



# Emulex Drivers for VMware ESXi

## User Manual

Versions 11.0 and 11.1

March 1, 2016

pub-005270

Corporate Headquarters  
San Jose, CA

Website  
[www.broadcom.com](http://www.broadcom.com)

Broadcom, the pulse logo, Connecting everything, the Connecting everything logo, Avago Technologies, and Emulex are among the trademarks of Broadcom Ltd. and/or its affiliates in the United States, certain other countries and/or the EU. Any other trademarks or trade names mentioned are the property of their respective owners.

Broadcom Ltd. 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 Ltd. is believed to be accurate and reliable. However, Broadcom Ltd. 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.

Copyright © 2003–2016 Broadcom Ltd. All Rights Reserved.

# Table of Contents

<b>Chapter 1: Introduction</b>	5
1.1 ESXi Compatibility	5
1.2 Supported Guest Operating Systems	5
1.3 Abbreviations	7
<b>Chapter 2: Installation</b>	10
2.1 ESXi 5.5 and 6.0 Overview	10
2.1.1 Migrating Driver Parameters When Upgrading to ESXi 5.5 or 6.0	10
2.1.2 Preparing to Upgrade to ESXi 5.5 or 6.0	11
2.2 Installing the FC/FCoE Driver	12
2.2.1 Installing the FC/FCoE Driver and Management Software	12
2.2.2 Uninstalling the FC/FCoE Driver	12
2.2.3 Installing the Native mode FC/FCoE Driver esxcli plugin for ESXi5.5 and ESXi6.0	12
2.3 Installing the NIC Driver	14
2.3.1 Installing the NIC Driver and Management Software	14
2.3.2 Uninstalling the NIC Driver	14
2.3.3 Installing the Native Mode Ethernet Driver esxcli Plug-in for ESXi 5.5 and 6.0	14
2.3.4 Virtualization Support	15
2.3.5 Obtaining Information on the Installed NIC Driver	15
2.4 Installing the iSCSI Driver	15
2.4.1 Installing the iSCSI Driver and Management Software	15
2.5 Updating Drivers with VMware Offline Bundle Media	16
<b>Chapter 3: Configuration</b>	17
3.1 ESXi Command Line Tool Transition	17
3.1.1 ESXi 5.1 Implementation	17
3.1.2 ESXi 5.5 and 6.0 Implementation	17
3.2 FC/FCoE Driver Configuration	19
3.2.1 FC/FCoE Driver Parameters Configuration Methods	19
3.2.2 Emulex ExpressLane Support	21
3.2.3 Dynamically Adding LUNs	22
3.2.4 Dynamically Adding Targets	22
3.2.5 FC/FCoE Driver Configuration Parameters	22
3.2.6 Creating an FC Remote Boot Disk	28
3.2.7 Managing Devices through the CIM Interface	29
3.2.8 Installing the Emulex CIM Provider	29
3.2.9 Creating, Deleting, and Displaying VPorts	29
3.2.10 Configuring VVols in ESXi 6.0	30
3.3 NIC Driver Configuration	30
3.3.1 Configuring ESXi 5.5 and 6.0 NIC Driver Parameters	31
3.3.2 Configuring ESXi 5.1 NIC Driver Parameters	32
3.3.3 NetQueue Support	32
3.3.4 Receive-Side Scaling	33
3.3.5 How an ESXi Server Creates and Names Interfaces	34
3.3.6 Enabling SR-IOV	35
3.3.7 Configuring Emulex Universal Multi-Channel	37
3.3.8 Configuring VLANs	38
3.3.9 Configuring NPar (Dell Only)	38
3.3.10 Configuring Network Heap Size in ESXi Server 5.1	39
3.3.11 Performance Tuning	40
3.4 Broadcom Emulex Drivers for VMware NSX 6.2	43
3.4.1 Module Parameters	43
3.4.2 Functionality	43
3.4.3 Usage	44

