernel.

2.2 Emulex FC Driver Kit

The Emulex FC driver kit contains the following items:

- A zipped tar file that includes the driver binary RPMs for a specific driver version and Linux distribution. On select driver kits, the zipped tar file might also include any of the following RPMs:
 - nvmefc-connect RPM
 - elx-lpfc-extras RPM
- **NOTE:** Use only officially released Linux distribution kernels. The binary RPM packages support only officially released Linux distribution kernels and do not support prerelease distribution kernels. Refer to the *Emulex Drivers for Linux Release Notes* for details of kernel versions that are supported by the driver.
- An installation script, `elx_lpfc_install.sh`, that installs the FC 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.

2.2.1 Installing the Emulex FC Driver Kit

NOTE: Installing a newer version of the FC driver kit automatically removes any previous versions of the FC driver kit.

After the operating system is installed, perform the following steps to install the Emulex FC driver:

1. Download the driver kit from www.broadcom.com for one of the following operating systems:
 - SLES 12 SP4 and later
 - SLES 15 SP4 and later
 - RHEL 8.8 and later
 - RHEL 9.2 and later
2. Log in to a terminal as `root`.
3. Untar the driver kit by typing the following command:
`tar -zvxf elx-lpfc-dd-<Linux_distribution_version>-<driver_version>-ds-n.tar.gz`
4. Change to the directory into which the driver kit was extracted by typing the following command:
`cd elx-lpfc-dd-<Linux_distribution_version>-<driver_version>-ds/`
 - To install only the Emulex FC driver kit:
 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 driver is loaded, and devices that are connected to the system are accessible.

- To install the Emulex FC driver kit with NVMe:

NOTE: For instructions on setting up boot from SAN on NVMe over FC in Linux, refer to the *Emulex Boot for the Fibre Channel and NVMe over FC Protocols User Guide*.

Ensure that `nvme-cli` is installed; otherwise the script `./elx_lpfc_install.sh -n` will fail.

To confirm whether `nvme-cli` is installed, run the `rpm -qa | grep nvme-cli` command.

Run the `elx_lpfc_install.sh` script with the `-n` option:

`./elx_lpfc_install.sh -n`

This installs the `lpfc` driver along with `nvmefc-connect` rpm.

After the `elx_lpfc_install.sh` script has finished running successfully, the NVMe over FC driver is installed.

5. Reboot the system to enable the newly added driver options in the ramdisk.
6. You can use the following command to verify whether the version of the newly installed LPFC driver is correct:
`cat /sys/module/lpfc/version`

2.2.2 Uninstalling the Emulex FC Driver Kit

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

To uninstall the Emulex FC driver, perform the following steps:

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`

NOTE: The `--uninstall` option also uninstalls `elx-lpfc-extras` and `nvmefc-connect`.

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

To configure SLES 12 SPx or SLES 15 SPx to boot from an FC-attached disk device other than a LUN 0 device, refer to the *Emulex Boot for the Fibre Channel and NVMe over FC Protocols*, which is available at www.broadcom.com.

2.4 Emulex HBA Manager Application

The Emulex HBA 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. For instructions on installing and using the Emulex HBA Manager application, refer to the *Emulex HBA Manager Application Command Line Interface User Guide*, which is available at www.broadcom.com.

2.5 Updating the Firmware

You can update the firmware on the adapter using any of the following options:

- Using the Emulex HBA Manager application. For information on how to update the firmware using the Emulex HBA Manager application, refer to the *Emulex HBA Manager Application Command Line Interface User Guide*.
- Using the `sysfs` command.

To update the firmware using the `sysfs` command, perform the following steps:

1. Before you update the firmware, procure the model name of the adapter.

```
cat /sys/class/scsi_host/<host_no>/modelname
```

2. Modify the firmware file name to the model name.

3. Copy the firmware file to `/lib/firmware` with the model name.

```
cp <firmware_file>.grp /lib/firmware/
```

4. Run the following command and note the value for `lpfc_log_verbose`. You must echo the same value to `lpfc_log_verbose` after firmware activation.

```
cat /sys/class/scsi_host/host<X>/lpfc_log_verbose
```

5. Set the driver parameter `lpfc_log_verbose=0x800`.

```
echo 0x800 > /sys/class/scsi_host/host<X>/lpfc_log_verbose
```

6. Download the firmware to the adapter.

```
echo 1 > /sys/class/scsi_host/<host_no>/lpfc_req_fw_upgrade
```

7. Check the `dmesg` log to determine the appropriate reset action.

8. Depending on the features available in the new firmware, the firmware update requires either a firmware reset or a PCI bus reset.

To perform a firmware reset, issue the following command:

```
echo fw_reset > /sys/class/scsi_host/host<X>/board_mode
```

To perform a PCI bus reset, perform the following steps:

- a. Run the following command and note the value for `lpfc_enable_hba_reset`. You must echo the same value to `lpfc_enable_hba_reset` after firmware activation.

```
cat /sys/class/scsi_host/host<X>/lpfc_enable_hba_reset
```

- b. Enable HBA reset on the adapter.

```
echo 2 > /sys/class/scsi_host/host<X>/lpfc_enable_hba_reset
```

- c. Set all ports of the adapter offline.

```
echo offline > /sys/class/scsi_host/host<X>/board_mode
```

NOTE: For HBAs with more than one port, run this command for each host (port) entry.

- d. Reset the PCI bus of the adapter.

```
echo pci_bus_reset > /sys/class/scsi_host/host<X>/board_mode
```

- e. Set all ports of the adapter online.

```
echo online > /sys/class/scsi_host/host<X>/board_mode
```

NOTE: For HBAs with more than one port, run this command for each host (port) entry.

9. Enter the value of `lpfc_enable_hba_reset` procured in [Step a](#).

```
echo <original_value> /sys/class/scsi_host/host<X>/lpfc_enable_hba_reset
```

10. Enter the value of `lpfc_log_verbose` procured in [Step 4](#).

```
echo <original_value> /sys/class/scsi_host/host<X>/lpfc_log_verbose
```

2.6 Installing the NVMe over FC Driver

This section describes how to install the NVMe OOB drivers.

NOTE: The NVMe over FC target drivers are meant for evaluation and testing during development only. They are not intended for production use.

2.6.1 Installing the NVMe over FC Driver on SLES Operating Systems

Install the SLES operating system on the computer following the instructions provided with the operating system. Installing the operating system automatically installs an inbox initiator and a target NVMe driver. For information about the target NVMe driver, see [Chapter 4, Configuring NVMe over FC](#).

Installation also allows you to install the `nvme-cli` utility, which is required on initiators that use NVMe disks. It supports initiation of new NVMe over fabric connections, and it functions as a general utility that allows you to query and manipulate an NVMe device.

During the operating system installation, configure the following items:

- Select the system role of **Default System**.
- In **Installation Settings**, select **NVMe CLI** to install the option.

2.6.2 Installing the NVMe over FC Driver on RHEL Operating Systems

Install the RHEL operating system on the computer following the instructions provided with the operating system. Installing the operating system automatically installs an inbox initiator NVMe driver. At this time, Emulex supports NVMe over FC on RHEL 8.8 and later and RHEL 9.2 and later.

During the installation, select the system role of **Default System**.

After the operating system is installed, extract the `nvme-cli` utility from the operating system ISO image. The `nvme-cli` utility supports initiation of new NVMe over fabric connections, and it functions as a general utility that allows you to query and manipulate an NVMe device. To install the `nvme-cli` utility, type this command:

```
rpm -ivh nvme-cli-x.x.x.rpm
```

You must update to the latest RHEL kernel and the latest RHEL release of the NVMe CLI package.

Chapter 3: Configuration

3.1 FC Driver Configuration

This section describes how to configure parameters for the FC driver.

3.1.1 FC Driver Parameters

The FC 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 changes to take effect. Changes to most dynamic parameters take effect immediately; some do not take effect until there is a link-down and link-up sequence.

For more information on driver parameters, see [Section 3.1.1.1, Static FC Driver Parameters](#), and [Section 3.1.1.2, Dynamic FC Driver Parameters](#).

3.1.1.1 Static FC Driver Parameters

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

Table 2: Static FC Driver Parameters

Parameter	Description	sysfs Visible	Activation
lpfc_discovery_threads	<p>Specifies the maximum number of ELS commands that can be outstanding for discovery.</p> <p>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.</p> <p>The minimum value is 1. The maximum value is 64. The default value is 32.</p>	No	Driver reload
lpfc_enable_bg	<p>Enables the BlockGuard (T10-DIF) feature. See Appendix A for more information about the BlockGuard feature.</p> <p>The minimum value is 0 (default). The maximum value is 1.</p>	Yes	Driver reload
lpfc_enable_da_id	<p>When enabled, the FC driver issues a <code>DA_ID CT</code> command to the fabric when vPorts log out of the fabric. The enabled value is 1 (default). The disabled value is 0.</p>	Yes	Driver reload
lpfc_enable_fc4_type	<p>When enabled, defines the FC4 types that are supported.</p> <p>The possible values are:</p> <ul style="list-style-type: none"> ■ 1 = Enables just FCP. ■ 3 = Enables both FCP and NVMe (default). <p>The default value is 3.</p>	Yes	Driver reload
lpfc_enable_hba_heartbeat	<p>When enabled, the heartbeat logic in the FC driver can detect whether the adapter is functional. If the heartbeat logic detects that the adapter is not functional, the driver shuts down the adapter.</p> <p>The disabled value is 0 (default). The enabled value is 1.</p>	Yes	Driver reload

Table 2: Static FC Driver Parameters (Continued)

Parameter	Description	sysfs Visible	Activation
lpfc_enable_hba_reset	When enabled, the FC 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	Driver reload
lpfc_enable_npiv	When enabled, the FC driver can use NPIV to create vPorts (if supported by the fabric). The enabled value is 1 (default). The disabled value is 0.	Yes	Driver reload
lpfc_enable_nvmet	Specifies the WWPN of the ports enabled for an NVMe target.	No	Driver reload
lpfc_enable_SmartsAN	Sets up FDMI support for HPE Smart SAN. The disabled value is 0 (default). The enabled value is 1.	Yes	Driver reload
lpfc_EnableXLane	On LPe31000-series and LPe32000-series adapters, specifies whether the ExpressLane™ feature is enabled or disabled. See Section 3.1.6, ExpressLane Support , for more information about the ExpressLane feature. The minimum value is 0 (default). The maximum value is 1. NOTE: If the lpfc_vmid_priority_tagging parameter is enabled, EnableXLane is disabled because both features conflict with regard to using the CS_CTL field in the FC header. NOTE: This parameter is not supported on LPe35000-series and LPe36000-series adapters.	Yes	Driver reload
lpfc_fcp_class	Specifies whether FC Class 2 or FC Class 3 is used for FCP data transmission. For Class 2, the value is 2. For Class 3, the value is 3 (default).	Yes	Driver reload
lpfc_fcp_wait_abts_rsp	Enables FCP completion to wait for an abort sequence response. The enabled value is 1. The disabled value is 0 (default).	Yes	Driver reload
lpfc_fDMI_on	Specifies whether FDMI support is enabled or disabled. The enabled value is 1 (default). The disabled value is 0.	Yes	Driver reload
lpfc_hba_queue_depth	Limits the number of outstanding commands per HBA. The supported range is 32 to 8192. NOTE: The maximum (default) value will be limited by the lesser of the number of XRLs for I/O and lpfc_hba_queue_depth.	Yes	Driver reload
lpfc_hw_queue	Defines the number of I/O channels supported by the driver. The default value is 0. A value of 0 means that the driver will automatically determine a recommended number of I/O channels to allocate based on system resources. NOTE: Nonzero values can be ignored when lpfc_irq_chann is set to its default value of 0. For more information, see the description of the lpfc_irq_chann parameter. The supported range is 0 to 256.	Yes	Driver reload
lpfc_irq_chann	Defines the number of EQs (and MSI-X vectors) requested. The default value is 0. A value of 0 means that the driver automatically determines a recommended number of EQs, MSI-X vectors, and hardware queues to allocate based on system resources. This value is limited by the lesser of the number of CPUs and hardware queues. NOTE: Depending on system resources, a default value of 0 may overwrite any values set for lpfc_hw_queue. This is to ensure proper driver recommended settings between hardware queues, EQs, and MSI-X vectors.	Yes	Driver reload

Table 2: Static FC Driver Parameters (Continued)

Parameter	Description	sysfs Visible	Activation
lpfc_lun_queue_depth	<p>Specifies the default maximum number of commands sent to a single logical unit (disk drive).</p> <p>The minimum value is 1. The maximum value is 512. The default value is 64.</p> <p>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.</p> <p>NOTE: This parameter can be dynamically changed using the Emulex HBA Manager application.</p>	Yes	Driver reload
lpfc_max_luns	<p>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.</p> <p>The minimum value is 0. The maximum value is 65535. The default value is 255.</p>	Yes	Driver reload
lpfc_max_scsicmpl_time	<p>Uses the specified command completion time to control queue depth. The units are in milliseconds.</p> <p>The minimum value is 0 (default). The maximum value is 6000.</p>	Yes	Driver reload
lpfc_max_vmid	<p>Indicates the number of VMIDs supported.</p> <p>Maximum number of VMs to be tagged.</p> <p>The minimum value is 4. The maximum value is 255. The default value is 8.</p>	Yes	Driver reload
lpfc_nvmet_mrq	Configures the allocation of multireceive queues (NVMe target mode only). The default value is 0 (the driver sets the appropriate value). The maximum value is 16.	Yes	Driver reload
lpfc_ras_fwlog_func	<p>Configures the function on which firmware logging support is enabled.</p> <p>The minimum value is 0 (default). The maximum value is 7.</p> <p>The default is logging enabled on Function 0.</p>	Yes	Driver reload
lpfc_restrict_login	<p>When enabled, this parameter restricts the vPorts login to remote initiators.</p> <p>The enabled value is 1 (default). The disabled value is 0.</p>	No	Driver reload
lpfc_scan_down	<p>When enabled, this parameter selects the <i>scan down</i> method (scanning the AL_PA from high to low) to assign a SCSI ID.</p> <p>The enabled value is 1 (default). The disabled value is 0.</p>	Yes	Driver reload
lpfc_sg_seg_cnt	<p>Controls the scatter-gather maximum segment count passed to the FC driver. This variable is applicable per SCSI command.</p> <p>The minimum value is 64 (default), and the maximum value is 510.</p>	Yes	Driver reload
lpfc_use_msi	<p>When enabled, this parameter 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	Driver reload
lpfc_vmid_app_header	<p>VMID application ID tagging.</p> <p>The minimum value is 0 (disabled). The maximum value is 1. The default value is 0.</p> <p>Brocade VMID uses the application service header field to communicate the capability.</p>	Yes	Driver reload

Table 2: Static FC Driver Parameters (Continued)

Parameter	Description	sysfs Visible	Activation
lpfc_vmid_inactivity_timeout	Inactivity timeout duration in hours. The minimum value is 0. The maximum value is 24. The default value is 4. VMID is an ID assigned per VM. The VMID is removed when a VM is inactive for lpfc_vmid_inactivity_timeout duration. The lpfc_vmid_inactivity_timeout value is in hexadecimal.	Yes	Driver reload
lpfc_vmid_priority_tagging	VMID CS_CTL tagging. Valid values are: <ul style="list-style-type: none">■ 0 = Disables priority tagging. Priority tagging is also disabled when the parameter is set to anything other than 1 or 2.■ 1 = Enables priority tagging only for targets that support it in their PLOGI LS_ACC response.■ 2 = Enables priority tagging for all targets, whether they support it in their PLOGI LS_ACC response. The minimum value is 0. The maximum value is 2. The default value is 0. Cisco VMID uses the priority tagging field to communicate the capability.	Yes	Driver reload

3.1.1.2 Dynamic FC Driver Parameters

Changes to most dynamic parameters take effect immediately; some do not take effect until a port reset occurs. All lpfc dynamic parameters are read and written using sysfs. The following table lists the dynamic FC driver parameters.

Table 3: Dynamic FC Driver Parameters

Parameter	Description	sysfs Visible	Activation
lpfc_devloss_tmo	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	Dynamic
lpfc_enable_auth	Enables or disables DH-CHAP authentication. The possible values are: 0 = Authentication is disabled (default). 1 = Authentication is enabled. NOTE: DH-CHAP and Dynamic D_Port are mutually exclusive. If both features are enabled simultaneously, warning messages result. You must either disable DH-CHAP using the lpfc_enable_auth parameter or disable Dynamic D_Port using the Emulex HBA Manager application.	Yes	Dynamic
lpfc_enable_bbcr	Specifies whether the Buffer-to-Buffer Credit Recovery feature is enabled or disabled. The minimum value is 0. The maximum value is 1 (default).	Yes	Port reset
lpfc_enable_mds_diags	Enables or disables support for Cisco MDS switch diagnostics. The minimum value is 0 (default). The maximum value is 1. NOTE: The parameter must be disabled (set to 0) after diagnostics are complete.	Yes	Port reset
lpfc_fcp_imax	Specifies the maximum number of fast-path FCP interrupts per second. The default value is 0. If set to 0, the parameter is OFF. The minimum value is 5000. The maximum value is 5,000,000. A value of 0 (default) provides dynamic interrupt coalescing that adapts to provide low latency or high IOPs according to I/O load.	Yes	Dynamic

Table 3: Dynamic FC Driver Parameters (Continued)

Parameter	Description	sysfs Visible	Activation
lpfc_fcp_io_sched	<p>Determines which algorithm to use when scheduling an FC I/O to an I/O channel. For more information, see Section 3.3, FC Driver Performance Tuning.</p> <p>The value 0 configures <i>schedule by hardware queue</i> (rather than round robin). A value of 1 (default) configures <i>schedule by CPU</i>.</p>	Yes	Dynamic
lpfc_first_burst_size	Configures the first burst size on targets that support the first burst. The minimum value is 0 (default). The maximum value is 65,536 bytes.	Yes	Dynamic
lpfc_link_speed	<p>Specifies the FC link speed. The possible values are:</p> <ul style="list-style-type: none"> ■ 0 = Auto-select (default) ■ 16 = 16Gb/s ■ 32 = 32Gb/s ■ 64 = 64Gb/s <p>NOTE:</p> <ul style="list-style-type: none"> ■ Setting this option incorrectly can cause the adapter to fail to initialize. ■ 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 x86 BootBIOS or UEFI utility. 	Yes	Port reset
lpfc_log_verbose	<p>Specifies the log verbosity level of the messages posted by the driver. When enabled, this parameter prints extra log messages for debugging purposes.</p> <p>The minimum value is 0x0 (default). The maximum value is 0xFFFFFFFF.</p> <p>NOTE: Setting the lpfc_log_verbose value to 0xffffffff might degrade system performance.</p>	Yes	Dynamic
lpfc_nvme_embed_cmd	<p>This is an optional parameter.</p> <p>On LPe31000-series and LPe32000-series adapters, when enabled, this parameter embeds an NVMe command in a WQE.</p> <p>The possible values are:</p> <ul style="list-style-type: none"> 0 = Embeds an NVMe command in an SGL. 1 = Embeds an NVMe command in a WQE (except for LPe35000-series and LPe36000-series adapters). 2 = Embeds an NVMe command in a WQE (force). <p>NOTE: This parameter is not supported on LPe35000-series and LPe36000-series adapters.</p>	Yes	Dynamic
lpfc_nvme_enable_fb	<p>Enables NVMe first burst on initiator and target functions.</p> <p>For an initiator, enabling this parameter means that an NVMET PRLI response with FBA enabled and an FB_SIZE set to a nonzero value are processed by the initiator for subsequent NVMe FCP I/O.</p> <p>Currently, this feature is not supported on NVMe targets.</p> <p>The possible values are:</p> <ul style="list-style-type: none"> 0 = Disabled (default) 1 = Enabled 	Yes	Dynamic
lpfc_nvme_oas	<p>When enabled, uses the oas bit when sending NVMe/NVMe target I/Os.</p> <p>The possible values are:</p> <ul style="list-style-type: none"> 0 = Disabled (default) 1 = Enabled 	Yes	Dynamic

Table 3: Dynamic FC Driver Parameters (Continued)

Parameter	Description	sysfs Visible	Activation
lpfc_nvmet_fb_size	When enabled, communicates the first burst size of the target to the NVMe initiator in the PRLI response. <code>lpfc_nvmet_fb_size</code> is used only when the <code>lpfc_nvme_enable_fb</code> is set and the target sends a PRLI response. The possible values are: 0 = Disabled (default) 1 = Enabled This parameter is supported only on physical ports.	Yes	Dynamic
lpfc_ras_fwlog_bufsize	Specifies the buffer size to be allocated in host memory for firmware logging. The possible values are: <ul style="list-style-type: none">■ 0 = Disables the firmware log written to the host memory (default).■ 1 = ¼ MB (256K) host memory to be allocated for the firmware log.■ 2 = ½ MB (512K) host memory to be allocated for the firmware log.■ 3 = ¾ MB (768K) host memory to be allocated for the firmware log.■ 4 = 1 MB host memory to be allocated for the firmware log.	Yes	Driver reload
lpfc_ras_fwlog_level	Configures the firmware logging level. Valid only if firmware logging is enabled using the <code>lpfc_ras_fwlog_bufsize</code> parameter. The minimum value is 0 (default). The maximum value is 4.	Yes	Driver reload
lpfc_req_fw_upgrade	Set this parameter to 1 to update the firmware.	Yes	Dynamic
lpfc_suppress_rsp	Negotiates whether to suppress the RSP from a target during a PLOGI. The minimum value is 0. The maximum value is 1. The default value is 1.	Yes	Dynamic
lpfc_tgt_queue_depth	Limits the number of outstanding commands per target port. The minimum value is 10. The maximum value is 65,535 (default). NOTE: The maximum value is determined by the adapter's capability.	Yes	Dynamic
lpfc_throttle_log_cnt	Specifies the number of messages logged within <code>throttle_log_time</code> . The minimum value is 1. The maximum value is 1000. The default value is 10.	Yes	Dynamic
lpfc_throttle_log_time	Specifies the time limit in seconds for <code>throttle_log_cnt</code> . The minimum value is 1. The maximum value is 60. The default value is 1.	Yes	Dynamic
lpfc_use_adisc	When enabled, an address discovery is sent instead of a PLOGI for device discovery or RSCN. The enabled value is 1. The default value is 0 (disabled).	Yes	Dynamic
lpfc_XLanePriority	For Emulex 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. NOTE: This parameter is not supported on LPe35000-series and LPe36000-series adapters.	Yes	Dynamic

3.1.2 Configuring FC Driver Parameters

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

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

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

- The Emulex HBA Manager application (refer to the *Emulex HBA Manager Application Command Line Interface User Guide* for more information)

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

3.1.2.1 Temporary Configuration with the modprobe Command

When you manually load the FC 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 driver is unloaded.

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

For example, if you want to temporarily set `lpfc_lun_queue_depth` to 20 (the default is 64) for all HBAs in your system, load the FC driver with the following command:

```
modprobe lpfc lpfc_lun_queue_depth=20
```

3.1.2.2 Persistent Configuration with the elx-lpfc.conf File

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

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


```
options lpfc lpfc_log_verbose=0x3FFF
```

If driver parameters are added to the `elx-lpfc.conf` file, the FC 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 3.1.3, Creating a New Ramdisk Image](#).

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

3.1.2.3 Configuring Parameters with the sysfs Command

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

NOTE: The following information applies to `sysfs`:

- `sysfs` changes exist only during driver load and are lost when the FC 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 affects only host 7.
modprobe lpfc lpfc_log_verbose=0x7F applies to all SCSI hosts (ports) managed by the FC 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 driver parameters. All `scsi_hosts` bound to the FC driver can be viewed with the following command:

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

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

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

Sample 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 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
```

3.1.3 Creating a New Ramdisk Image

The `elx_lpfc_install.sh` script creates a ramdisk image that contains the FC driver for the currently running kernel.

NOTE: You must create a new ramdisk image whenever the `lpfc` options in `/etc/modprobe.d/elx-lpfc.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 drivers and installed binary RPM FC driver kits, type the following command:

```
dracut -f /boot/initramfs-<kernel_version>.img <kernel_version>
```

3.1.4 Dynamically Recognizing LUNs and Targets (Using Scan)

The FC driver enables you to dynamically recognize LUNs and targets without unloading or reloading the FC 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 the following command:

```
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 the following command:

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

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

3.1.5 Persistent Naming

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

3.1.5.1 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 `sdc`, type the following command:

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

Sample output follows:

```
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
```

3.1.5.2 Configuring the System to Boot Using Persistent Names

For SLES Operating Systems:

NOTE: SLES operating systems 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, perform the following steps:

1. Do one of the following:
 - For x86 boot:
In the `/boot/grub/menu.lst` file, find the kernel line for the default boot. For example:
`kernel /boot/vmlinuz root=/dev/sda2 vga=0x314`
 - For UEFI boot:
In the `/boot/efi/EFI/sles/grub.cfg` file, 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 3.1.5.1, Using udev to Discover Logical-to-Physical Mappings for sd Devices](#).
3. In the file you updated for x86 boot or UEFI boot, 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 Operating Systems:

To use a persistent name for a boot device, perform the following steps:

1. Do one of the following:
 - For x86 boot:
In the `/boot/grub/grub.conf` file, find the kernel line for the default boot. For example:
`kernel /boot/vmlinuz -<kernel version> ro root=/dev/sda2`
 - For UEFI boot:
In the `/boot/efi/EFI/redhat/grub.cfg` file, 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 3.1.5.1, Using udev to Discover Logical-to-Physical Mappings for sd Devices](#).
3. In the file you updated for x86 boot or UEFI boot, 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.

3.1.5.3 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, the WWPN is unique and can be used to create the persistent name. In this case, the `scsi_id` command should return the WWPN 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 approved. 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 are 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: Because 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 `[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 an FC controller with multiple SCSI tape devices behind it (FC-SCSI tape devices). 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 a udev Persistent Name for a 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 follows:

```

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 follows:

```

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

```

3.1.5.4 Persistent Naming References

See the following reference by Daniel Drake (dsd) for more information on persistent naming:
www.reactivated.net/udevrules.php

3.1.6 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,

NOTE: The ExpressLane feature is supported only on Emulex LPe31000-series and LPe32000-series adapters.

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

An ExpressLane LUN must be enabled in the driver before it can be used by the Emulex HBA Manager application. Additionally, the priority value to mark each of the ExpressLane frames must be specified to the FC driver through the appropriate driver parameters. For additional information, refer to the *Emulex HBA Manager Application Command Line Interface User Guide*.

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

NOTE: On inbox drivers, the ExpressLane configuration is not persistent. To make it persistent you must install the `elx-lpfc-extras` RPM (`elx-lpfc-vector-map` in case of older distributions) from the out-of-box `lpfc` driver package. The `lpfc_configure_oas.sh` script that is a part of this RPM takes care of making the ExpressLane configuration persistent.

3.1.7 VMID Support

You can view the VMID information only if VMID support has been enabled in the switch and the driver parameters are set up to enable VMID on the ESXi host. VMID support is enabled on the ESXi host using the `lpfc_vmid_app_header` parameter for a Brocade switch and the `lpfc_vmid_priority_tagging` parameter for a Cisco switch.

For a Cisco Switch:

You can view the VMID of the VM whose LUN you mapped using the following command:

```
cat /sys/class/scsi_host/host<X>/lpfc_vmid_info
```

Output similar to the following is shown:

```
Key 'vmid':
lpfc VMID Page: ON
VMID priority ranges:
[x97 - x97], qos: x0
[x98 - xff], qos: x0
VEM ID: 10:00:00:90:fa:5e:68:9c:20:00:00:90:fa:5e:68:9c
ID00 READs:2ff1 WRITEs:0
UUID:
35 34 33 36 61 62 39 30 38 31 66 61 63 31 32 30
String (5436ab9081fac120)
CS_CTL VMID: 151
Last Access Time :2022-9-28T06:33:42
```

For a Brocade Switch :

You can view the VMID of the VM whose LUN you mapped using the following command:

```
cat /sys/class/scsi_host/host<Y>/lpfc_vmid_info
Output similar to the following is shown:
```

```
Key 'vmid':
lpfc VMID Page: ON
ID00 READs:f8 WRITEs:0
UUID:
35 34 33 36 61 62 39 30 38 31 66 61 63 31 32 30
String (5436ab9081fac120)
Application id: 521
Last Access Time :2022-9-27T06:53:04
```

3.2 vPort Configuration

This section describes how to create, delete, and display vPorts. You can also configure the vPorts using the Emulex HBA Manager application. For more information, refer to the *Emulex HBA Manager Application Command Line Interface User Guide*.

3.2.1 vPort Configuration Prerequisites

Before configuring vPorts, consider the following points:

- Ensure that you are using the latest firmware for vPort functionality. For the latest firmware, go to www.broadcom.com.

NOTE: Before performing a firmware update, driver installation is required. For more information on installing the driver, see [Chapter 2, 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.

3.2.2 Creating, Deleting, and Displaying vPorts

vPorts are created through `sysfs` entries that are presented in the physical port's `sysfs` directory. Scripts that create, delete, and display vPorts are provided in the `/usr/sbin/lpfc` directory and are part of the Emulex HBA 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 during discovery, set the `lpfc_devloss_tmo` parameter to 60 seconds when NPIV is enabled.

3.2.2.1 Creating vPorts Using the `mkvport.sh` Script

You can use the `mkvport.sh` 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 a port name of 10000000c94ac63a and a node name of 20010000c94ac63a on the physical port with `scsi_host` name `host7`, type:

```
./mkvport.sh host7 10000000c94ac63a 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.

3.2.2.2 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.sh` 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_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 the following command:

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

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

3.2.2.3 Displaying vPorts Using the `lsvport.sh` Script

You can use the `lsvport.sh` 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.

3.2.3 vPort sysfs

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

3.2.3.1 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 trees:

```
/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, a `hostY` directory 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 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
```

3.2.3.2 vPort sysfs Entries

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

Table 4: vPort sysfs Entries

vPort sysfs 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) logs in only to N_Ports that are physically located on a different port. The port still attempts 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 N_Ports, even if they are physically located on the same port.</p> <p>NOTE: This parameter was created to reduce the number 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>

Table 4: vPort sysfs Entries (Continued)

vPort sysfs Entries	Type	Range/Input	Location and Description
lpfc_restrict_login	Read/Write	0 = Off 1 = On (default)	<p>/sys/class/scsi_host/hostX/</p> <p>lpfc_restrict_login</p> <p>(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 in to 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 sysfs entry is turned off, the driver attempts to log in to every port that it can access in the SAN, and it accepts logins from all ports.</p> <p>NOTE: This parameter was created to reduce the number 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>
max_npiv_vports	Read-only	integers	<p>/sys/class/fc_host/hostX/</p> <p>max_npiv_vports</p> <p>This entry displays the maximum number of vPorts that are supported by the underlying fc_host hardware.</p> <p>This sysfs entry exists only if the sysfs entries for vport_create and vport_delete exist. If an fc_host does not support NPIV, this sysfs entry might not exist.</p> <p>Use this sysfs entry with the npiv_vports_inuse entry to determine whether the maximum number of vPorts have been created on this fc_host.</p>
node_name	Read-only	16-byte hexadecimal value	<p>/sys/class/fc_host/hostX/node_name</p> <p>/sys/class/fc_vports/vport-X:0-Z/</p> <p>node_name</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>/sys/class/fc_host/hostX/npiv_vports_inuse</p> <p>This entry displays the number of vPorts that were created on this fc_host.</p> <p>This sysfs entry exists only if the sysfs entries for vport_create and vport_delete exist. If an fc_host does not support NPIV, this sysfs entry might not exist.</p> <p>Use this sysfs entry with max_npiv_vports to determine whether the maximum number of vPorts have been created on this fc_host.</p>

Table 4: vPort sysfs Entries (Continued)

vPort sysfs Entries	Type	Range/Input	Location and Description
port_name	Read-only	16-byte hexadecimal value	/sys/class/fc_host/hostX/port_name /sys/class/fc_vports/vport-X:0-Z/port_name This entry displays the physical or vPort port name. You assign this value when the vPort is created, and it is transmitted to the fabric upon fabric login.
vport_create	Write-only	WWPN:WWNN	/sys/class/fc_host/hostX/vport_create This entry creates a vPort on the physical port on which hostX is located. The new vPort has a WWPN and WWNN present on the fabric based on the WWPN and WWNN that are entered with this sysfs entry. This entry returns a 0 if the vPort creation was successful. A nonzero value indicates that the vPort was not created. If an fc_host does not support NPIV, this sysfs entry might not exist. NOTE: It is possible for the vPort creation to succeed, fail, or be in an inoperative state. Use the new sysfs tree created by the new vPort to check the state of the new vPort.
vport_delete	Write-only	WWPN:WWNN	/sys/class/fc_host/hostX/vport_delete This entry deletes a vPort on the physical port on which hostX is located. The vPort that matches the WWPN and WWNN is immediately deleted. This entry returns a 0 if the vPort deletion was successful. A nonzero value indicates that the vPort was not deleted. If an fc_host does not support NPIV, this sysfs entry might not exist. 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.

3.2.4 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 appears 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`).

3.2.5 vPort Configuration Limits

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

NOTE: 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 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.

- 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 a 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 too many vPorts exist 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.

3.3 FC Driver Performance Tuning

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

3.3.1 Overview

The configurable parameters `lpfc_hw_queue` and `lpfc_irq_chann` can enhance performance on supported RHEL and SLES operating systems. These features are available through module parameters that are defined in the FC driver as well as `sysfs` entries defined by the Linux kernel.

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

3.3.1.1 `lpfc_hdw_queue`

The `lpfc_hdw_queue` module parameter can be configured at driver load time. It defines the number of hardware queues supported by the driver for each port. The driver is capable of supporting parallel I/O paths, and each I/O path is capable of posting and completing FCP and NVMe commands independent of the other. Each hardware queue is composed of a unique pair of completion queue and work queue.

By default, `lpfc_hdw_queue` is configured for an automatically determined recommended amount based on system resources. The driver also limits the number of hardware queues to not exceed the number of online logical CPUs (as reported by `/proc/cpuinfo`). It is highly desirable, for performance, to have one hardware queue per CPU.

3.3.1.2 `lpfc_irq_chann`

The `lpfc_irq_chann` module parameter can be configured at driver load time. It defines the number of MSI-X vectors supported by the driver for each port. The driver is capable of supporting parallel I/O paths, and each I/O path is capable of posting and completing FCP and NVMe commands independent of the other.

Each IRQ channel is composed of a unique MSI-X vector and an event queue.

By default, `lpfc_irq_chann` is configured for an automatically determined recommended amount based on system resources. The driver also limits `lpfc_irq_chann` to not exceed the number of online logical CPUs (as reported by `/proc/cpuinfo`) and to not exceed `lpfc_hdw_queue`. If `lpfc_irq_chann` is less than `lpfc_hdw_queue`, multiple hardware queues must share the same IRQ channel. It is highly desirable, for performance, to have one IRQ channel per hardware queue.

Chapter 4: Configuring NVMe over FC

NOTE: The NVMe over FC target drivers are meant for evaluation and testing during development only. They are not intended for production use.

NOTE: Creation of NPIV connections on initiator ports that are configured for NVMe over FC is not supported. However, initiator ports can connect to FCP and NVMe targets simultaneously.

To configure NVMe over FC, you must perform the following procedures in the indicated sequence:

1. Attach NVMe over FC-based storage to the FC SAN. Alternatively, an Emulex adapter can be configured as an NVMe over FC target and used by the Linux NVMeT subsystem (see [Appendix B, Configuring NVMe over FC on a Target](#)).
2. Configure NVMe over FC on initiator systems, as described in this chapter.
3. Verify that the operating system has the dependency `nvme-cli` RPM installed to enable the NVMe feature.

NOTE: The following information applies to configuring NVMe over FC:

- The `nvmefc-connect` RPM is not installed by default. To install the RPM, the `elx_lpfcc_install.sh` script includes a `-n` option. When the `-n` option is used, the install script installs the driver RPM and the `nvmefc-connect` RPM. The `nvmefc-connect` RPM is not needed on the latest releases of SLES operating systems such as SLES 12 SP5 and SLES 15 variants. It is also not needed on the latest releases of RHEL operating systems such as RHEL 8.8 and later.
- All operating system distributions will be releasing the upstream version of the NVMe over fabrics auto-connect scripts, which will replace this `nvmefc-connect` RPM.
- Before configuring NVMe over FC using native NVMe CLI commands, ensure that you have installed the latest Emulex firmware for the adapters.
- The target system must be configured and running before you connect the host system to the target.

4.1 Configuring an NVMe over FC Initiator Using Native CLI Commands

NVMe over FC on an initiator is configured automatically with the `lpfc` installer script `elx_lpfcc_install.sh`.

To install the `lpfc` driver with NVMe on the initiator, perform the following steps:

1. Type the following command:

```
# elx_lpfcc_install.sh -n
-n installs the nvmefc-connect RPM if the running operating system is supported.
```

NOTE: The following information applies to configuring NVMe over FC on an initiator system:

- The `nvmefc-connect` RPM is not needed on the latest releases of SLES operating systems such as SLES 12 SP5 and SLES 15 variants. It is also not needed on the latest releases of RHEL operating systems such as RHEL 8.8 and later.
- The existing `lpfc` driver parameters in the `elx-lpfcc.conf` file are not affected.

2. By default, the SCSI and NVMe protocols are enabled on the port; and, by default, the port is an initiator for the SCSI/ NVMe protocol. Reboot the system.

After the system is rebooted, the configured protocols are enacted on the FC ports.

The `nvmefc-connect.rpm` file, which is installed with the OOB Linux driver, automatically causes the initiator to connect to the target devices. Connections are established whenever any of the following events occur, as long as the target is running:

- A host system reboot
- A linkup event on the initiator port
- A fabric zone event

After configuring the initiator, you can run the `nvme_info` command to verify that your configuration is working by typing the following command on the initiator:

```
cat /sys/class/scsi_host/host<X>/nvme_info
```

where `<X>` is the host number.

Information similar to the following is displayed:

```
NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc0 WWPN x10000090fa9429e4 WWNN x20000090fa9429e4 DID x011d00 ONLINE
NVME RPORT      WWPN x10000090fa942f0c WWNN x20000090fa942f0c DID x010a00 TARGET DISCSRVC ONLINE

NVME Statistics
LS: Xmt 0000000348 Cmpl 0000000348 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 00000000000019cb Issue 00000000000019cf OutIO 0000000000000004
    abort 00000000 noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr 00000000 err 00000000
FCP CMPL: xb 00000010 Err 00000015
```

3. If multipathing is used in SLES 15 SP4 and later operating systems, skip to [Step 4](#). For operating systems with a single path, use the `nvme list` command to display target connections by typing the following command on the initiator:

```
nvme list
```

NOTE: There are two similar commands to list all NVMe namespace devices:

- `nvme list` (with a space) is a native NVMe CLI command.
- `nvme-list-linux` (with a hyphen) is an Emulex HBA Manager CLI command.