---

3.4.4 Limitations .....	44
3.4.5 Limitations Exclusive to OCe11000-Series Adapters .....	44
3.5 iSCSI Driver Configuration .....	45
3.5.1 Configuring iSCSI Driver Options .....	45
3.5.2 Interrupt Moderation Policy Settings .....	46
3.5.3 Configuring LDTO and ETO on ESXi Server .....	48
3.5.4 Multipath I/O Support .....	48
3.5.5 Reading the Driver Statistics for a Specified Port .....	49
3.5.6 Setting the Log Level .....	49
<b>Chapter 4: Troubleshooting .....</b>	<b>51</b>
4.1 Troubleshooting the FC/FCoE Driver .....	51
4.1.1 FC/FCoE Driver Log Messages .....	53
4.1.2 Extended Link Service Events (0100 to 0199) .....	56
4.1.3 Link Discovery Events (0200 to 0299) .....	59
4.1.4 Mailbox Events (0300 to 0339) .....	65
4.1.5 Initialization Events (0400 to 0599) .....	73
4.1.6 FCP Traffic History (0700 to 0799) .....	84
4.1.7 Node Table Events (0900 to 0999) .....	86
4.1.8 Security Events (1000 to 1099) .....	86
4.1.9 Miscellaneous and FCoE Events (1200 to 1299) .....	90
4.1.10 Link Events (1300 to 1399) .....	93
4.1.11 Port Setup Events (1400 to 1499) .....	94
4.1.12 IOCTL Events (1600 to 1699) .....	99
4.1.13 VPort Events (1800 to 1832) .....	99
4.1.14 Extended Link Service Events (1833 to 2800) .....	102
4.1.15 New Events .....	120
4.2 Troubleshooting the NIC Driver .....	133
4.2.1 NIC Event/Error Logging for ESXi 5.1 .....	134
4.2.2 NIC Adapter Firmware Error .....	136
4.2.3 NIC Informational Log Groups for ESXi 5.5 and 6.0 .....	136
4.3 Troubleshooting the iSCSI Driver .....	137
4.3.1 Retrieving iSCSI Error Log Codes .....	137
4.3.2 iSCSI Error Log Code Entries .....	138
4.3.3 Additional iSCSI Driver Messages .....	139
4.4 ESXi 5.5 and ESXi 6.0 Native Mode NIC Driver Troubleshooting Support .....	142
4.5 ESXi 5.5 and ESXi 6.0 Native Mode NIC Driver Support for ethtool Commands .....	143
<b>Appendix A: Configuring iSCSI Through DHCP .....</b>	<b>145</b>
A.1 Recommendations .....	145
A.1.1 Vendor-Specific Option 43 .....	145

# Chapter 1: Introduction

This product supports the Emulex® OneConnect® and LightPulse® families of adapters.

The VMware ESXi drivers support the Fibre Channel (FC), FC over Ethernet (FCoE), Ethernet (or network interface card–NIC), and Internet Small Computer System Interface (iSCSI) protocols.

This document explains how to install the VMware ESXi drivers on your system and how to configure the drivers' capabilities. Below is a partial list of configurable parameters by protocol:

- FC/FCoE
  - Adding logical unit numbers (LUNs) and targets
  - Configuring driver parameters
  - Creating an FC remote boot disk
  - Managing devices through the CIM interface
  - Working with virtual ports (vPorts)
  - Troubleshooting FC/FCoE issues
- NIC
  - Configuring NIC driver parameters
  - Improving driver performance
  - Using vmxnet emulation, including TCP segmentation offload (TSO) and jumbo frames
  - Troubleshooting NIC issues
- iSCSI
  - Configuring iSCSI driver parameters, iSCSI error handling, and link down timeout (LDTO) and extended timeout (ETO)
  - Implementing Multipath I/O (MPIO) support for non-boot and boot targets
  - Configuring iSCSI through Dynamic Host Configuration Protocol (DHCP)
  - Troubleshooting iSCSI issues

Refer to the *OneCommand Manager User Manual* and *OneCommand Manager for VMware vCenter User Manual* for complete lists of supported technology.

## 1.1 ESXi Compatibility

For a list of adapters that are compatible with this driver, see the driver Downloads page on the Avago website. For compatible firmware versions, see the Downloads page for the specific adapter.

## 1.2 Supported Guest Operating Systems

The Broadcom Emulex driver supports the following guest operating systems:

- CentOS 5.x
- CentOS 6.x
- RHEL 5.9, 5.10, and 5.11
- RHEL 6.4, 6.5, and 6.6
- RHEL 6.7 inbox
- RHEL 7.0, 7.1 and 7.2

- 
- SLES 11 SP3 and 11 SP4
  - SLES 12 and SLES 12 SP1
  - XenServer® Citrix® 6.5
  - Ubuntu® 14.03 (3.13) and 14.04 (3.19)
  - Debian® 8.0
  - FreeBSD® 10.1
  - OVM 3.3
  - OVM 3.2
  - Oracle® UEK R3 and R4
  - Oracle S10, S11, and S12
  - Windows Server® 2012 and Windows Server 2012 R2