This chapter refers to the `nvme list` command.

Information similar to the following is displayed:

Node	SN	Model	Namespace	Usage	Format	FW	Rev
/dev/nvme0n1	981fad118adb0fa	Linux	1	268.44 GB / 268.44 GB	512 B + 0 B	4.4.70-2	
/dev/nvme0n2	981fad118adb0fa	Linux	2	268.44 GB / 268.44 GB	512 B + 0 B	4.4.70-2	

[Skip to Step 5.](#)

4. For multipathing in SLES 15 SP4 and later operating systems, type the following command to display target connections on the initiator:

```
nvme list-subsys
```

Information similar to the following is displayed:

```
NVME-SUBSYS0 - NQN=NQN.2014-08.ORG.NVMEXPRESS:NVMF:UUID:19D48A6D-3644-4F65-8046-80BD88FE45A6
 \
 +- NVME0 FC TRADDR=NN-0x20000090FA942E16:PN-0x10000090FA942E16 HOST_TRADDR=NN-0x200000109B346F1F:PN-
0x100000109B346F1F
 +- NVME3 FC TRADDR=NN-0x20000090FA942E16:PN-0x10000090FA942E16 HOST_TRADDR=NN-0x200000109B346F20:PN-
0x100000109B346F20
NVME-SUBSYS1 - NQN=NQN.2014-08.ORG.NVMEXPRESS:NVMF:UUID:2c6d862f-b77e-43b1-b345-7f990125088E
```

```

\-
+- NVME1 FC TRADDR=NN-0x20000090FA942E16:PN-0x10000090FA942E16 HOST_TRADDR=NN-0x200000109B346F1F:PN-
0x100000109B346F1F
+- NVME6 FC TRADDR=NN-0x20000090FA942E16:PN-0x10000090FA942E16 HOST_TRADDR=NN-0x200000109B346F20:PN-
0x100000109B346F20
NVME-SUBSYS2 - NQN=NQN.2014-08.ORG.NVMEXPRESS:NVMF:UUID:4BBBD877-6DE2-4A67-AD56-E0333E2437B5
\
+- NVME2 FC TRADDR=NN-0x20000090FA942E16:PN-0x10000090FA942E16 HOST_TRADDR=NN-0x200000109B346F1F:PN-
0x100000109B346F1F
+- NVME7 FC TRADDR=NN-0x20000090FA942E16:PN-0x10000090FA942E16 HOST_TRADDR=NN-0x200000109B346F20:PN-
0x100000109B346F20
NVME-SUBSYS3 - NQN=NQN.2014-08.ORG.NVMEXPRESS:NVMF:UUID:C8DA7263-FEFE-4563-A14C-A9F0362DCB03
\
+- NVME5 FC TRADDR=NN-0x20000090FA942E16:PN-0x10000090FA942E16 HOST_TRADDR=NN-0x200000109B346F1F:PN-
0x100000109B346F1F
+- NVME9 FC TRADDR=NN-0x20000090FA942E16:PN-0x10000090FA942E16 HOST_TRADDR=NN-0x200000109B346F20:PN-
0x100000109B346F20
NVME-SUBSYS4 - NQN=NQN.2014-08.ORG.NVMEXPRESS:NVMF:UUID:5F1D7788-FAB1-45CF-AF9C-D30DC9A9CE89
\
+- NVME4 FC TRADDR=NN-0x20000090FA942E16:PN-0x10000090FA942E16 HOST_TRADDR=NN-0x200000109B346F1F:PN-
0x100000109B346F1F
+- NVME8 FC TRADDR=NN-0x20000090FA942E16:PN-0x10000090FA942E16 HOST_TRADDR=NN-0x200000109B346F20:PN-
0x100000109B346F20

```

NOTE: For additional information on multipathing, see [Appendix B, Native NVMe Multipathing](#).

5. Run the `lsblk` command on the initiator to display all the block devices, including newly discovered ones, by typing the following command:

`lsblk`

Information similar to the following is displayed:

NAME	MAJ:MIN	RM	SIZE	RO	TYPE	MOUNTPOINT
sda	8:0	0	136.7G	0	disk	
└─sda1	8:1	0	2G	0	part	
└─sda2	8:2	0	40G	0	part	
└─sda3	8:3	0	94.7G	0	part	
└─3600508b1001037383941424344450400	254:0	0	136.7G	0	mpath	
├─3600508b1001037383941424344450400-part1	254:1	0	2G	0	part	[SWAP]
├─3600508b1001037383941424344450400-part2	254:2	0	40G	0	part	/var/opt
└─3600508b1001037383941424344450400-part3	254:3	0	94.7G	0	part	/home
sdb	8:16	0	3G	0	disk	
└─3600000e00d2a000002a017300340000	254:4	0	3G	0	mpath	
sdc	8:32	0	3G	0	disk	
└─3600000e00d2a000002a017300350000	254:5	0	3G	0	mpath	
sdd	8:48	0	3G	0	disk	
└─3600000e00d2a000002a017300360000	254:6	0	3G	0	mpath	
nvme0n1	259:0	0	3G	0	disk	
nvme0n2	259:1	0	3G	0	disk	

NOTE: The `/dev/sde` and `/dev/sdf` block devices, which were previously referred to on the NVMe target system, are now discovered by the NVMe initiator system, and they appear as the NVMe block devices `/dev/nvme0n1` and `/dev/nvme0n2`.

4.2 Emulex Support for NVMe CLI Commands

Supported RHEL and SLES operating systems include many native NVMe CLI commands for initiators. Of these, Emulex has tested, and thus supports, only the following commands:

- `connect`
- `connect-all`
- `disconnect`
- `disconnect-all`
- `gen-hostnqn`
- `list`
- `list-subsy` (SLES12 SP5 only)

Documentation for these commands is available with the operating system.

Chapter 5: Troubleshooting

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

5.1 Hardware Situations and Resolutions

[Table 5](#) lists the hardware situations and their resolutions.

Table 5: 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 the following events if an unapproved optical transceiver is detected:</p> <ul style="list-style-type: none"> ■ Unsupported module ■ Optics faulted, incorrectly installed, or not installed ■ Incompatible optics ■ Unknown port SFP 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 <i>Emulex Fibre Channel Host Bus Adapters Installation Guide</i> available at www.broadcom.com.</p>

5.2 FC Driver Situations and Resolutions

[Table 6](#) lists the FC driver situations and their resolutions.

Table 6: FC Driver Situations and Resolutions

Situation	Resolution
An FC link fails to come up.	If an FC link fails to come up, verify that an 8Gb/s or a 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 4Gb/s, 8Gb/s, and 16Gb/s devices are supported on 16Gb/s adapters. Only 8Gb/s, 16Gb/s, and 32Gb/s devices are supported on 32Gb/s adapters.
If a SAN configuration has 256 targets mapped by the FC driver, any additional added targets do not get a target ID mapping by the driver and cause target discovery to fail.	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.

Table 6: FC Driver Situations and Resolutions (Continued)

Situation	Resolution
<code>rmmmod</code> fails to unload an FC driver module due to ERROR: Module <code>lpfc</code> is in use.	<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 the following 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 HBA Manager application. 3. Stop Device Mapper.
<code>rmmmod</code> stops responding and the module reference count is 0.	<p>Due to a small race condition in the kernel, it is possible for an <code>rmmmod</code> command to stop responding. Issue the <code>rmmmod -w</code> command. If this does not help, reboot the computer.</p>
<code>rmmmod</code> fails to unload the driver due to device or resource busy.	<p>This situation occurs when you attempt to unload the driver without first stopping the Emulex HBA Manager application, when the Emulex HBA Manager daemons are running, or when FC disks connected to an Emulex adapter are mounted. Use the following steps to resolve this situation:</p> <p>Stop the Emulex HBA Manager application before attempting to unload the driver.</p> <ol style="list-style-type: none"> 1. Run the following command: <code>/usr/sbin/ocmanager/stop_ocmanager</code> Or <code>/opt/emulex/ocmanager/scripts/stop_ocmanager.sh</code> 2. Unmount any disks connected to the adapter. 3. Unload the driver. Type: <code>rmmmod lpfc</code>
An <code>lspci</code> shows new Emulex adapters as unknown.	<p>This situation occurs because of the delay in getting new product IDs into the Linux development cycle.</p> <p>There is no resolution.</p>
The FC driver fails to recognize an adapter and logs unknown IOCB messages in the system log during driver load.	<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 2.2.1, Installing the Emulex FC Driver Kit.</p>
Loading the FC driver on SLES operating systems reports unsupported module, tainting kernel in the system log.	<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>
The system panics when it is booted with a failed adapter installed.	<p>Remove the failed adapter and reboot the system.</p>
Driver installation fails.	<p>The <code>elx_lpfc_install.sh</code> script fails to install the driver. The install script can fail for the following reasons:</p> <ul style="list-style-type: none"> ■ A previous version of the driver is installed. Run the <code>elx_lpfc_install.sh --uninstall</code> script, and then try to install the driver. ■ The current driver is already installed. ■ The RHEL or SLES kernel is not supported. Run a supported RHEL or SLES kernel. ■ Required dependencies are not installed.

Table 6: FC Driver Situations and Resolutions (Continued)

Situation	Resolution
The No module lpfc found for kernel KERNELVERSION RPM error message appears 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 that you use depends on whether you are updating the driver. <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.
A recently upgraded kernel cannot find the ramdisk. After upgrading the kernel, the kernel cannot find the ramdisk, which halts or panics the system.	For more information on upgrading the kernel, see Chapter 2, Installing and Uninstalling .
The driver is not loaded after a system reboot after upgrading the kernel.	
Driver uninstall fails.	<p>The <code>elx_lpfc_install.sh --uninstall</code> script fails with an error. Try the following solutions:</p> <ul style="list-style-type: none"> ■ Stop the Emulex HBA Manager application before attempting to unload the driver. <ol style="list-style-type: none"> 1. Run the following command: <code>/usr/sbin/ocmanager/stop_ocmanager</code> Or <code>/opt/emulex/ocmanager/scripts/stop_ocmanager.sh</code> 2. Unmount any disks connected to the adapter. 3. Unload the driver. Type: <code>rmmmod lpfc</code> ■ Unmount all FC disk drives. ■ Unload the <code>lpfc</code> driver. ■ Use <code>rpm -e `rpm -qa grep lpfc`</code> to uninstall the lpfc driver RPM.
The <code>elx_lpfc_install.sh</code> script exit message appears.	The <code>elx_lpfc_install.sh</code> script contains exit messages that can be useful in diagnosing installation issues.
The Emulex driver for Linux does not load in ramdisk for a custom-built kernel.	<p>Emulex does not support custom-built kernels. However, the Emulex installation script attempts to install the driver into a ramdisk that follows the naming scheme used by RHEL or SLES kernels.</p> <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> <p>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:</p> <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.
The Linux SCSI subsystem sees only eight LUNs when more are present.	<p>Some SCSI drivers do not scan past eight LUNs when the target reports itself as a SCSI-2 device.</p> <p>To resolve this situation, force a SCSI bus scan with the following command: <code>/usr/sbin/lpfc/lun_scan</code></p> <p>SUSE supplies a <code>/bin/rescan-scsi-bus.sh</code> script, which can be changed to scan everything.</p>

Table 6: FC Driver Situations and Resolutions (Continued)

Situation	Resolution
During RHEL operating system installation, the LPFC driver update disk (DUD) might fail to load the out-of-box driver on a multipath installation or stop responding when configured with dual paths to a target.	<p>To load the out-of-box driver, perform the following steps:</p> <ol style="list-style-type: none"> 1. Boot up the operating system installer, and press e to edit the GRUB. 2. Bring down the multipath adapter links by appending the <code>linuxefi</code> with the following driver parameter: <code>lpfc.lpfc_suppress_link_up=1</code> 3. Press CTRL+X to continue the installer. 4. Press CTRL+ALT+F2 to go to the System Console. 5. Remove the inbox driver by typing the following command: <code>rmmmod lpfc</code> 6. Install the out-of-box driver by setting the suppress link to 0. <code>modprobe lpfc lpfc_suppress_link_up=0</code> 7. Press CTRL+ALT+F6 to return to the installer GUI. 8. Continue with the operating system installation.

5.3 Firmware Update Error Messages

The following table lists error messages that pertain to firmware updates.

Table 7: Firmware Update Situations

Message	Resolution
Incompatible object	A component of the firmware image is incompatible with the HBA. Install a later version of the firmware from www.broadcom.com .
Firmware incompatible with flash part	A new flash device on the HBA is not supported by the version of firmware you are installing. Install a later version of firmware from www.broadcom.com .
Incorrect ASIC version	The firmware you are installing does not support the HBA. Install a firmware image that supports the HBA from www.broadcom.com .
Incompatible object - general error	The firmware image has an unknown error. Contact Broadcom® Technical Support.

5.4 Dynamic Firmware Logging

For firmware-related issues, you might want to refer to the firmware logs. Dynamic firmware logging allows you to collect the firmware logs when required or when you encounter a firmware issue. You can also enable dynamic firmware logging in the absence of the Emulex HBA Manager application to dynamically generate firmware logs directly from the driver.

To dynamically generate firmware logs, perform the following steps:

1. Enable `lpfc_ras_fwlog_bufsize`.

```
echo 1 > /sys/class/scsi_host/host#/lpfc_ras_fwlog_bufsize
```

`lpfc_ras_fwlog_bufsize` is the firmware logging host buffer size. For more information on `lpfc_ras_fwlog_bufsize`, see [Table 3, Dynamic FC Driver Parameters](#).

2. Mount debugfs.

```
mount -t debugfs - /debug
```

Use `debugfs` in the absence of the Emulex HBA Manager application.

3. Collect firmware logs.

```
cat /debug/lpfc/fn#/fw_log
```

where `#` is the port for which you want to generate firmware logs.

5.5 NVMe Driver Situations and Resolutions

Table 6 lists the NVMe driver situations and their resolutions.

Table 8: NVMe Driver Situations and Resolutions

Situation	Resolution
Unloading the Fibre Channel Protocol (FCP) driver using the <code>modprobe -r</code> command might cause issues on the initiator before NVMe devices are disconnected.	Unload the driver with <code>rmmmod lpfc</code> , or if <code>modprobe -r lpfc</code> is required, wait for the device loss period of 60 seconds before unloading the driver.
On the Oracle Linux 7.7 unbreakable enterprise kernel (UEK) R6 operating system, unloading a Linux FCP driver with NVMe connections might fail.	Disconnect the NVMe connections before unloading the Linux FCP driver.
On an NVMe controller, reconnectivity issues might occur during switch port toggles. This issue is an operating system issue.	This issue is fixed in RHEL 8.6. For SLES, contact SUSE technical support.
In the <code>fstab</code> file, when an NVMe namespace mount point entry is included and the <code>defaults</code> option is specified, if the operating system is rebooted, the operating system might go into emergency mode.	<p>For SLES 15 SP4: Specify <code>nofail</code> as the mount option when creating the NVMe namespace <code>fstab</code> mount point entry.</p> <p>For SLES 12 SP5 and RHEL8.x: Specify <code>nofail</code> as the mount option when creating the NVMe namespace <code>fstab</code> mount point entry.</p> <p>Install the <code>nvme</code> boot <code>dracut</code> module provided in the <code>nvme</code> <code>fc-dm</code> kit.</p>
On SLES 12 SP5, the <code>nvme</code> - <code>boot-connections</code> <code>systemd</code> unit is disabled after the operating system is installed using NVMe BFS. This issue leads to the operating system failing to discover all NVMe subsystems when a large number of NVMe subsystems exist.	<p>Use <code>systemctl enable</code> to manually enable the <code>nvme</code>-<code>boot-connections</code> <code>systemd</code> unit to start during boot.</p> <pre># systemctl enable nvme Created symlink /etc/systemd/system/default.target.wants/nvme-boot-connections.service → /usr/lib/systemd/system/nvme-boot-connections.service.</pre>
An unrecoverable operating system fault might occur when generating a crash dump after successful installation of the operating system in the NVMe BFS namespace.	<p>From the NVMe BFS kit, extract and install the <code>nvme</code>-<code>dm</code> RPM that matches the installed operating system.</p> <pre># tar zxf elx-nvme<code>fc-dm-<version>-<release>/nvme<code>fc-dm-<version>-<release>.OS.tar.gz</code></code></pre>
If the operating system was installed using NVMe over FC BFS, and you are updating the firmware from version 12.8.xxx.xx to version 14.0.xxx.xx or later, NVMe storage might not be found, and the UEFI boot menu might not boot the operating system.	In the UEFI configuration utility, enable the Generate and use a Port Based HostNQN option. Refer to the <i>Emulex Boot for the Fibre Channel and NVMe over FC Protocols User Guide</i> for details.
If the operating system was installed using NVMe over FC BFS, and you are updating the firmware from version 14.0.xxx.xx or later to version 12.8.xxx.xx, NVMe storage might not be found, and the UEFI boot menu might not boot the operating system.	Update the SAN storage configuration and add an adapter WWPN-based HostNQN to the storage array.
NVMe disks might not reconnect after a device timeout greater than 60 seconds has occurred.	You must perform a manual scan or connect using the <code>nvme connect-all</code> CLI command.

5.6 FC Driver Log Messages

The following sections describe retrieving and interpreting FC driver log messages.

5.6.1 Retrieving FC Driver Log Messages

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

An example of an FC message follows:

```
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 FC hardware port.
- `0:` – Indicates Emulex adapter 0.
- `1305` – Indicates a log message number of 1305.

NOTE: The following information applies to FC driver log messages:

- 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 2, Installing and Uninstalling](#).

[Table 9](#) lists the groups and defines the associated number ranges.

Table 9: Message Log Table

LOG Message Verbose Mask Definition	Verbose Bit	Verbose Description
LOG_ELS	0x00000001	ELS events
LOG_DISCOVERY	0x00000002	Link discovery events
LOG_MBOX	0x00000004	Mailbox events
LOG_INIT	0x00000008	Initialization events
LOG_LINK_EVENT	0x00000010	Link events
LOG_IP	0x00000020	IP traffic history
LOG_FCP	0x00000040	FCP traffic history
LOG_NODE	0x00000080	Node table events
LOG_TEMP	0x00000100	Temperature sensor events
LOG_BG	0x00000200	BlockGuard events
LOG_MISC	0x00000400	Miscellaneous events
LOG_SLI	0x00000800	SLI TM events
LOG_FCP_ERROR	0x00001000	Log errors, not underruns
LOG_LIBDFC	0x00002000	Libdfc events
LOG_VPORT	0x00004000	NPIV events
LOG_SECURITY	0x00008000	Security events
LOG_EVENT	0x00010000	CT, TEMP, DUMP, logging
LOG_FIP	0x00020000	FIP events
LOG_FCP_UNDER	0x00040000	FCP underruns errors

Table 9: Message Log Table (Continued)

LOG Message Verbose Mask Definition	Verbose Bit	Verbose Description
LOG_SCSI_CMD	0x00080000	ALL SCSI commands
LOG_NVME	0x00100000	NVME general events
LOG_NVME_DISC	0x00200000	NVME discovery/connect events
LOG_NVME_ABTS	0x00400000	NVME abort sequence events
LOG_NVME_IOERR	0x00800000	NVME I/O Error events
LOG_EDIF	0x01000000	External DIF events
LOG_AUTH	0x02000000	Authentication events
LOG_ALL_MSG	0xFFFFFFFF	LOG all messages

5.6.2 FC Error Log Messages and Their Descriptions

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

Table 10: FC Error Log Messages and Their Descriptions

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.
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 the device DID.
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.
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.
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.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

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.	
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.	
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.	
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.	
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.	
0137: No retry ELS command <ELS_CMD> to remote	
Data: (1) ulpStatus (2) ulpWord[4]	
Severity: Error	
Log: LOG_ELS	
Action: None required.	
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.	
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.	

Table 10: FC Error Log Messages and Their Descriptions (Continued)

0140: PLOGI Reject: invalid name Invalid node WWN provided. Data: None Severity: Error Log: LOG_ELS Action: None required.
0141: PLOGI Reject: invalid pname Invalid port WWN provided. Data: None Severity: Error Log: LOG_ELS Action: None required.
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.
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.
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.
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.
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.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

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.
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.
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.
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.
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.
2559: 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.
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 is 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.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

0222: Initial FLOG/FDISC timeout
The driver sent the initial FLOGI or FDISC to the fabric and never received a response.
Data: None
Severity: Error
Log: Always
Action: Check fabric configuration. The driver recovers from this situation and continues with device discovery.
0223: Timeout while waiting for NameServer login
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.
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.
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.
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.
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.
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.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

0233: Nodelist not empty Driver unloaded or hot plug detected a node still in use. Data: None Severity: Error Log: LOG_DISCOVERY Action: None required.
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.
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.
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.
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.
0251: NameServer login: no memory Could not allocate memory for the NDLP structure. Data: None Severity: Error Log: LOG_ELS Action: None required.
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.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

<p>0253: Register VPI: Can't send mbox\ Could not issue the REG_VPI mailbox command for this vPort.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_MBOX</p> <p>Action: None required.</p>
<p>0254: Register VPI: no memory Could not allocate memory for the REG_VPI mailbox command.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_MBOX</p> <p>Action: None required.</p>
<p>0255: Issue FDISC: no IOCB All of the preallocated IOCBs are in use.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_ELS</p> <p>Action: None required.</p>
<p>0256: Issue FDISC: Cannot send IOCB\ Unable to send the fabric IOCB.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_ELS</p> <p>Action: None required.</p>
<p>0257: GID_FT Query error The GID_FT CT request for the NameServer has failed.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_ELS</p> <p>Action: Check the switch configuration.</p>
<p>0258: Register Fabric login error: The REG_LOGIN for the fabric has failed.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_MBOX</p> <p>Action: Check the port connection and the switch configuration.</p>
<p>0259: No NPIV Fabric support The switch to which the port is connected does not support NPIV.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_ELS</p> <p>Action: Check the switch configuration.</p>

Table 10: FC Error Log Messages and Their Descriptions (Continued)

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
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.
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.
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.
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.
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.
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.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

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.
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.
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 midlayer.
0273: Unexpected discovery timeout, vport State x% The discovery process has timed out. Data: None Severity: Error Log: LOG_DISCOVERY Action: Verify that all targets are still visible.
0274: lpfc_nlp_put: ndlp:x%pusgmap:x%x refcnt:%d, void *)ndlp, ndlp->nlp_usg_map, atomic_read(&ndlp->kref.refcount) Data: None Severity: Warning Log: LOG_NODE Action: None required.
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.
0276: lpfc_nlp_get: ndlp:x%pusgmap:x%x refcnt:%d, void *)ndlp, ndlp->nlp_usg_map, atomic_read(&ndlp->kref.refcount) A kref_get was attempted on a node that was being released. Data: None Severity: Warning Log: LOG_NODE Action: None required.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

0277: lpfc_enable_node: ndlp:x%pusgmap:x% refcnt:%d, void *)ndlp, ndlp->nlp_usg_map, atom-ic_read(&ndlp->kref.refcount)
Enable node was attempted on a node that has already been marked for memory free.
Data: None
Severity: Warning
Log: LOG_NODE
Action: None required.
0278: lpfc_enable_node: ndlp:x%pusgmap:x% refcnt:%d, void *)ndlp, ndlp->nlp_usg_map, atom-ic_read(&ndlp->kref.refcount)
Enable node was attempted on an inactive node.
Data: None
Severity: Warning
Log: LOG_NODE
Action: None required.
0280: lpfc_cleanup_node: ndlp:x%pusgmap:x% refcnt:%d, void *)ndlp, ndlp->nlp_usg_map, atom-ic_read(&ndlp->kref.refcount)
Node cleanup was attempted on a node that has already been marked for memory free.
Data: None
Severity: Warning
Log: LOG_NODE
Action: None required.
0281: lpfc_cleanup_node: ndlp:x%pusgmap:x% refcnt:%d, void *)ndlp, ndlp->nlp_usg_map, atom-ic_read(&ndlp->kref.refcount)
Node cleanup was called to prepare the node for release.
Data: None
Severity: Warning
Log: LOG_NODE
Action: None required.
0282: ldid:x%x ndlp:x%pusgmap:x% refcnt:%d, ndlp->nlp_DID, (void *)ndlp, lpfc_init.c-ndlp->nlp_usg_map,
Driver clean-up has found a node that is still on the node list during driver unload or PCI hot plug removal.
Data: None
Severity: Error
Log: LOG_NODE
Action: None required.
06119: Failed to allocate mbox cmd memory
Mailbox allocation error.
Data: None
Severity: Error
Log: LOG_INIT
Action: None required.
0285: Allocated DMA memory size <allocflen> is less than the requested DMA memory size <reqflen>
Memory allocation was truncated.
Data: None
Severity: Error
Log: LOG_INIT
Action: None required.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

0286: lpfc_nlp_state_cleanup failed to allocate statistical data buffer <nlp_DID> Memory allocation failed for node's statistical data. Data: None Severity: Error Log: LOG_INIT Action: None required.
0287: lpfc_alloc_bucket failed to allocate statistical data buffer DID <nlp_DID> Memory allocation failed for node's statistical data. Data: None Severity: Error Log: LOG_NODE Action: None required.
0289: Issue Register VFI failed: Err <rc> The driver could not register the Virtual Fabric Index for the FCFI. Data: None Severity: Error Log: LOG_ELS Action: Check the switch and port configurations.
0290: The SLI4 DCBX asynchronous event is not handled yet The SLI-4 DCBX asynchronous event is not handled yet. Data: None Severity: Error Log: LOG_SLI Action: None required.
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. Data: None Severity: Error Log: LOG_INIT Action: Check the switch configuration.
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.
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.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

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.
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.
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.
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.
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.
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.
0303: Ring <ringno> handler: portRspPut <portRspPut> is bigger than 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.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

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.	
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.	
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.	
0312: Ring <ringno> handler: portRspPut <rspPutInx> is bigger than rsp ring <numRiocb>	The IOCB command rings put pointer is ahead of the get pointer.
Data: None	
Severity: Error	
Log: LOG_SLI	
Action: None required.	
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.	
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.
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.	
0317: iotag <ulp_IoTag> is out of range: max iotag <max_iotag> wd0 <wd0>	The IoTag in the completed IOCB is out of range.
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.	

Table 10: FC Error Log Messages and Their Descriptions (Continued)

0318: Failed to allocate IOTAG. last IOTAG is <last_allocated_iotag> The driver cannot allocate an IOTAG. Display the last value used. 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.
0319: READ_SPARAM mbxStatus error <mbxStatus> hba state <hba_state> The driver issued a READ_SPARAM 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.
0320: CLEAR_LA mbxStatus error <mbxStatus> hba state <hba_state> The driver issued a CLEAR_LA 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.
0322: Ring <ringno> handler: unexpected completion IoTag <IoTag> The driver could not find a matching command for the completion received on the specified ring. 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.
0323: Unknown Mailbox command <mbxCommand> Cmpl An unknown mailbox command completed. 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.
0324: Config port initialization error, mbxCmd <mbxCommand> READ_NVPARM, mbxStatus <mbxStatus> The READ_NVPARMS mailbox command failed during port configuration. 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.
0328: Rsp Ring <ring number> error: IOCB Data: The firmware has returned an error for this IOCB. Data: (1) <iocb word[0]:iocb word[7]>, (2) <rsp word[0]:rsp[word[7]> Severity: Warning Log: LOG_SLI Action: None required.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

0330: IOCB wake NOT set
The completion handler associated with the IOCB was never called.
Data: (1) timeout (2) timeleft/jiffies
Severity: Error
Log: Always
Action: This error could indicate a software driver, firmware, or hardware issue. If this issue persists, report the error to Broadcom Technical Support.
0334: Unknown IOCB command
Received an unknown IOCB command completion.
Data: (1) type (2) ulpCommand (3) ulpStatus (4) ulploTag (5) ulpContext
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.
0335: Unknown IOCB command
Received an unknown IOCB command completion.
Data: (1) ulpCommand (2) ulpStatus (3) ulploTag (4) ulpContext
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
0336: Rsp Ring <ringno> error: IOCB
An IOCB error has occurred on the specified ring.
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
Severity: Warning
Log: LOG_SLI verbose
Action: If this issue persists, check the targets. If the targets are not the problem, report the error to Broadcom Technical Support.
0340: Adapter temperature is OK now
Adapter temperature has reverted to normal range.
Data: Temperature in Celsius
Severity: Error
Log: LOG_TEMP verbose
Action: No action needed, informational.
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.
0342: Ring <ringno> Cannot find buffer for an unsolicited iocb tag <unsli3.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.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

<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 <unsli3.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% 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 <code>lpfc_log_verbose</code> 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: Shut down and remove the HBA. Contact Broadcom Technical Support.</p>
<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>

Table 10: FC Error Log Messages and Their Descriptions (Continued)

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.
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.
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.
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.
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.
0360: Unsupported EQ count. <entry_count>
Cannot create an event queue of this size.
Data: None
Severity: Error
Log: LOG_SLI
Action: None required.
0361: Unsupported CQ count. <entry_count>
Cannot create an completion queue of this size.
Data: None
Severity: Error
Log: LOG_SLI
Action: None required.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

<p>0362: Unsupported MQ count. <entry_count> Cannot create a mailbox queue 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 identifier passed in the event queue entry does not reference a valid completion queue. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>0366: Not a valid completion event: majorcode=<major code hex>, minor-code=<minor code hex> The major or minor code in the Event Queue field is invalid. 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 completion queue identifier in the event queue entry does not match the <code>fcp_cqid</code> that was passed into the routine. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>
<p>0369: No entry from completion queue qid No completions in the completion queue were referenced by <code>fcp_cqid</code>. 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. Data: None Severity: Error Log: LOG_SLI Action: None required.</p>

Table 10: FC Error Log Messages and Their Descriptions (Continued)

0372: iotag <iotag> lookup error: max iotag (<sli.last_iotag>) iocb_flag <iocb_flag> The IOCB lookup cannot be performed because the iotag is out of range. Data: None Severity: Error Log: LOG_SLI Action: None required.
0373: FCP complete error: status=<status> hw_status=<hw status>, total_data_specified=<total data transferred>, parameter=<rsp word[4]>, word3=<wcqe word 3> Logs the FCP failure. Status and parameter are equivalent to ulpStatus and ulpWord[4]. Data: None Severity: Warning Log: LOG_SLI Action: None required.
0374: FCP complete with no corresponding cmdiocb: iotag <iocb iotag> No IOCB was on the in-progress list that matched this iotag. Data: None Severity: Warning Log: LOG_SLI Action: None required.
0375: FCP cmdiocb not callback function iotag: <iocb iotag> The IOCB found for this iotag does not have a completion handler set in it. Data: None Severity: Warning Log: LOG_SLI Action: None required.
0377: Error <rc> parsing vpd. Using defaults. Could not parse the VPD data, so the driver is using the default values. Data: None Severity: Error Log: Always Action: None required.
0378: No support for fcpi mode. Could not configure the port to run in FCP initiator mode. Data: None Severity: Warning Log: LOG_MBOX, LOG_SLI Action: None required.
0379: Feature Mismatch Data: <req_ftr word2 hex> <req_ftr word3 hex> <cfg_enable_npiv> <max_vpi hex> The features passed in to the driver as module parameters do not match what the firmware can do. Setting to default values. Data: None Severity: Warning Log: LOG_MBOX, LOG_SLI Action: None required.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

<p>0381: Error %d during queue setup. Could not set up all the queues that the driver requires to exchange I/Os with the HBA. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: Reload the driver.</p>
<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>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 cmdiocb: 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: NO IOCBQ data: txq_cnt=<txq_cnt> iocb_cnt=<iocb_cnt> els_txcmplq_cnt=<els_txcmplq_cnt> 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 10: FC Error Log Messages and Their Descriptions (Continued)

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.
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.
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.
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.
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.
0398 Invalid link fault code: <hex link_fault> The attempt to read the link attention register returned an unknown value. Data: None Severity: Error Log: LOG_INIT Action: None required.
0399 Invalid link attention type: <hex link_type> The READ_LA mailbox command returned an invalid link type. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

<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>
<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>

Table 10: FC Error Log Messages and Their Descriptions (Continued)

0421: MSI-X slow-path request_irq failed <rc>
The kernel API to request an IRQ has failed.
Data: None
Severity: Warning
Log: LOG_INIT
Action: Use module parameter <code>lpfc_use_msi=0 (INTx)</code> .
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)
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.
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.
0425:lpfc_restrict_login attribute cannot be set to %d, allowed range is [0, 1]
The module parameter <code>lpfc_restrict_login</code> can only be set to 0 (off) or 1 (on).
Data: None
Severity: Error
Log: LOG_INIT
Action: Set <code>lpfc_restrict_login = [0,1]</code> .
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.
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.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

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.
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.
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.
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.
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.
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.
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.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

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.
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.
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.
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.
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.
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.
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.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

<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>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>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 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 still outstanding:last compl time was <time> ms. The Heartbeat mailbox command failed. Data: None Severity: Error Log: LOG_INIT Action: Ensure that the adapter firmware is current. Reload the driver.</p>

Table 10: FC Error Log Messages and Their Descriptions (Continued)

0460: Bring adapter offline
The FC driver received a request to bring the adapter offline. This issue might occur when running <code>lputil</code> .
Data: None
Severity: Warning
Log: LOG_INIT verbose
Action: None required.
0466: Outstanding I/O when waiting for I/O to drain
The I/O is still pending while attempting to stop the driver.
Data: None
Severity: Warning
Log: LOG_INIT
Action: None required.
0467: <code>lpfc_topology</code> attribute cannot be set to %d, allowed range is [0, 6], <code>phba->brd_no</code> , <code>val</code> .
The <code>lpfc_topology</code> module parameter is invalid.
Data: None
Severity: Error
Log: LOG_INIT
Action: Use a value in the valid range.
0468: <code>lpfc_restrict_login</code> must be 0 for Physical ports.
Cannot restrict the login for the physical port.
Data: None
Severity: Error
Log: LOG_INIT
Action: None required.
0469: <code>lpfc_link_speed</code> attribute cannot be set to <val>, allowed values are [list of val]
The link speed module parameter is invalid.
Data: None
Severity: Error
Log: LOG_INIT
Action: Use a link speed parameter in the valid range.
0472: Unknown PCI error state: <state>
The PCI bus detected an error.
Data: None
Severity: Error
Log: LOG_INIT
Action: Issue an HBA reset.
0474: Unable to allocate memory for issuing <code>MBOX_CONFIG_MSI</code> command
Mailbox memory pool allocation error.
Data: None
Severity: Error
Log: LOG_INIT
Action: None required.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

0475: Not configured for supporting MSI-X cfg_use_msi: 0x% 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.
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.
0478: MSI request_irq failed (<rc>). The request_irq kernel API failed. Data: None Severity: Warning Log: LOG_INIT Action: Set lpfc_use_msi = 0.
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.
0486: MSI-X fast-path (<index>) request_irq failed (<rc>). The request_irq kernel API failed. Data: None Severity: Warning Log: LOG_INIT Action: Set module parameter lpfc_use_msi = 0.
0490: MSI request_irq failed (<rc>). The request_irq kernel API failed. Data: None Severity: Warning Log: LOG_INIT Action: Set module parameter lpfc_use_msi = 0.
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.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

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.
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.
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.
0497: Failed allocate fast-path IO WQ (<wq index>) The event queue for the fast path was not allocated. Data: None Severity: Error Log: LOG_INIT Action: Unload and reload the driver.
0499: Failed allocate fast-path FCP CQ (<fcq_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.
0500: Failed allocate slow-path mailbox CQ Failed to allocate slow-path mailbox CQ. Data: None Severity: Error Log: LOG_INIT Action: None required.
0501: Failed allocate slow-path ELS CQ Failed to allocate slow-path ELS CQ. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

0503: Failed allocate fast-path FCP Failed to allocate fast-path FCP. Data: None Severity: Error Log: LOG_INIT Action: None required.
0504: Failed allocate slow-path ELS WQ Failed to allocate slow-path ELS WQ. Data: None Severity: Error Log: LOG_INIT Action: None required.
0505: Failed allocate slow-path ELS MQ Data: None Severity: Error Log: LOG_INIT Action: None required.
0506: Failed allocate receive HRQ\n Data: None Severity: Error Log: LOG_INIT Action: None required.
0507: Failed allocate receive DRQ Failed to allocate receive DRQ. Data: None Severity: Error Log: LOG_INIT Action: None required.
0523: Failed setup of fast-path EQ <fc_eq_idx>, rc = <rc> The fast-path EQ setup failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
0528: Mailbox CQ not allocated The mailbox CQ is not allocated. Data: None Severity: Error Log: LOG_INIT Action: None required.
0530: ELS CQ not allocated The ELS CQ is not allocated. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

<p>0535: Failed setup of fast-path FCP WQ <fcp_wqidx>, rc = <rc> The fast-path FCP WQ setup failed. 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 10: FC Error Log Messages and Their Descriptions (Continued)

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.
0602: Failed to allocate CQ_EVENT entry Failed to allocate a CQ_EVENT entry. Data: None Severity: Error Log: LOG_SLI Action: None required.
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.
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.
0704: At limitation of <total> preallocated command buffers The maximum number of command buffers has already been allocated. Data: None Severity: Warning Log: LOG_FCP verbose Action: None required.
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.
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 is allocated. Data: None Severity: Warning Log: LOG_FCP Action: None required.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

0713: SCSI layer issued Device Reset (%d, %d)
A device reset was issued.
Data: None
Severity: Error
Log: LOG_FCP
Action: None required.
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.
0721: Device Reset rport failure: rdata <rdata>
The reset of the R_Port failed.
Data: None
Severity: Error
Log: LOG_FCP
Action: None required.
0722: Target Reset rport failure: rdata <rdata>
The reset of the target failed.
Data: None
Severity: Error
Log: LOG_FCP
Action: None required.
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.
0724: I/O flush failure for context <"LUN", "TGT", "HOST", "Unknown">: cnt <cnt>
The I/O flush to the LUN, target, or host failed.
Data: None
Severity: Error
Log: LOG_FCP
Action: None required.
0727: TMF <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.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

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.
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.
0915 Register VPI failed:<mbxStatus> Could not register the VPI. Data: None Severity: Error Log: LOG_MBOX Action: None required.
1268 - Failed to allocate ndlp for oxid:x% SID:x% Failed to allocate NDLP. Data: None Severity: Error Log: LOG_ELS Action: None required.
1302: Invalid speed for this board: Reset link speed to auto: <cfg_link_speed> The driver is reinitializing the link speed to autodetect. Data: None Severity: Warning Log: LOG_LINK_EVENT verbose Action: None required.
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) AL_PA granted, (3) Link speed, (4) number of entries in the loop init LILP AL_PA map. An AL_PA map message is also recorded if LINK_EVENT verbose mode is set. Each AL_PA 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.
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.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

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.
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.
1308- Link Down Event in loop back mode x%x received Data: x%x x%x x%x Link Down Event received in loop back mode. Data: eventTag, fc_eventTag, port_state, fc_flag Severity: Error LOG: LOG_LINK_EVENT Action: None required.
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.
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.
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.
1400 - Failed to update xri-sgl size and mapping: %d Updating of host els xri-sgl sizes and mappings failed. Data: None Severity: Error LOG: LOG_MBOX LOG_SLI Action: None required.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

<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>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>
<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% "perr=x% sfi=x% nip=x% ipc=x% xrom=x% "dl=x% pstatus=x%\\n", sta_reg.word0, bf_get(lpfc_hst_state_perr, &sta_reg), The HBA POST failed. Data: None Severity: Error Log: LOG_INIT Action: None required.</p>

Table 10: FC Error Log Messages and Their Descriptions (Continued)

<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>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>
<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>

Table 10: FC Error Log Messages and Their Descriptions (Continued)

1419: Invalid HBA PCI device group: <dev_grp> Invalid HBA PCI device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.
1420: Invalid HBA PCI device group:<dev_grp> Invalid HBA PCI device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.
1421: Failed to set up hba Initialization failed to set up the HBA. Data: None Severity: Error Log: LOG_INIT Action: None required.
1422: HBA Unrecoverable error: uerr_lo_reg=<ue lo>, uerr_hi_reg=<ue hi>, online0_reg=<Online0>, online1_reg=<Online1> The HBA notified the driver that it encountered an unrecoverable error. Data: None Severity: Error Log: LOG_INIT Action: Perform a dump from the Emulex HBA Manager application. Then, unload and reload the driver.
1423: HBA Unrecoverable error: uerr_lo_reg=<ue lo>, uerr_hi_reg=<ue hi>, online0_reg=<Online0>, online1_reg=<Online1> The HBA notified the driver that it encountered an unrecoverable error. Data: None Severity: Error Log: LOG_INIT Action: Perform a dump from the Emulex HBA Manager application. Then, unload and reload the driver.
1424: Invalid PCI device group:<pci_dev_grp> Invalid HBA PCI device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.
1425: Invalid PCI device group: <pci_dev_grp> Invalid HBA PCI device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

1426: Invalid PCI device group: <pci_dev_grp> Invalid HBA PCI device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.
1427: Invalid PCI device group: <pci_dev_grp> Invalid HBA PCI device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.
1428: Invalid PCI device group: <pci_dev_grp> Invalid HBA PCI device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.
1429: Invalid PCI device group: <pci_dev_grp> Invalid HBA PCI device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.
1430: Failed to initialize sql list Failed to initialize SQL list. Data: None Severity: Error Log: LOG_INIT Action: None required.
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.
1432: Failed to initialize rpi headers. Failed to initialize RPI headers. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

1476: Failed to allocate sysfs attr Failed to allocate sysfs attributes. Data: None Severity: Error Log: LOG_INIT Action: None required.
1477: Failed to set up hba Failed to set up the HBA. Data: None Severity: Error Log: LOG_INIT Action: None required.
1603: Loopback test did not receive expected data length. actual length <len>expected length <full_size>. The loopback test did not receive the same amount of data that it transmitted. Data: None Severity: Error Log: LOG_LIBDFC Action: None required.
1800: Could not issue unreg_vpi Driver attempt to unregister VPI failed. Data: None Severity: Error Log: LOG_VPORT verbose Action: Software driver error. If this issue persists, report these errors to Broadcom Technical Support.
1801: Create vport work array FAILED: cannot do scsi_host_get The driver was unable to get a reference to a SCSI host. Data: None Severity: Warning Log: LOG_VPORT verbose Action: Software driver warning. If this issue persists, report these errors to Broadcom Technical Support.
1816: FLOGI NPIV supported, response data <port> The fabric reports support for NPIV upon FLOGI. Data: (1) response_multiple_NPort Severity: Warning Log: LOG_VPORT verbose Action: No action needed, informational.
1817: Fabric does not support NPIV - configuring single port mode The fabric reports no support for NPIV upon FLOGI. Data: None Severity: Warning Log: LOG_VPORT verbose Action: No action needed, informational.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

1818: VPort failed init, mbxCmd <mailbox command> READ_SPARM mbxStatus <mailbox status> , rc = <status>
A pending mailbox command that was issued to initialize the port failed.
Data: (1) mbxCommand (2) mbxStatus (3) rc
Severity: Error
Log: LOG_VPORT verbose
Action: Software driver error. If this issue persists, report these errors to Broadcom Technical Support.
1819: Unrecognized lpfc_sli_mode parameter: <mode>
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.
Data: (1) lpfc_sli_mode
Severity: Error
Log: LOG_VPORT verbose
Action: Correct the lpfc_sli_mode driver parameter setting. Valid values are 0, 2, and 3.
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 use the auto setting (0).
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.
1822: Invalid <name>: <xx: 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.
1823: Create VPORT failed. Duplicate WWN on HBA.
The port could not be created because it would duplicate an existing WWN HBA address. The resources for the port were discarded.
Data: None
Severity: Error
Log: LOG_VPORT verbose
Action: Provide a unique WWN.
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.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

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.
1826: Vport Disabled.
The port had to be disabled in the system.
Data: None
Severity: Error
Log: LOG_VPORT verbose
Action: No action, informational.
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.
1828: Vport Deleted
A vPort was deleted.
Data: None
Severity: Error
Log: LOG_VPORT verbose
Action: No action, informational.
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.
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.
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.
1832: No pending MBOX command to handle
Data: None
Severity: Error
Log: LOG_MBOX
Action: None required.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

1835: Vport discovery quiesce failed: state <port_state> fc_flags <fc_flag> wait msecs <jiffies_to_mssecs(jiffies - start_time)> Could not pause discovery on this vPort. Data: None Severity: Error Log: LOG_VPORT Action: None required.
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.
1837: Vport_delete failed: Cannot delete static vport Static vPorts cannot be deleted. Data: None Severity: Error Log: LOG_VPORT Action: None required.
1838: Failed to INIT_VPI on vpi <vpi> status <rc> Failed to INIT_VPI. Data: None Severity: Error Log: LOG_VPORT Action: None required.
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.
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.
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.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

2004: Failed to allocate XRI.last XRI TAG is <XRI> Max XRI is <MAX_XRI>, Used XRI is <USED_XRI>. All XRI are in use.
Data: None
Severity: Warning
Log: LOG_SLI
Action: None required.
2005 - Reducing Queues - FW resource limitation: WQ %d CQ %d EQ %d: min %d: IRQ %d HDWQ %d Reducing queues due to resource limitations.
Data: None
Severity: Error
Log: LOG_SLI
Action: None required.
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.
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.
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.