## 1.3 Abbreviations

ACL	Access Control List
AL_PA	Arbitrated Loop Physical Address
API	application programming interface
ARI	alternative routing-ID interpretation
BIOS	basic input/output system
CHAP	Challenge Handshake Authentication Protocol
CLI	command line interface
CPU	central processing unit
CRC	cyclic redundancy check
CT	Common Transport
DCBX	Data Center Bridging Capabilities Exchange
DCC	device control channel
DHCP	Dynamic Host Control Protocol
DID	device ID
DMA	direct memory access
DRSS	default queue receive-side scaling
ELS	extended link service
ETO	extended timeout
FC	Fibre Channel
FCF	Fibre Channel over Ethernet Forwarder
FCoE	Fibre Channel over Ethernet
FCP	Fibre Channel Protocol
FDMI	Fabric-Device Management Interface
FIP	FCoE Initialization Protocol
FLOGI	fabric login
GB	gigabyte
Gb	gigabit
GbE	gigabit Ethernet
Gbps	gigabits per second
GFC	gigabit Fibre Channel
GUI	Graphical User Interface
HBA	host bus adapter
HTTP	Hypertext Transfer Protocol
HW	hardware
IHV	Independent Hardware Vendors
I/O	input/output
IMA	iSCSI Management API
IOCB	input/output control block

---

IOCTL	input/output control
IODM	I/O Device Management
IOPs	I/O operations per second
IP	Internet Protocol
IQN	iSCSI qualified name
iSCSI	Internet Small Computer System Interface
KB	Kilobyte (1024 bytes)
LACP	Link Aggregation Control Protocol
LDTO	link down timeout
LUN	logical unit number
MAC	Media Access Control
MB	megabyte
Mb	megabit
Mbps	megabits per second
MPIO	multipath I/O
MSI	message signaled interrupt
MSI-X	message signaled interrupt-extended
MTU	maximum transmission unit
NIC	network interface card
NFS	Network File System
NLP	Natural Language Processing
NPar	NIC partitioning
NPIV	N_Port ID Virtualization
N_Port	node port
PCI	Peripheral Component Interconnect (interface)
PCIe	Peripheral Component Interconnect Express
PE	Protocol Endpoints
PF	physical function
PLOGI	port login
POST	power-on self-test
R_A_TOV	resource allocation timeout
RFC	Request for Comments
RHEL	Red Hat® Enterprise Linux®
ROM	read-only memory
RPI	remote port indicator
RRQ	Reinstate Recovery Qualifier
RSCN	registered state change notification
RSS	receive-side scaling
SAN	storage area network
SCSI	Small Computer System Interface

---

SFP	small form-factor pluggable
SLES	SUSE Linux Enterprise Server
SLI	Service Level Interface
SR-IOV	single root I/O virtualization
TCP	Transmission Control Protocol
TSO	TCP segmentation offload
UEFI	Unified Extensible Firmware Interface
UMC	Universal Multi-Channel
VASA	vSphere® APIs for Storage Awareness
VF	virtual function
VGT	virtual guest tagging
VLAN	virtual local area network
VLAN ID	VLAN identifier
VM	virtual machine
VPD	vital product data
vPort	virtual port
vSwitch	virtual switch
VST	VLAN Switch Tagging
VVols	Virtual Volumes
VXLAN	Virtual extensible LAN
WLAN	wireless LAN
WWN	World Wide Name
WWNN	World Wide Node Name
WWPN	World Wide Port Name
XRI	extensible resource indicator

## Chapter 2: Installation

This section provides information for installing the ESXi driver for the FC/FCoE, NIC, and iSCSI protocols.

### 2.1 ESXi 5.5 and 6.0 Overview

With the release of ESXi 5.5 (vSphere 2013) and ESXi 6.0 (vSphere 2015), VMware introduced a new driver model called *native mode*. Broadcom has endorsed the native mode driver model for ESXi 5.5 with the following drivers:

- For FC/FCoE functionality, the out-of-box driver is the native mode *lpfc* driver for OneConnect OCe11100-series, OCe14000-series, LPe12000-series, LPe16202/OCe15100, LPe16000-series, LPe31000-series, and LPe32000-series adapters.
- For Ethernet (NIC) functionality, the out-of-box driver for OCe11100-series, OCe14000-series, LPe16202/OCe15100, LPe31000-series, and LPe32000-series adapters is the native mode *elxnet* driver.

**NOTE** For iSCSI functionality on OCe11100 and OCe14000-series adapters, the out-of-box driver continues to be the vmklinux-based *be2iscsi* driver.