2010: Resume RPI Mailbox failed status <status>, mbxStatus <mbx status>
Data: None
Severity: Error
Log: LOG_SLI
Action: None required.
2011: Unable to allocate memory for issuing SLI_CONFIG_SPECIAL mailbox command
Data: None
Severity: Error
Log: LOG_SLI
Action: None required.
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.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

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.
2014: Invalid command <iocb.ulpCommand> The IOCB command is invalid. Data: None Severity: Error Log: LOG_SLI Action: None required.
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.
2017: REG_FCFI mbxStatus error <mbx status> HBA state <port_state> The REG_FCFI mailbox command failed. Data: None Severity: Error Log: LOG_MBOX Action: None required.
2018: REG_VFI mbxStatus error <mbx status> HBA state <port_state> The REG_VFI mailbox command failed. Data: None Severity: Error Log: LOG_MBOX Action: None required.
2022: VPI Mailbox failed status <status>, mbxStatus <mbxStatus> The INIT-VPI mailbox command failed. Data: None Severity: Error Log: LOG_SLI Action: None required.
2500: EQ_CREATE mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc> The mailbox command sent to create the event queue failed. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

2501: CQ_CREATE mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>
The mailbox command sent to create the completion queue failed.
Data: None
Severity: Error
Log: LOG_INIT
Action: None required.
2502: MQ_CREATE mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>
The mailbox command sent to create the mailbox queue failed.
Data: None
Severity: Error
Log: LOG_INIT
Action: None required.
2503: WWQ_CREATE mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>
The mailbox command sent to create the work queue failed.
Data: None
Severity: Error
Log: LOG_INIT
Action: None required.
2504: RQ_CREATE mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>
The mailbox command sent to create the RQ failed.
Data: None
Severity: Error
Log: LOG_INIT
Action: None required.
2505: EQ_DESTROY mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>
The mailbox command sent to delete the event queue failed.
Data: None
Severity: Error
Log: LOG_INIT
Action: None required.
2506: CQ_DESTROY mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>
The mailbox command sent to delete the completion queue failed.
Data: None
Severity: Error
Log: LOG_INIT
Action: None required.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

2507: MQ_DESTROY mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>
The mailbox command sent to delete the mailbox queue failed.
Data: None
Severity: Error
Log: LOG_INIT
Action: None required.
2508: WQ_DESTROY mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>
The mailbox command sent to delete the work queue failed.
Data: None
Severity: Error
Log: LOG_INIT
Action: None required.
2509: RQ_DESTROY mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>
The mailbox command sent to delete the work queue failed.
Data: None
Severity: Error
Log: LOG_INIT
Action: None required.
2510: RQ_DESTROY mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>
The mailbox command sent to delete the work queue failed.
Data: None
Severity: Error
Log: LOG_INIT
Action: None required.
2511: POST_SGL mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>n
The mailbox command sent to post the SGL pages to the firmware failed.
Data: None
Severity: Error
Log: LOG_INIT
Action: None required.
2512 - MODIFY_EQ_DELAY mailbox failed with status x%x add_status x%x, mbx status x%x
The mailbox command sent to modify EQ delay failed.
Data: None
Severity: Error
Log: LOG_INIT
Action: None required.
2513: POST_SGL_BLOCK mailbox command failed status <shdr_status> add_status <shdr_add_status> mbx status <rc>
The mailbox command sent to post the SGL pages to the firmware failed.
Data: None
Severity: Error
Log: LOG_INIT
Action: None required.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

2514: POST_RPI_HDR mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>
The mailbox command sent to post the RPI header pages to the firmware failed.
Data: None
Severity: Error
Log: LOG_INIT
Action: None required.
2515: ADD_FCF_RECORD mailbox failed with status <rc>
The mailbox command to add the FCF record failed.
Data: None
Severity: Error
Log: LOG_INIT
Action: None required.
2516 - CGN SET_FEATURE mbox failed with status x%x add_status x%x, mbx status x%x
Reset congestion to FPINs only.
Data: None
Severity: Error
LOG: LOG_INIT LOG_CGN_MGMT
Action: None required.
2517 - CGN Signal Event - Warn: x%x Alarm: x%x
Data: None
Severity: Error
LOG: LOG_CGN_MGMT
Action: None required.
2518 - EEPROM failure - Event Data1: x%08x Event Data2:
EEPROM failure occurred.
Data: event_data1, event_data2
Severity: Error
LOG: LOG_SLI
Action: Contact Broadcom Technical Support.
2521: READ_FCF_RECORD mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx
The READ_FCF_RECORD mailbox command failed.
Data: None
Severity: Error
Log: LOG_INIT
Action: None required.
2523: Allocated DMA memory size (<alloc_len>) is less than the requested DMA memory size (<req_len>)
The ADD_FCF_RECORD mailbox command failed to retrieve the length required from the firmware.
Data: None
Severity: Error
Log: LOG_INIT
Action: None required.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

<p>2524: Failed to get the non-embedded SGE virtual address The READ_FCF_RECORD mailbox command could not retrieve the SGE that was requested. Data: None Severity: Error Log: LOG_MBOX Action: None required.</p>
<p>2527: Failed to allocate non-embedded SGE array. Failed to allocate the nonembedded SGE array. Data: None Severity: Error Log: LOG_MBOX Action: None required.</p>
<p>2528: Mailbox command <vpi> cannot issue The mailbox command could not be issued because the mailbox interrupt is disabled. Data: (1) mbxCommand (2) sli_flag (3) flag Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.</p>
<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 <cpi> 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>

Table 10: FC Error Log Messages and Their Descriptions (Continued)

<p>2535: Unsupported RQ count. (<entry_count>) The RQ 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 RQ 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>
<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>

Table 10: FC Error Log Messages and Their Descriptions (Continued)

<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 is found. Data: None Severity: Error Log: LOG_DISCOVERY Action: None required.</p>
<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 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 failed. Data: None Severity: Error Log: LOG_DISCOVERY, LOG_MBOX Action: None required.</p>

Table 10: FC Error Log Messages and Their Descriptions (Continued)

2551: UNREG_FCFI mbox allocation failed HBA state <port_state> The allocation for the UNREG_FCFI mailbox command failed. Data: None Severity: Error Log: LOG_DISCOVERY, LOG_MBOX Action: None required.
2552: UNREG_FCFI issue mbox failed rc <rc> HBA state <port_state> The UNREG_FCFI mailbox command failed. Data: None Severity: Error Log: LOG_DISCOVERY, LOG_MBOX Action: None required.
2553: lpfc_unregister_unused_fcf failed to read FCF record HBA state Data: None Severity: Error Log: LOG_DISCOVERY, LOG_MBOX Action: None required.
2554: Could not allocate memory for fcf record Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.
2555: UNREG_VFI mbxStatus error <u.mb.mbxStatus> HBA state <port_state> The UNREG_VFI mailbox command failed. Data: None Severity: Error Log: LOG_DISCOVERY, LOG_MBOX Action: None required.
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.
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.
2558: ADD_FCF_RECORD mailbox failed with status<shdr_status> add_status <shdr_add_status> The ADD_FCF_RECORD mailbox command failed. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

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.
2562 - Failure to allocate an ELS sgl entry:%d Memory allocation failed for SGL entry. Data: None Severity: Error LOG: LOG_SLI Action: None required.
2563 - Failure to allocate an ELS mbuf:%d Memory allocation failed for ELS buffer. Data: None Severity: Error LOG: LOG_SLI Action: None required.
2566: Failed to allocate table entry Failed to allocate connection table entry. Data: None Severity: Error Log: LOG_INIT Action: None required.
2567: Config region 23 has bad signature The driver was unable to read Config Region 23 because it has an invalid signature. Data: None Severity: Error Log: LOG_INIT Action: None required.
2568: Config region 23 has bad version The driver was unable to read Config Region 23 because it is an invalid version. Data: None Severity: Error Log: LOG_INIT Action: None required.
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.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

2574 - IO channels: hdwQ %d IRQ %d MRQ: %d Information about HWQ, interrupts, and MRQ. Data: None Severity: Error LOG: LOG_INIT Action: None required.
2579: Slow-path wqe consume event carries miss-matched qid: wcqe-qid=<wcqe_quid>, sp-qid=<sp_quid> The consumed entry does not have the queue ID of the slow-path. Data: None Severity: Warning Log: LOG_SLI Action: None required.
2580: Fast-path wqe consume event carries miss-matched qid: wcqe-qid=<fcp_wqid>. The consumed entry does not have the queue ID of the fast path. Data: None Severity: Warning Log: LOG_SLI Action: None required.
2597 - Sync Mailbox command x%x (x%x/x%x) failure: mqe_sta: x%x mcqe_sta: x%x/x%x Data: x%x x%x Issue of synchronous mailbox command failed. Data: None Severity: Error LOG: LOG_MBOX LOG_SLI Action: None required.
2598: Adapter Link is disabled. The adapter link is disabled. Data: None Severity: Error Log: LOG_INIT Action: None required.
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.
2600 - failed to allocate mailbox memory Memory allocation failed. Severity: Error Data: None LOG: LOG_INIT Action: None required.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

2605: lpfc_dump_static_vport: memory allocation failed Failed to allocate mailbox memory. Data: None Severity: Error Log: LOG_MBOX Action: None required.
2606: No NPIV Fabric support No NPIV Fabric support. Data: None Severity: Error Log: LOG_ELS Action: None required.
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.
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.
2609: Init VPI mailbox failed <u.mb.mbxStatus> The INIT_VPI mailbox command failed. Data: None Severity: Error Log: LOG_MBOX Action: None required.
2610 - FCF (x% x) reached driver's book keeping dimension:x% x FCF index exceeded maximum FCF table index. Data: None Severity: Error Log: LOG_FIP Action: None required.
2611 - FLOGI failed on FCF (x% x), status:x% x/x% x, tmo:x% x, perform roundrobin FCF failover FLOGI failed on FCF. Data: None Severity: Error Log: LOG_FIP LOG_ELS Action: None required.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

2619: Config region 23 has bad signature Configuration region 23 has an invalid signature. Data: None Severity: Error Log: LOG_INIT Action: None required.
2620: Config region 23 has bad version Configuration region 23 has an invalid version. Data: None Severity: Error Log: LOG_INIT Action: None required.
2621 - IO completion: Not an active IO The I/O was not an active I/O in the completion routine. Data: None Severity: Error Log: LOG_INIT Action: None required.
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 configuration mailbox command. Data: None Severity: Error Log: LOG_SLI Action: None required.
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.
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.
2719: Invalid response length: tgt <TGT_ID> lun <LUN> cmnd <CMD> rsplen <RSPLLEN> The response length for this FCP command is not supported. Data: None Severity: Error Log: LOG_FCP Action: None required.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

<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.</p> <p>Data: None</p> <p>Severity: Warning</p> <p>Log: LOG_ELS</p> <p>Action: None required.</p>
<p>2726: READ_FCF_RECORD Indicates empty FCF table The driver requested the firmware for a list of FCF entries to connect, to and the firmware responded indicating that the FCF table is empty.</p> <p>Data: None</p> <p>Severity: Error</p> <p>Log: LOG_INIT</p> <p>Action: None required.</p>
<p>6045: Issue GEN REQ WQE to NPORT <nlp_DID> Data <iotag> <port_state> Not able to queue NVMe LS command.</p> <p>Data: (1) nlp_DID, (2) iotag and (3) port_state.</p> <p>Severity: Error</p> <p>Log: LOG_ELS</p> <p>Action: Reload the driver.</p>
<p>6088: DID x<port_id> not ready for IO. State <nlp_state>, Type <nlp_type> Target port it not ready.</p> <p>Data: (1) port_id, (2) nlp_state and (3) nlp_type.</p> <p>Severity: Error</p> <p>Log: LOG_NODE</p> <p>Action: Verify that all targets are still visible.</p>
<p>6051: Remoteport x<pnvme_rport>, rport has invalid ndlp. Failing LS Req Target port it not ready.</p> <p>Data: pnvme_rport</p> <p>Severity: Error</p> <p>Log: LOG_NODE</p> <p>Action: Verify that all targets are still visible.</p>
<p>6044: Could not find node for DID <port_id> Not able to allocate memory for NVMe LS request.</p> <p>Data: port_id</p> <p>Log: LOG_NVME_DISC</p> <p>Severity: Error</p> <p>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>6081: NVME Completion Protocol Error: xri <sli4_xritag> status x<status> result x<result> placed x<total_data_placed> NVMe I/O completed with a protocol error.</p> <p>Data: (1) sli4_xritag, (2) status, (3) result and (4) total_data_placed.</p> <p>Log: LOG_NVME_IOERR</p> <p>Severity: Error</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>

Table 10: FC Error Log Messages and Their Descriptions (Continued)

6072: NVME Completion Error: xri <sli4_xritag> status x<status> result x<result> [x<parameter>] placed x<total_data_placed> NVMe I/O completed with an error. Data: (1) sli4_xritag, (2) status, (3) result, (4) parameter, and (5) total_data_placed. Log: LOG_NVME_IOERR Severity: Info Action: This error could indicate a software driver, firmware, or hardware issue. If this issue persists, report the error to Broadcom Technical Support.
6058: Too many sg segments from NVME Transport. Max <cfg_nvme_seg_cnt>, nvmeIO sg_cnt <seg_cnt> Too many sg segments for an NVMe command request. Data: (1) cfg_nvme_seg_cnt and (2) seg_cnt. Log: LOG_NVME_IOERR Severity: Error Action: This error could indicate a software driver issue. If this issue persists, report the error to Broadcom Technical Support.
6052: EXIT. issue ls wqe failed lport x<pnvme_lport>, rport x<pnvme_rport> lsreq x<pnvme_lsreq> Status <ret> DID <nlp_DID> Not able to queue NVMe LS command. Data: (1) pnvme_lport, (2) pnvme_rport, (3) pnvme_lsreq, (4) ret and (5) nlp_DID. Log: LOG_NVME_DISC Severity: Error Action: This error could indicate a software driver, firmware, or hardware issue. If this issue persists, report the error to Broadcom Technical Support.
6049: Could not find node for DID <port_id> Not able to abort LS request. Data: port_id Log: LOG_NVME_ABTS Severity: Error Action: This error could indicate a software driver issue. If this issue persists, report the error to Broadcom Technical Support.
6071: Null lpfc_ncmd pointer. No release, skip completion Skip NVMe I/O completion because of null lpfc_ncmd pointer. Data: None Log: LOG_NVME_IOERR Severity: Error Action: This error could indicate a software driver issue. If this issue persists, report the error to Broadcom Technical Support.
6066: Missing cmpl ptrs: lpfc_ncmd x<lpfc_ncmd>, nvmeCmd x<nvmeCmd> Missing completion pointer. Data: (1) lpfc_ncmd and (2) nvmeCmd. Log: LOG_NVME_IOERR Severity: Error Action: This error could indicate a software driver issue. If this issue persists, report the error to Broadcom Technical Support.
6062 Ignoring NVME cmpl. No ndlp Ignoring NVMe completion, no NDLP. Data: None Log: LOG_NVME_IOERR Severity: Error Action: This error could indicate a software driver issue. If this issue persists, report the error to Broadcom Technical Support.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

6059: dptr err <i>, nseg <nseg> Scatter-gather data pointer is null. Data: (1) i and (2) nseg Log: LOG_NVME_IOERR Severity: Error Action: This error could indicate a software driver issue. If this issue persists, report the error to Broadcom Technical Support.
6063: NVME DMA Prep Err: sg_cnt <sg_cnt> payload_length x<payload_length> NVMe DMA preparation failed. Data: (1) sg_cnt and (2) payload_length Log: LOG_NVME_IOERR Severity: Error Action: This error could indicate a software driver issue. If this issue persists, report the error to Broadcom Technical Support.
6139: Driver in reset cleanup - flushing NVME Req now. hba_flag <hba_flag> Driver in reset cleanup. Data: hba_flag Log: LOG_NVME_ABTS Severity: Error Action: None required.
6140: NVME IO req <freqpriv> has no matching lpfc nvme io buffer. Skipping abort req. No matching NVMe I/O buffer to abort, skipping abort request. Data: freqpriv Log: LOG_NVME_ABTS Severity: Error Action: This error could indicate a software driver issue. If this issue persists, report the error to Broadcom Technical Support.
6141: lpfc NVME IO req has no nvme_fcreq io buffer. Skipping abort req No nvmeCmd for NVMe I/O buffer to abort, skipping abort request. Log: LOG_NVME_ABTS Severity: Error Action: None required.
6143: NVME req mismatch: lpfc_nbuf x<lpfc_nbuf> nvmeCmd x<nvmeCmd>, pnvme_fcreq x<pnvme_fcreq>. Skipping Abort xri x<sli4_xritag> NVMe I/O abort request mismatch between FC NVMe command and transport pnvme_fcreq. Data: (1) lpfc_nbuf, (2) nvmeCmd, (3) pnvme_fcreq and (4) sli4_xritag. Log: LOG_NVME_ABTS Severity: Error Action: This error could indicate a software driver issue. If this issue persists, report the error to Broadcom Technical Support.
6142: NVME IO req x<pnvme_fcreq> not queued - skipping abort req xri x<sli4_xritag> NVMe I/O request pnvme_fcreq not queued, skipping abort request. Data: (1) pnvme_fcreq and (2) sli4_xritag. Log: LOG_NVME_ABTS Severity: Error Action: None required.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

6144: Outstanding NVME I/O Abort Request still pending on nvme_fcreq x<pnvme_fcreq>, lpfc_ncmd <lpfc_nbuf> xri x<sli4_xritag> Outstanding NVMEe I/O abort request still pending. Data: (1) pnvme_fcreq and (2) sli4_xritag. Log: LOG_NVME_ABTS Severity: Error Action: None required.
6176: Lport x<lport> Localport x<localport> wait timed out. Pending <pending>. Renewing. NVMe port is waiting for I/O completion. Data: (1) lport, (2) localport and (3) pending. Log: LOG_NVME_IOERR Severity: Error Action: This error could indicate a software driver issue. If this issue persists, report the error to Broadcom Technical Support.
6031: RemotePort Registration failed err: <ret>, DID x<nlp_DID> RemotePort registration with NVMe transport failed. Data: (1) ret and (2) nlp_DID. Log: LOG_NVME_DISC Severity: Error Action: This error could indicate a software driver issue. If this issue persists, report the error to Broadcom Technical Support.
6172: NVME rescanned DID x<nlp_DID> port_state x<port_state> NVMe rescanned for an NVMe target. Data: (1) nlp_DID and (2) port_state. Log: LOG_NVME_DISC Severity: Error Action: None required.
6167: NVME unregister failed <ret> port_state x<port_state> Unregister NVMe target port with NVMe transport failed. Data: (1) ret and (2) port_state. Log: LOG_NVME_DISC Severity: Error Action: None required.
6168: State error: lport x<localport>, rport x<rport> FCID <nlp_DID> Unregister NVMe target port failed. Data: (1) localport, (2) rport and (3) nlp_DID. Log: LOG_NVME_DISC Severity: Error Action: None required.
6244 No available WQEs for CMF_SYNC_WQE WQEs for CMF_SYNC_WQE are not available. Data: None Severity: KERN_ERR Log: LOG_CGN_MGMT Action: If this issue persists, report the error to Broadcom Technical Support.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

6154 Drop NVME LS: SID <source_id> OXID <original_exchange_id>: <failure_reason> NVME target dropped NVMe link service command. Data: (1) source_id (2) original_exchange_id and (3) failure_reason Severity: KERN_ERR Log: LOG_TRACE_EVENT Action: This could be a configuration issue. Please check the configuration. If this issue persists, report the error to Broadcom Technical Support.
6216 NVME Unsol rcv: No ndlp: NPort_ID <port_id> oidx <original exchange id> Received link service command with unsolicited response. Data: (1) port_id and (2) original_exchange_id Severity: KERN_ERR Log: LOG_NVME_DISC Action: Report the error to Broadcom Technical Support.
6206 NVMET unsol ls_req ndlp <port_id> DID <remote_port_id> xflags <flag> refcnt <reference_count> NVMe target dropped the unsolicited link service. Data: (1) port_id (2) remote_port_id (3) flag and (4) reference_count Severity: KERN_ERR Log: LOG_NODE Action: Report the error to Broadcom Technical Support.
6155 Drop NVME LS from DID <destination_port_id>: SID <source_port_id> OXID <original_exchange_id> NVMe<T/I> handler failed <return_code> NVMe target dropped the unsolicited link service. Data: (1) destination_port_id (2) source_port_id (3) original_exchange_id and (4) return_code Severity: KERN_ERR Log: LOG_TRACE_EVENT Action: Report the error to Broadcom Technical Support.
2640 Failed to alloc memory for Congestion Info Failed to allocate memory for congestion information. Data: None Severity: KERN_ERR Log: LOG_INIT Action: Check the system memory.
6242 Error reading Cgn Params <status code> Error reading congestion parameter. Data: (1) status_code Severity: KERN_ERR Log: LOG_CGN_MGMT LOG_INIT Action: Check the congestion parameter values. Report the error to Broadcom Technical Support.
6243 CGN Event empty object Congestion event received with empty object. Data: None Severity: KERN_ERR Log: LOG_CGN_MGMT LOG_INIT Action: None

Table 10: FC Error Log Messages and Their Descriptions (Continued)

2644 Failed to alloc memory for RX Monitor Buffer Failed to allocate memory for receive monitor buffer. Data: None Severity: KERN_ERR Log: LOG_INIT Action: Check the system memory.
2904 Firmware Dump Image Present on Adapter Firmware dump image is present on the adapter. Data: None Severity: KERN_ERR Log: LOG_SLI Action: Report the error to Broadcom Technical Support.
2352 Mailbox command <command> (<subsystem>/<opcode>) sli_flag <sli_flag> could not complete Mailbox command did not complete. Data: (1) command (2) subsystem (3) opcode and (4) sli_flag Severity: KERN_ERR Log: LOG_TRACE_EVENT Action: Report the error to Broadcom Technical Support.
3268 WQ[<wqe_queue_id>] failed to memmap pci barset: <dpp_barset> Failed to memory map PCI barset. Data: (1) wqe_queue_id and (2) dpp_barset Severity: KERN_ERR Log: LOG_TRACE_EVENT Action: Report the error to Broadcom Technical Support.
2122 <function_name>: Failed to get nlp ref Invalid mailbox command received. Data: None Severity: KERN_ERR Log: LOG_TRACE_EVENT Action: None
6207 Invalid command <command_code> Invalid mailbox command received. Data: (1) command_code Severity: KERN_ERR Log: LOG_TRACE_EVENT Action: None
6209 Buffer List Empty. Job entry buffer is empty. Data: None Severity: KERN_ERR Log: LOG_ELS Action: If this issue persists, report the error to Broadcom Technical Support.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

6208 Invalid Command : <wqe_command_code> Invalid WQE command. Data: (1) wqe_command_code Severity: KERN_ERR Log: LOG_TRACE_EVENT Action: None
0933 <function_name>: Register FC login error: <status> Register FC login mailbox command failed. Data: (1) status Severity: KERN_ERR Log: LOG_TRACE_EVENT Action: Report the error to Broadcom Technical Support.
2910 Async FC Trunking Event <event_tag> - Speed<speed> Logical speed:<logical_speed> port0: <port0_status> port1: <port1_status> port2: <port2_status> port3: <port3_status> Asynchronous FC trunking event received. Data: (1) event_tag (2) speed (3) logical_speed (4) port0_status (5) port1_status (6) port2_status and (7) port3_status Severity: KERN_ERR Log: LOG_LINK_EVENT Action: None
3296 LPFC_SLI_EVENT_TYPE_MISCONFIGURED event: Invalid link <link_number> Link event received with an invalid link number. Data: (1) link_number Severity: KERN_ERR Log: LOG_TRACE_EVENT Action: If this issue persists, report the error to Broadcom Technical Support.
6225 CMF mode param out of range: <cgn_param_mode> CMF mode parameter value is out of range. Data: (1) cgn_param_mode Severity: KERN_ERR Log: LOG_CGN_MGMT Action: If this issue persists, report the error to Broadcom Technical Support.
2999 Get SLI4 Parameters failed Failed to get SLI-4 parameter. Data: None Severity: KERN_ERR Log: LOG_INIT Action: Report the error to Broadcom Technical Support.
3331 Failed allocating per cpu cgn stats Failed to allocate per CPU congestion states. Data: None Severity: KERN_ERR Log: LOG_TRACE_EVENT Action: If this issue persists, report the error to Broadcom Technical Support.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

0936 <function_name>: no memory for reg_login Data: <port_destination_id> <port_state> <flag> <remote_port_indicator> No memory available for register login. Data: (1) port_destination_id (2) port_state (3) flag and (4) remote_port_indicator Severity: KERN_ERR Log: LOG_NODE Action: Check the system memory.
0938 <function_name>: failed to format reg_login Data: <port_destination_id> <port_state> <flag> <remote_port_indicator> Failed to format register login. Data: (1) port_destination_id (2) port_state (3) flag and (4) remote_port_indicator Severity: KERN_ERR Log: LOG_NODE Action: Report the error to Broadcom Technical Support.
0937 <function_name>: Failed to reg fc node, rc <return_code> Failed to register FC node. Data: (1) return_code Severity: KERN_ERR Log: LOG_NODE Action: Report the error to Broadcom Technical Support.
3334 Delay fc port discovery for <timeout_value> secs FC port discovery delayed for timeout value secs. Data: (1) timeout_value Severity: KERN_ERR Log: LOG_TRACE_EVENT Action: None
6529 QFPA failed with status <status_code> <word_value> VMID failed. Data: (1) status_code and (2) word_value Severity: KERN_ERR Log: LOG_SLI Action: Report the error to Broadcom Technical Support.
9032 Null lpfc_cmd pointer. No release, skip completion Work queue command is received as null. Data: None Severity: KERN_ERR Log: LOG_TRACE_EVENT Action: None
0725 SCSI layer TGTRST failed & LOGO TMO (<target_id>, <lun_id>) return <status> Failed to issue target reset and logo timed out. Data: (1) target_id (2) lun_id and (3) status Severity: KERN_ERR Log: LOG_FCP Action: Report the error to Broadcom Technical Support.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

6051 NVME LS REQ: Bad NDLP <port_id>, Failing LS Req NVMe link server request failed. Data: (1) port_id Severity: KERN_ERR Log: LOG_TRACE_EVENT Action: Report the error to Broadcom Technical Support.
2724 GET CGNBUF error: <error_code> Error in getting congestion buffer. Data: (1) error_code Severity: KERN_ERR Log: LOG_LIBDFC Action: If this issue persists, report the error to Broadcom Technical Support.
6460 E2E FDMI RPA failure End-to-end management FDMI registration failed. Data: None Severity: KERN_ERR Log: LOG_DISCOVERY LOG_ELS Action: Report the error to Broadcom Technical Support.
4576 PLOGI ACC fails pt2pt discovery: DID <destination_id> Data: <return_code> PLOGI accept failed for point-to-point discovery. Data: (1) destination_id and (2) return_code Severity: KERN_ERR Log: LOG_TRACE_EVENT Action: If this issue persists, report the error to Broadcom Technical Support.
6324 IO oxid <original_exchange_id> aborted I/O original exchange ID is aborted. Data: (1) original_exchange_id Severity: KERN_ERR Log: LOG_TRACE_EVENT Action: Report the error to Broadcom Technical Support.
0321 Rsp Ring <ring_no> error: Job Data: <word0_value> <total_data_placed> <comp_paramter> <word3_value> Error in response ring. Data: (1) ring_no (2) word0_value (3) total_data_placed (4) comp_value and (5) word3_value Severity: KERN_WARNING Log: LOG_SLI Action: If this issue persists, report the error to Broadcom Technical Support.
6219 Enable CMF Mailbox <command> (<subsystem>/<opcode>) failed, rc:<return_code> dd:<cmf_value> CMF mailbox command failed to enable CMF. Data: (1) command (2) subsystem (3) opcode (4) return_code and (5) cmf_value Severity: KERN_WARNING Log: LOG_CGN_MGMT LOG_INIT Action: Report the error to Broadcom Technical Support.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

4199 Firmware write failed: image incompatible with flash <flash_id> Firmware write failure due to image incompatibility with flash. Data: flash_id Severity: KERN_WARNING Log: LOG_MBOX LOG_SLI Action: Report the error to Broadcom Technical Support.
4200 Firmware write failed: image incompatible with ASIC architecture <asic_revision> Firmware write failed due to image incompatibility with ASIC architecture. Data: (1) asic_revision Severity: KERN_WARNING Log: LOG_MBOX LOG_SLI Action: Report the error to Broadcom Technical Support.
4210 Firmware write failed: add_status_2 <status> Firmware write failed. Data: (1) status Severity: KERN_WARNING Log: LOG_MBOX LOG_SLI Action: Report the error to Broadcom Technical Support.
6329 Invalid CMF mode Invalid CMF mode. Data: None Severity: KERN_WARNING Log: LOG_CGN_MGMT Action: Try to set CMF with one of the valid CMF mode. If issue persist report the error to Broadcom Technical Support.
6461 Truncated TLV hdr on Diagnostic descriptor[<descriptor_count>] Invalid Exchange Diagnostic Capabilities payload received. Data: (1) descriptor_count Severity: KERN_WARNING Log: LOG_CGN_MGMT Action: Report the error to Broadcom Technical Support.
6462 Truncated Link Fault Diagnostic descriptor[<descriptor_count>]: <bytes_remain> vs <descriptor_length> <size_of_congestion_signal_descriptor> Invalid link fault Exchange Diagnostic Capabilities payload received. Data: (1) descriptor_count (2) bytes_remain (3) descriptor_length and (4) size_of_congestion_signal_descriptor Severity: KERN_WARNING Log: LOG_CGN_MGMT Action: Report the error to Broadcom Technical Support.
6463 Truncated Cgn Signal Diagnostic descriptor[<descriptor_count>]: <bytes_remain> vs <descriptor_length> <size_of_congestion_signal_descriptor> Invalid Congestion Exchange Diagnostic Capabilities payload received Data: (1) descriptor_count (2) bytes_remain (3) descriptor_length (4) size_of_congestion_signal_descriptor Severity: KERN_WARNING Log: LOG_CGN_MGMT Action: Report the error to Broadcom Technical Support.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

4919 unknown Diagnostic Descriptor[<descriptor_count>]: tag <descriptor_tag_value> (<descriptor_tag_string>) Invalid Exchange Diagnostic Capabilities payload received. Data: (1) descriptor_count (2) descriptor_tag_value and (3) descriptor_tag_string Severity: KERN_WARNING Log: LOG_CGN_MGMT Action: Report the error to Broadcom Technical Support.
4202 EDC rsp error - sending RDF for FPIN only. Error in Exchange Diagnostic Capabilities response. Data: None Severity: KERN_WARNING Log: LOG_ELS LOG_CGN_MGMT Action: Report the error to Broadcom Technical Support.
6464 Truncated TLV hdr on Diagnostic descriptor[<descriptor_count>] Invalid Exchange Diagnostic Capabilities payload received. Data: (1) descriptor_count Severity: KERN_WARNING Log: LOG_CGN_MGMT Action: Report the error to Broadcom Technical Support.
6465 Truncated Link Fault Diagnostic descriptor[<descriptor_count>]: <bytes_remain> vs <descriptor_length> <size_of_congestion_signal_descriptor> Invalid link fault Exchange Diagnostic Capabilities payload received. Data: (1) descriptor_count (2) bytes_remain (3) descriptor_length and (4) size_of_congestion_signal_descriptor Severity: KERN_WARNING Log: LOG_CGN_MGMT Action: Report the error to Broadcom Technical Support.
6466 Truncated cgn signal Diagnostic descriptor[<descriptor_count>]: <bytes_remain> vs <descriptor_length> <size_of_congestion_signal_descriptor> Invalid Congestion Exchange Diagnostic Capabilities payload received. Data: (1) descriptor_count (2) bytes_remain (3) descriptor_length and (4) size_of_congestion_signal_descriptor Severity: KERN_WARNING Log: LOG_CGN_MGMT Action: Report the error to Broadcom Technical Support.
6467 unknown Diagnostic Descriptor[<descriptor_count>]: tag <descriptor_tag> (<descriptor_tag_string>) Invalid Exchange Diagnostic Capabilities payload received. Data: (1) descriptor_count (2) descriptor_tag_value and (3) descriptor_tag_string Severity: KERN_WARNING Log: LOG_CGN_MGMT Action: Report the error to Broadcom Technical Support.
4671 Bad ELS FPIN length <length>: <fpin_length> Bad ELS FPIN length. Data: (1) length and (2) fpin_length Severity: KERN_WARNING Log: LOG_CGN_MGMT Action: Report the error to Broadcom Technical Support.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

4672 Bad FPIN descriptor TLV length <count>: <length> <fpin_length> <descriptor_tag_string> BAD FPIN descriptor TLV length. Data: (1) count (2) length (3) fpin_length and (4) descriptor_tag_string Severity: KERN_WARNING Log: LOG_CGN_MGMT Action: Report the error to Broadcom Technical Support.
1115 Received RDF from invalid DID <destination_id> Received register diagnostic function from an invalid destination port id. Data: (1) destination_id Severity: KERN_WARNING Log: LOG_CGN_MGMT Action: Report the error to Broadcom Technical Support.
3277 QFPA LS_RJT <data0_value> <data1_value> VMID registration rejected by the switch. Data: (1) data0_value and (2) data1_value Severity: KERN_WARNING Log: LOG_SLI Action: Report this issue to switch vendor.
4532 UVEMLS_RJT <data0_value> <data1_value> VMID registration rejected by the switch. Data: (1) data0_value and (2) data1_value Severity: KERN_WARNING Log: LOG_SLI Action: Report this issue to switch vendor.
4533 UVEMLS_RJT <ulp_status>: <ulp_word4_value>\ VMID error by the switch. Data: (1) ulp_status and (2) ulp_word4_value Severity: KERN_WARNING Log: LOG_SLI Action: Report this issue to switch vendor.
0711 detected queue full - lun queue depth adjusted to <queue_depth> Detected queue full. The LUN queue depth is adjusted to a new value. Data: (1) queue_depth Severity: KERN_WARNING Log: LOG_FCP Action: None
9035 Fabric/Node busy FCP cmd <command> failed <> device_id <> lun <> status: <status> result: <result> sid: <source_port_id> did: <destination_port_id> oxid: <original_exchange_id> Data: <data1> <data2> <data3> FCP command failure due to Fabric node being busy. Data: (1) command (2) device_id (3) lun (4) status (5) result (6) source_port_id (7) destination_port_id (8) original_exchange_id (9) data1 (10) data2 and (11) data3 Severity: KERN_WARNING Log: LOG_FCP Action: If this issue persists, report the error to Broadcom Technical Support.

Table 10: FC Error Log Messages and Their Descriptions (Continued)

9040 non-zero BGSTAT on unprotected cmd Error in BlockGuard command. Data: None Severity: KERN_WARNING Log: LOG_BG Action: Check the target configuration. the error to Broadcom Technical Support.
9036 Local Reject FCP cmd <command> failed <> device_id>/<lun>> status: <status> result: <result> sid: <source_port_id> did: <destination_port_id> oxid: <original_exchange_id> Data: <data1> <data2> <data3> FCP command rejected. Data: (1) command (2) device_id (3) lun (4) status (5) result (6) source_port_id (7) destination_port_id (8) original_exchange_id (9) data1 (10) data2 (11) data3 Severity: KERN_WARNING Log: LOG_FCP Action: If this issue persists, report the error to Broadcom Technical Support.

Appendix A: FCP Driver BlockGuard Functionality

This appendix describes how to enable the BlockGuard feature and set FCP driver module parameters.

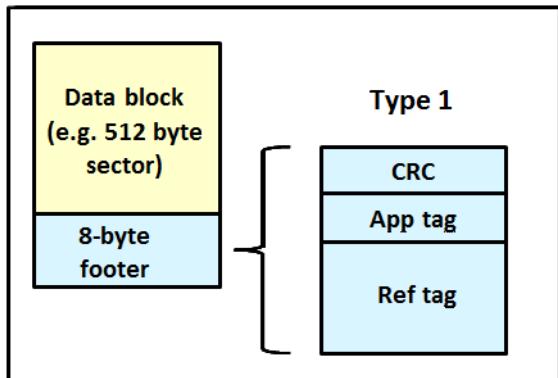
A.1 Overview

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 DIF defined in the ANSI T10 standard.

The Emulex FCP driver supports T10 DIF Type 1. In the Type 1 implementation, the 8-byte DIF consists of a Ref tag (or 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 an 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 FCP driver.

Figure 1: Data Block Showing a 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 LBA. 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 the data block and DIF on the physical media. If the CRC does not match the data, the data was corrupted during the write. A Check Condition is returned 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, along with the data block, is returned 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 the T10 DIF. This setting 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 translates the IP checksum into a CRC, or vice versa, on data being sent and received on the wire. The CRC is called a DIF protection type, and the IP checksum is referred to as a DIX protection type.

A.2 Enabling BlockGuard

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

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

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

```
options lpfc lpfc_enable_bg=1
```

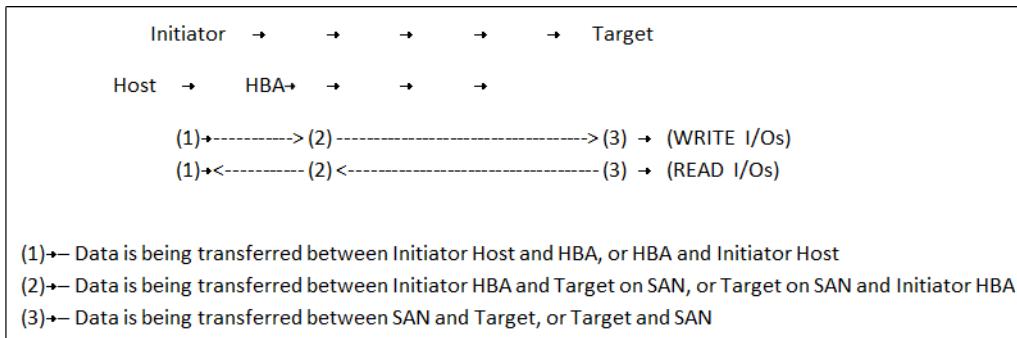
Additional module parameters can be added to this line, separated by spaces.

A.3 SCSI Command Support – SCSI Operation Codes

When there are both READ and WRITE requests, the CDB passed to the adapter from the initiator host has a read protect/write protect (RDPRETECT/WRPROTECT) field that indicates to the target whether to perform data integrity verification. It also indicates whether to transfer protection data between the 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, always prepares a CDB with the appropriate RDPRETECT/WRPROTECT information, depending on target format and capabilities. The request also includes information about the protection type with which the target has been formatted.

In addition, the initiator host provides 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 the Initiator Host and the Adapter



The initiator operations are listed in the following table.

Table 11: 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 a CRC or an IP checksum.
READ_PASS	Protected	Protected	Reads the data and protection data from the target on the SAN. The adapter verifies data integrity and transfers both data and protection data to the initiator host. The adapter can convert the protection data Guard tag from a CRC to an IP checksum.
READ_STRIP	Unprotected	Protected	Reads the data and protection data from the target. The adapter verifies data integrity, discards protection data, and transfers only the data to the initiator host. It does not send the protection data to the initiator host. Protection data is sent only on the SAN.
WRITE_INSERT	Unprotected	Protected	Transfers the data from the initiator host. The adapter then generates protection data and writes both the data and protection data to the target. Protection data is sent only on the SAN.
WRITE_PASS	Protected	Protected	Transfers the data and protection data from the initiator host to the adapter. The adapter verifies protection data and writes both data and protection data to the target on the SAN. The adapter can convert the protection data Guard tag from an IP checksum to a CRC.
WRITE_STRIP	Protected	Unprotected	Transfers data and protection data from the initiator host. The adapter verifies the data integrity, discards the protection data, and writes the data only to the target. No protection data is sent on the SAN.

A.4 FCP Driver Module Parameters

The FCP driver has two module parameters: `lpfc_prot_mask` and `lpfc_prot_guard`. Using these parameters, you can control which DIF capabilities the FCP 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.

A.4.1 lpfc_prot_mask

The `lpfc_prot_mask` 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 typically uses 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 12: `lpfc_prot_mask` Protection Types

Flag	Value	Indicates	Description
<code>SHOST_DIF_TYPE1_PROTECTION</code>	1	Adapter supports T10 DIF Type 1	Adapter-to-target Type 1 protection
<code>SHOST_DIX_TYPE0_PROTECTION</code>	8	Adapter supports DIX Type 0	Host-to-adapter protection only
<code>SHOST_DIX_TYPE1_PROTECTION</code>	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 13: Protection Data Placement for Supported Profiles

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

A.4.2 lpfc_prot_guard

The `lpfc_prot_guard` parameter correlates to the `SCSI_host_guard_type` of the Linux kernel. This parameter specifies the type of CRC that the Linux operating system passes to the FCP driver. The following table shows the two guard types: CRC and IP checksum, with values of `0x1` and `0x2`, respectively.

Table 14: `lpfc_prot_guard` Guard Types

Flag	Value	Indicates
<code>SHOST_DIX_GUARD_CRC</code>	<code>0x1</code>	Adapter supports T10 DIF CRC
<code>SHOST_DIX_GUARD_IP</code>	<code>0x2</code>	Adapter supports both T10 DIF CRC and IP checksum

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 the 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 typically uses an IP checksum as the method for computing the protection data guard tag because it uses less CPU overhead.

Appendix B: Native NVMe Multipathing

Two multipathing options can be used with NVMe storage: DM and native NVMe multipathing.

DM is a general block-device multipathing framework that can work with devices of any protocol. DM supports the addition of hardware-specific handlers to help manage the storage device and its path management. DM has been used extensively with SCSI storage. DM can be used in its generic form with NVMe storage. However, DM currently does not support advanced NVMe features, such as ANA support. Refer to the operating system release documentation for the level of DM and NVMe storage integration supported.

Native NVMe multipathing is a kernel-level feature that can be enabled to add an NVMe-specific multipath module to the kernel, which automatically detects subsystems and manages paths to the subsystems. The NVMe-specific module supports the NVMe ANA feature. Inclusion and enablement of the Native NVMe Multipath module is operating-system-specific. Refer to the operating system release documentation for the level of native NVMe multipathing supported.

Native NVMe multipathing is included, but is not enabled by default, in SLES 12 SP4 and SLES 15 SP2 and later operating systems. To enable native NVMe multipathing on SLES 12 SP4 and SLES 15 SP2 and later initiators, add `nvme-core.multipath=Y` in the kernel command line, either in the grub boot loader or during bootup.

To verify whether native NVMe multipathing is enabled, run the following command:

```
# cat /sys/module/nvme_core/parameters/multipath
```

If native NVMe multipathing is enabled, the output displayed is `Y`. If native NVMe multipathing is not enabled, the output displayed is `N`.

When enabled, native NVMe multipathing changes the manner in which NVMe devices are presented and related in `/sys/class/nvme`, especially in cases where multiple paths to target namespaces exist. These changes significantly affect the views generated by the kernel and system utilities, such as `nvme-cli`. This appendix describes these new views.

When native NVMe multipathing is not enabled, if an NVMe device is connected to a subsystem, a controller element is created. If the subsystem is connected using multiple paths, multiple controller elements are created. Each controller element is considered a unique and separate storage entity, even though the subsystem and namespaces that the controller could access might be the same.

When native NVMe multipathing is enabled, controller elements are still created for each connection to a subsystem. However, when a controller is created, the subsystem is compared to the list of subsystems that have already been found in the system, and one of the following actions occurs:

- If the subsystem is not found, a new subsystem element is created and added to the system list, with the controller linked to it as a path.
- If the subsystem is found, the controller is linked to the subsystem as an additional path.

Namespaces are elements of the NVMe subsystem and not of the NVMe controller. A controller acts as a communication port to a subsystem. If a subsystem contains two controllers (communication ports), and namespace 1 is changed by controller A, controller B sees this change as an action on namespace 1 because the namespace is a singular component of the subsystem.

When native NVMe multipathing is not enabled, because each controller is seen as unique, namespaces for the subsystem are created for every controller on which the namespace is found. Processing of I/O requests by the device name is handled only by the controller, and they are blindly passed on to the storage device. If concurrent actions occur on namespaces on different controllers connected to the same subsystem, the application or administrator must realize that the namespace is a single entity behind the controllers, and the application or administrator must access the two controllers accordingly, so as to provide data coherency and to avoid data corruption.

When native NVMe multipathing is enabled, namespace device names are created when the device initially connects to the subsystem, and the controller is the first path to the subsystem. The device name reflects the subsystem of which the controller is a part. If multiple controllers connect to the same subsystem, no additional namespace device names are created because the controllers are simply paths that provide access to the same storage device. To maintain compatibility with utilities that are operating on NVMe devices on which native NVMe multipathing is not enabled, the name format used for the namespace device name has been kept the same. This can present some confusion because the namespace device name (for example, `/dev/nvme0n1`) contains an `nvme<x>` prefix that, when native NVMe multipathing is not enabled, is the controller name, but on systems on which native NVMe multipathing is enabled, is a prefix for the subsystem instance. In addition, for compatibility with management of NVMe controllers, the device name used for NVMe controllers (for example, `/dev/nvme2`) was left the same in systems that have native NVMe multipathing enabled. So, when native NVMe multipathing is enabled, you might see both a `/dev/nvme3` controller name and a `/dev/nvme3n1` namespace name, but no correlation exists between the `nvme3` portion of both names. In summary, when native NVMe multipathing is enabled, the name `/dev/nvme<A>n<x>` means *namespace number <x> on subsystem instance <A>*, whereas the name `/dev/nvme<A>` means *controller instance <A>*. The only way that the two names correlate is if controller instance `<A>` actually connects to subsystem instance `<A>` and is therefore a path to subsystem instance `<A>`.

Beyond the device name difference, the other place that this change in behavior is seen is in system utilities, such as `nvme-cli`. For example, the `nvme list` command lists all NVMe namespace devices.

NOTE: Two similar commands can be used to list all NVMe namespace devices:

- `nvme list` (with a space) is a native NVMe CLI command.
- `nvme-list-linux` (with a hyphen) is an Emulex HBA Manager CLI command.

This appendix refers to the `nvme list` command.

Thus, if two controllers are connected to a single subsystem with a single namespace, information similar to the following is displayed when native NVMe multipathing is not enabled:

Node	SN	Model	Namespace	Usage	Format	FW	Rev
<code>/dev/nvme0n1</code>	0000000000000000	Linux	1	268.44 GB / 268.44 GB	512 B + 0 B	4.4.131-	
<code>/dev/nvme1n1</code>	0000000000000000	Linux	1	268.44 GB / 268.44 GB	512 B + 0 B	4.4.131-	

On the other hand, information similar to the following is displayed when native NVMe multipathing is enabled:

Node	SN	Model	Namespace	Usage	Format	FW	Rev
<code>/dev/nvme0n1</code>	0000000000000000	Linux	1	268.44 GB / 268.44 GB	512 B + 0 B	4.4.131-	

When native NVMe multipathing is enabled, you can display the list of subsystems that have been detected by the system using the `nvme list-subsy`s command. For example:

```
nvme list-subsy
nvme-subsy0 - NQN=nqn.2014-08.org.nvmexpress:NVMf:uuid:19d48a6d-3644-4f65-8046-80bd88fe45a6
 \
 +- nvme0 fc traddr=nn-0x20000090fa942e16:pn-0x10000090fa942e16 host_traddr=nn-0x20000109b346f1f:pn-
0x10000109b346f1f
 +- nvme3 fc traddr=nn-0x20000090fa942e16:pn-0x10000090fa942e16 host_traddr=nn-0x20000109b346f20:pn-
0x10000109b346f20
nvme-subsy1 - NQN=nqn.2014-08.org.nvmexpress:NVMf:uuid:2c6d862f-b77e-43b1-b345-7f990125088e
 \
 +- nvme1 fc traddr=nn-0x20000090fa942e16:pn-0x10000090fa942e16 host_traddr=nn-0x20000109b346f1f:pn-
0x10000109b346f1f
 +- nvme6 fc traddr=nn-0x20000090fa942e16:pn-0x10000090fa942e16 host_traddr=nn-0x20000109b346f20:pn-
0x10000109b346f20
nvme-subsy2 - NQN=nqn.2014-08.org.nvmexpress:NVMf:uuid:4bbbd877-6de2-4a67-ad56-e0333e2437b5
 \
 +- nvme2 fc traddr=nn-0x20000090fa942e16:pn-0x10000090fa942e16 host_traddr=nn-0x20000109b346f1f:pn-
0x10000109b346f1f
 +- nvme7 fc traddr=nn-0x20000090fa942e16:pn-0x10000090fa942e16 host_traddr=nn-0x20000109b346f20:pn-
0x10000109b346f20
nvme-subsy3 - NQN=nqn.2014-08.org.nvmexpress:NVMf:uuid:c8da7263-fefe-4563-a14c-a9f0362dc03
 \
 +- nvme5 fc traddr=nn-0x20000090fa942e16:pn-0x10000090fa942e16 host_traddr=nn-0x20000109b346f1f:pn-
0x10000109b346f1f
 +- nvme9 fc traddr=nn-0x20000090fa942e16:pn-0x10000090fa942e16 host_traddr=nn-0x20000109b346f20:pn-
0x10000109b346f20
nvme-subsy4 - NQN=nqn.2014-08.org.nvmexpress:NVMf:uuid:5f1d7788-fab1-45cf-af9c-d30dc9a9ce89
 \
 +- nvme4 fc traddr=nn-0x20000090fa942e16:pn-0x10000090fa942e16 host_traddr=nn-0x20000109b346f1f:pn-
0x10000109b346f1f
 +- nvme8 fc traddr=nn-0x20000090fa942e16:pn-0x10000090fa942e16 host_traddr=nn-0x20000109b346f20:pn-
0x10000109b346f20
```

In the previous example, five subsystems are displayed. Each subsystem contains multiple controllers that represent multiple paths. Note that the `nvme-subsy`s numbering is disassociated from the component controller identification, as demonstrated by `nvme-subsy3` and controllers `nvme5` and `nvme9`.

In the preceding example configuration, each of the five subsystems contains four namespaces. The `nvme list-subsy`s view shows two controllers. When native NVMe multipathing is not enabled, each namespace and path are uniquely presented to the operating system. When native NVMe multipathing is enabled, each namespace is displayed once.

nvme list	Node	SN	Model	Namespace	Usage	Format	FW	Rev
	/dev/nvme0n1	0000000000000000	Linux	1	268.44 GB / 268.44 GB	512 B + 0 B	4.4.131-	
	/dev/nvme0n2	0000000000000000	Linux	1	268.44 GB / 268.44 GB	512 B + 0 B	4.4.131-	
	/dev/nvme0n3	0000000000000000	Linux	1	268.44 GB / 268.44 GB	512 B + 0 B	4.4.131-	
	/dev/nvme0n4	0000000000000000	Linux	1	268.44 GB / 268.44 GB	512 B + 0 B	4.4.131-	
	/dev/nvme1n1	0000000000000000	Linux	1	268.44 GB / 268.44 GB	512 B + 0 B	4.4.131-	
	/dev/nvme1n2	0000000000000000	Linux	1	268.44 GB / 268.44 GB	512 B + 0 B	4.4.131-	
	/dev/nvme1n3	0000000000000000	Linux	1	268.44 GB / 268.44 GB	512 B + 0 B	4.4.131-	
	/dev/nvme1n4	0000000000000000	Linux	1	268.44 GB / 268.44 GB	512 B + 0 B	4.4.131-	

Node	SN	Model	Namespace	Usage	Format	FW	Rev
/dev/nvme2n1	00000000000000000000	Linux	1	268.44 GB / 268.44 GB	512 B + 0 B	4.4.131-	
/dev/nvme2n2	00000000000000000000	Linux	1	268.44 GB / 268.44 GB	512 B + 0 B	4.4.131-	
/dev/nvme2n3	00000000000000000000	Linux	1	268.44 GB / 268.44 GB	512 B + 0 B	4.4.131-	
/dev/nvme2n4	00000000000000000000	Linux	1	268.44 GB / 268.44 GB	512 B + 0 B	4.4.131-	
/dev/nvme3n1	00000000000000000000	Linux	1	268.44 GB / 268.44 GB	512 B + 0 B	4.4.131-	
/dev/nvme3n2	00000000000000000000	Linux	1	268.44 GB / 268.44 GB	512 B + 0 B	4.4.131-	
/dev/nvme3n3	00000000000000000000	Linux	1	268.44 GB / 268.44 GB	512 B + 0 B	4.4.131-	
/dev/nvme3n4	00000000000000000000	Linux	1	268.44 GB / 268.44 GB	512 B + 0 B	4.4.131-	
/dev/nvme4n1	00000000000000000000	Linux	1	268.44 GB / 268.44 GB	512 B + 0 B	4.4.131-	
/dev/nvme4n2	00000000000000000000	Linux	1	268.44 GB / 268.44 GB	512 B + 0 B	4.4.131-	
/dev/nvme4n3	00000000000000000000	Linux	1	268.44 GB / 268.44 GB	512 B + 0 B	4.4.131-	
/dev/nvme4n4	00000000000000000000	Linux	1	268.44 GB / 268.44 GB	512 B + 0 B	4.4.131-	
/dev/nvme5n1	00000000000000000000	Linux	1	268.44 GB / 268.44 GB	512 B + 0 B	4.4.131-	
/dev/nvme5n2	00000000000000000000	Linux	1	268.44 GB / 268.44 GB	512 B + 0 B	4.4.131-	
/dev/nvme5n3	00000000000000000000	Linux	1	268.44 GB / 268.44 GB	512 B + 0 B	4.4.131-	
/dev/nvme5n4	00000000000000000000	Linux	1	268.44 GB / 268.44 GB	512 B + 0 B	4.4.131-	

Note that the naming includes the subsystem and namespace numbering in the form

`/dev/nvme{subsystem instance}n{namespace}`. The controller (or path) is not shown.

You can also view the device nodes using the `ls` command:

```
ls -l /dev/nvme*
0 crw----- 1 root root 10, 57 Jun 10 16:02 /dev/nvme-fabrics
0 crw----- 1 root root 246, 0 Jun 10 16:02 /dev/nvme0
0 brw-rw---- 1 root disk 259, 231 Jun 10 16:02 /dev/nvme0n1
0 brw-rw---- 1 root disk 259, 271 Jun 10 16:02 /dev/nvme0n2
0 brw-rw---- 1 root disk 259, 303 Jun 10 16:02 /dev/nvme0n3
0 brw-rw---- 1 root disk 259, 311 Jun 10 16:02 /dev/nvme0n4
0 crw----- 1 root root 246, 1 Jun 10 16:02 /dev/nvme1
0 brw-rw---- 1 root disk 259, 323 Jun 10 16:02 /dev/nvme1n1
0 brw-rw---- 1 root disk 259, 343 Jun 10 16:02 /dev/nvme1n2
0 brw-rw---- 1 root disk 259, 347 Jun 10 16:02 /dev/nvme1n3
0 brw-rw---- 1 root disk 259, 351 Jun 10 16:02 /dev/nvme1n4
0 crw----- 1 root root 246, 2 Jun 10 16:02 /dev/nvme2
0 brw-rw---- 1 root disk 259, 359 Jun 10 16:02 /dev/nvme2n1
0 brw-rw---- 1 root disk 259, 363 Jun 10 16:02 /dev/nvme2n2
0 brw-rw---- 1 root disk 259, 381 Jun 10 16:02 /dev/nvme2n3
0 brw-rw---- 1 root disk 259, 383 Jun 10 16:02 /dev/nvme2n4
0 crw----- 1 root root 246, 3 Jun 10 16:02 /dev/nvme3
0 brw-rw---- 1 root disk 259, 397 Jun 10 16:02 /dev/nvme3n1
0 brw-rw---- 1 root disk 259, 399 Jun 10 16:02 /dev/nvme3n2
0 brw-rw---- 1 root disk 259, 401 Jun 10 16:02 /dev/nvme3n3
0 brw-rw---- 1 root disk 259, 403 Jun 10 16:02 /dev/nvme3n4
0 crw----- 1 root root 246, 4 Jun 10 16:02 /dev/nvme4
0 brw-rw---- 1 root disk 259, 385 Jun 10 16:02 /dev/nvme4n1
0 brw-rw---- 1 root disk 259, 387 Jun 10 16:02 /dev/nvme4n2
0 brw-rw---- 1 root disk 259, 389 Jun 10 16:02 /dev/nvme4n3
0 brw-rw---- 1 root disk 259, 391 Jun 10 16:02 /dev/nvme4n4
0 crw----- 1 root root 246, 5 Jun 10 16:02 /dev/nvme5
0 crw----- 1 root root 246, 6 Jun 10 16:02 /dev/nvme6
```

```
0 crw----- 1 root root 246,    7 Jun 10 16:02 /dev/nvme7
0 crw----- 1 root root 246,    8 Jun 10 16:02 /dev/nvme8
0 crw----- 1 root root 246,    9 Jun 10 16:02 /dev/nvme9
```

This example demonstrates that one device node is created per namespace, with the following form:

```
/dev/nvme{subsystem instance}n{namespace}
```

Additionally, nodes are created for the controller instances `/dev/nvme{controller instance}`, but it is important to note that the subsystem instance and the controller instance are not correlated, as demonstrated by the previous `nvme list-subsy` example.

You can also list all of the block devices that are NVMe devices using the `ls` command. The file names are displayed as symbolic links to the actual block devices.

```
ls -l /sys/block/nvme*
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme0c222n1 -> ../../devices/virtual/nvme-fabrics/ctl/
nvme0/nvme0c222n1
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme0c222n2 -> ../../devices/virtual/nvme-fabrics/ctl/
nvme0/nvme0c222n2
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme0c222n3 -> ../../devices/virtual/nvme-fabrics/ctl/
nvme0/nvme0c222n3
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme0c222n4 -> ../../devices/virtual/nvme-fabrics/ctl/
nvme0/nvme0c222n4
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme0c227n1 -> ../../devices/virtual/nvme-fabrics/ctl/
nvme3/nvme0c227n1
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme0c227n2 -> ../../devices/virtual/nvme-fabrics/ctl/
nvme3/nvme0c227n2
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme0c227n3 -> ../../devices/virtual/nvme-fabrics/ctl/
nvme3/nvme0c227n3
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme0c227n4 -> ../../devices/virtual/nvme-fabrics/ctl/
nvme3/nvme0c227n4
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme0n1 -> ../../devices/virtual/nvme-subsystem/nvme-
subsys0/nvme0n1
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme0n2 -> ../../devices/virtual/nvme-subsystem/nvme-
subsys0/nvme0n2
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme0n3 -> ../../devices/virtual/nvme-subsystem/nvme-
subsys0/nvme0n3
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme0n4 -> ../../devices/virtual/nvme-subsystem/nvme-
subsys0/nvme0n4
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme1c223n1 -> ../../devices/virtual/nvme-fabrics/ctl/
nvme1/nvme1c223n1
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme1c223n2 -> ../../devices/virtual/nvme-fabrics/ctl/
nvme1/nvme1c223n2
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme1c223n3 -> ../../devices/virtual/nvme-fabrics/ctl/
nvme1/nvme1c223n3
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme1c223n4 -> ../../devices/virtual/nvme-fabrics/ctl/
nvme1/nvme1c223n4
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme1c229n1 -> ../../devices/virtual/nvme-fabrics/ctl/
nvme6/nvme1c229n1
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme1c229n2 -> ../../devices/virtual/nvme-fabrics/ctl/
nvme6/nvme1c229n2
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme1c229n3 -> ../../devices/virtual/nvme-fabrics/ctl/
nvme6/nvme1c229n3
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme1c229n4 -> ../../devices/virtual/nvme-fabrics/ctl/
nvme6/nvme1c229n4
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme1n1 -> ../../devices/virtual/nvme-subsystem/nvme-
subsys1/nvme1n1
```

```
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme1n2 -> ../devices/virtual/nvme-subsystem/nvme-
subsys1/nvme1n2
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme1n3 -> ../devices/virtual/nvme-subsystem/nvme-
subsys1/nvme1n3
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme1n4 -> ../devices/virtual/nvme-subsystem/nvme-
subsys1/nvme1n4
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme2c224n1 -> ../devices/virtual/nvme-fabrics/ctl/
nvme2/nvme2c224n1
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme2c224n2 -> ../devices/virtual/nvme-fabrics/ctl/
nvme2/nvme2c224n2
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme2c224n3 -> ../devices/virtual/nvme-fabrics/ctl/
nvme2/nvme2c224n3
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme2c224n4 -> ../devices/virtual/nvme-fabrics/ctl/
nvme2/nvme2c224n4
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme2c230n1 -> ../devices/virtual/nvme-fabrics/ctl/
nvme7/nvme2c230n1
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme2c230n2 -> ../devices/virtual/nvme-fabrics/ctl/
nvme7/nvme2c230n2
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme2c230n3 -> ../devices/virtual/nvme-fabrics/ctl/
nvme7/nvme2c230n3
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme2c230n4 -> ../devices/virtual/nvme-fabrics/ctl/
nvme7/nvme2c230n4
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme2n1 -> ../devices/virtual/nvme-subsystem/nvme-
subsys2/nvme2n1
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme2n2 -> ../devices/virtual/nvme-subsystem/nvme-
subsys2/nvme2n2
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme2n3 -> ../devices/virtual/nvme-subsystem/nvme-
subsys2/nvme2n3
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme2n4 -> ../devices/virtual/nvme-subsystem/nvme-
subsys2/nvme2n4
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme3c228n1 -> ../devices/virtual/nvme-fabrics/ctl/
nvme5/nvme3c228n1
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme3c228n2 -> ../devices/virtual/nvme-fabrics/ctl/
nvme5/nvme3c228n2
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme3c228n3 -> ../devices/virtual/nvme-fabrics/ctl/
nvme5/nvme3c228n3
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme3c228n4 -> ../devices/virtual/nvme-fabrics/ctl/
nvme5/nvme3c228n4
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme3c232n1 -> ../devices/virtual/nvme-fabrics/ctl/
nvme9/nvme3c232n1
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme3c232n2 -> ../devices/virtual/nvme-fabrics/ctl/
nvme9/nvme3c232n2
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme3c232n3 -> ../devices/virtual/nvme-fabrics/ctl/
nvme9/nvme3c232n3
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme3c232n4 -> ../devices/virtual/nvme-fabrics/ctl/
nvme9/nvme3c232n4
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme3n1 -> ../devices/virtual/nvme-subsystem/nvme-
subsys3/nvme3n1
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme3n2 -> ../devices/virtual/nvme-subsystem/nvme-
subsys3/nvme3n2
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme3n3 -> ../devices/virtual/nvme-subsystem/nvme-
subsys3/nvme3n3
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme3n4 -> ../devices/virtual/nvme-subsystem/nvme-
subsys3/nvme3n4
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme4c226n1 -> ../devices/virtual/nvme-fabrics/ctl/
nvme4/nvme4c226n1
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme4c226n2 -> ../devices/virtual/nvme-fabrics/ctl/
nvme4/nvme4c226n2
```

```
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme4c226n3 -> ../devices/virtual/nvme-fabrics/ctl1/nvme4/nvme4c226n3
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme4c226n4 -> ../devices/virtual/nvme-fabrics/ctl1/nvme4/nvme4c226n4
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme4c231n1 -> ../devices/virtual/nvme-fabrics/ctl1/nvme8/nvme4c231n1
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme4c231n2 -> ../devices/virtual/nvme-fabrics/ctl1/nvme8/nvme4c231n2
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme4c231n3 -> ../devices/virtual/nvme-fabrics/ctl1/nvme8/nvme4c231n3
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme4c231n4 -> ../devices/virtual/nvme-fabrics/ctl1/nvme8/nvme4c231n4
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme4n1 -> ../devices/virtual/nvme-subsystem/nvme-
subsys4/nvme4n1
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme4n2 -> ../devices/virtual/nvme-subsystem/nvme-
subsys4/nvme4n2
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme4n3 -> ../devices/virtual/nvme-subsystem/nvme-
subsys4/nvme4n3
lrwxrwxrwx 1 root root 0 Jun 11 09:56 /sys/block/nvme4n4 -> ../devices/virtual/nvme-subsystem/nvme-
subsys4/nvme4n4
```

In the previous example, the symbolic links illustrate that the available namespaces are assigned a virtual controller (or path) number that is uniquely associated with the subsystem.

Appendix C: Trunking Support

Trunking enables you to combine multiple physical FC links to form a single logical link (aggregated port). The aggregated port's maximum link speed is the sum of the maximum link speeds of the individual physical links comprising the aggregated port. The actual link speed of the aggregated port depends on the states (active or non-active) of the individual physical links comprising the aggregated port.

The physical links comprising an aggregated port are referred to as lanes. Both 2-lane and 4-lane aggregated ports are supported.

The LPe3X002 adapters support 2-lane port aggregation. If 2-lane port aggregation is configured on a dual-port adapter, the two physical links are combined to form a single 2-lane aggregated port whose aggregate speed is two times the maximum speed supported, assuming both physical links are active.

For example:

- On an LPe35002-series adapter, an aggregated port comprised of two physical links running at 32Gb/s each has a potential logical (aggregate) link speed of 64Gb/s.
- On an LPe36002-series adapter, an aggregated port comprised of two physical links running at 64Gb/s each has a potential logical (aggregate) link speed of 128Gb/s.

An LPe3X004 trunking supported adapter supports both 2-lane port aggregation and 4-lane port aggregation as follows:

- If 2-lane port aggregation is configured on an LPe35004 adapter, the four physical links on the adapter are divided among two separate aggregated ports. The two lowest-numbered physical links form the first aggregated port, and the two highest-numbered physical links form the second aggregated port.
- If 4-lane port aggregation is configured on an LPe35004 adapter, all four physical links are combined to form a single 4-lane trunk whose aggregate speed is potentially 128Gb/s, assuming all four links are active, each at 32Gb/s link speed.

NOTE: The following information applies to trunking support:

- Trunking is not supported on LPe31000-series and LPe32000-series adapters. It is supported only on LPe35000-series, LPe36000-series, LPe37000-series, and LPe38000-series adapters.
- Before you enable or disable trunking on the Emulex adapter, follow the instructions from Brocade to enable or disable trunking on the switch.
- Dynamic D_Port cannot co-exist with the trunking feature. If trunking is enabled, the firmware automatically disables Dynamic D_Port.
- Trunking is not supported at 8Gb/s speeds, and the link will not come up at this speed.

Trunking can be configured using the UEFI or x86 BIOS. Trunking is a one-time configuration, and the configuration is performed through the Emulex HBA Manager application. For more information on configuring trunking, refer to the *Emulex HBA Manager Application Command Line Interface User Guide*.

When trunking is configured for an adapter, internal loopback and external loopback tests are supported. For more information on running loopback tests, refer to the *Emulex HBA Manager Application Command Line Interface User Guide*.

You can view the existing trunk links using `sysfs`.

Following is an example of an LPe35002 adapter showing 2-lane port aggregation with a single SCSI host:

```
# cat /sys/class/scsi_host/host14/link_state
Link Up - Ready
  Fabric
Trunk port 0: Link Up
Trunk port 1: Link Up
```

Following is an example of an LPe35004 adapter showing 2-lane port aggregation with a single SCSI host:

```
# cat /sys/class/scsi_host/host13/link_state
Link Up - Ready
  Fabric
Trunk port 0: Link Up
Trunk port 1: Link Up
# cat /sys/class/scsi_host/host14/link_state
Link Up - Ready
  Fabric
Trunk port 2: Link Up
Trunk port 3: Link Up
```

Following is an example of an LPe35004 adapter showing 4-lane port aggregation with a single SCSI host:

```
# cat /sys/class/scsi_host/host13/link_state
Link Up - Ready
  Fabric
Trunk port 0: Link Up
Trunk port 1: Link Up
Trunk port 2: Link Up
Trunk port 3: Link Up
```

Appendix D: FC-SP-2 Authentication (DH-CHAP) Support

Fibre Channel Security Protocol-2 (FC-SP-2) DH-CHAP is a secret-based authentication and key management protocol that uses the CHAP algorithm augmented with an optional Diffie-Hellmann algorithm with hash priority (MD5 and SHA1).

NOTE: The following information applies to DH-CHAP authentication:

- DH-CHAP is available only for physical ports, not for virtual ports.
- The authentication driver parameters are available only on local hosts and not for remote hosts.
- Boot from SAN is not supported when DH-CHAP authentication is enabled.
- DH-CHAP is not supported if Dynamic D_Port is enabled.

You can configure FC-SP-2 DH-CHAP authentication between an adapter and a switch.

FC-SP-2 DH-CHAP authentication is disabled by default. To enable DH-CHAP on Linux systems, the `lpfc_enable_auth` parameter must be passed to the driver by typing the following command:

```
insmod lpfc.ko lpfc_enable_auth=1
```

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

```
options lpfc lpfc_enable_auth=1
```

After DH-CHAP has been activated and configured, you can manually initiate authentication per adapter, and you can get, set, update, or delete a password using the Emulex HBA Manager application. For more information on manually initiating authentication, refer to the *Emulex HBA Manager Application Command Line Interface User Guide*.

Appendix E: Arm Support

Emulex HBAs support Arm servers. This appendix describes information pertaining to Arm support.

The following inbox Linux operating system distributions are the minimum versions supported on Arm servers:

- RHEL 8.8, 8.9, 8.10, 9.2, 9.3, 9.4
- OpenEuler 20.03 SP3, SP4; 22.03 SP2, SP3
- Kylin V10 SP3, SP3-2403
- UnionTechOS V20 1060E, 1070E
- Ubuntu 22.04, 24.04
- Fusion OS 22.0.4.9, 23.0.1

The following Emulex utilities are supported on Arm servers:

- Certain commands from the Emulex HBA Manager CLI. Refer to the *Emulex HBA Manager Application Command Line Interface User Guide* for details.
- The Emulex HBA Capture utility.

To update the firmware, perform one of the following tasks:

- Use the EHM CLI utility download command. For details, refer to the *Emulex HBA Manager Application Command Line Interface User Guide*.
- Use the Linux `sysfs` command. For details, see [Section 2.5, Updating the Firmware](#).

To configure driver parameters, use one of the following Linux tools. See [Chapter 3, Configuration](#), for details on FC driver parameters and how to configure them.

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

To create a vPort for NPIV, use the SCSI transport device `/sys/class/fc_host/<hostx>/vport_create` attribute, where `<hostx>` is the HBA. Echo a string formatted as `<wwpn>:<wwnn>`, where `<wwpn>` and `<wwnn>` are the WWPN and WWNN for the vPort to be created. When specifying the WWPN and WWNN, use colons only between the WWPN and WWNN as shown in the following example:

For example:

```
echo 30000000c9111111:50000000c9111111 > /sys/class/fc_host/host5/vport_create
```

To remove a vPort for NPIV, use the SCSI transport device `/sys/class/fc_host/<hostx>/vport_delete` attribute.

For example:

```
echo 30000000c9111111:50000000c9111111 > /sys/class/fc_host/host5/vport_delete
```