ESXi 5.5 and 6.0 also require changes to the installation process and tools:

- **Driver parameters migration** – If you upgrade to ESXi 5.5 or 6.0, you must document the existing driver parameter values for the initial ram disk for the server being upgraded. After upgrading, you must then manually reprogram those initial ram disk values. See [Section 2.1.1, Migrating Driver Parameters When Upgrading to ESXi 5.5 or 6.0](#), for additional information.
- **Command line tool transition** – For the ESXi 5.5 release and all subsequent operating system releases, VMware is transitioning from esxcfg commands to esxcli commands. Both sets of commands are supported in the ESXi 5.5 and 6.0 release. See [Section 3.1, ESXi Command Line Tool Transition](#), for additional information.
- **Native mode driver management support** – With the introduction of the new native mode driver, ethtool is no longer supported. Instead, you can use either the esxcli commands or the VmkMgmtKeyVal interface to provide troubleshooting support. See [Section 4.4, ESXi 5.5 and ESXi 6.0 Native Mode NIC Driver Troubleshooting Support](#), and [Section 4.5, ESXi 5.5 and ESXi 6.0 Native Mode NIC Driver Support for ethtool Commands](#), for additional information.

#### 2.1.1 Migrating Driver Parameters When Upgrading to ESXi 5.5 or 6.0

VMware's ESXi 5.5 and 6.0 releases have two driver models:

- **vmklinux** – the legacy driver model
- **native mode** – the new driver model

vmklinux is a Linux compatibility layer for device drivers in all ESX releases from ESX 2.x to ESXi 5.1. The vmklinux layer allows IHVs to use their Linux drivers in ESXi with only a few modifications. This consistent driver model usage along the ESXi release train ensures that upgrades automatically copy the initial RAM disk driver parameters from the original operating system to the upgraded operating system. However, this is not the case for ESXi 5.5 and 6.0 systems.

VMware recommends loading the drivers by default in native mode because this is the preferred driver model for all subsequent releases. However, because native mode is different than vmklinux and because there is no native mode in prior ESXi releases, there is no driver parameter migration from vmklinux to native mode when migrating to ESXi 5.5 or 6.0.

If you install ESXi 5.5 or 6.0 over your existing ESXi system, you will notice after updating and rebooting the system that the driver parameter settings did not migrate. For example, if you set the *lpfc\_lun\_queue\_depth* driver parameter

---

to 8 on ESXi 5.1 and then update from ESXi 5.1 to ESXi 5.5, you will see the following when the driver parameters are verified after the update:

```
~ # esxcfg-module -g lpfc
Unknown module name lpfc
```

In summary, when updating an existing ESXi system to ESXi 5.5 or 6.0, you must manually set all of the driver parameter values.

## 2.1.2 Preparing to Upgrade to ESXi 5.5 or 6.0

Upgrading your server to ESXi 5.5 or 6.0 requires an additional planning step. You must evaluate the initial ram disk for the server that will be upgraded and document the existing driver parameter values. After the upgrade, you must then manually reprogram those initial ram disk values.

### 2.1.2.1 FC/FCoE Driver Example

The server administrator executes the following commands on the ESXi 5.1 server prior to the ESXi 5.5 or 6.0 upgrade:

```
~ # esxcfg-module -g lpfc820
lpfc820 enabled = 1 options = 'lpfc_lun_queue_depth=8'
```

After upgrading to ESXi 5.5 or 6.0, the server administrator must run the following command.

**NOTE**

The driver binary name is now the native mode Broadcom Emulex FC/FCoE driver.

```
~ # esxcfg-module -s "lpfc_lun_queue_depth=8" lpfc
```

And to verify that the value has been reprogrammed, type:

```
~ # esxcfg-module -g lpfc
lpfc enabled = 1 options = 'lpfc_lun_queue_depth=8'
```

After the verification is complete, reboot the server to activate the command and burn the new driver parameters into the initial ram disk.

### 2.1.2.2 Ethernet Driver Example

This example assumes a single OCe11100-series adapter in the host, since the max\_vfs parameter receives the number of VFs for each PF.

Before upgrading to ESXi 5.5 or 6.0, you would enter the following command:

```
~ # esxcfg-module -g be2net
be2net enabled = 1 options = 'max_vfs=2,2'
```

After upgrading to ESXi 5.5 or 6.0, when the server administrator runs this command for the out-of-box native mode Ethernet driver for the OCe11100-series adapter, there are no module parameters.

**NOTE**

The driver binary name is now elxnet for the native mode Broadcom Emulex Ethernet driver for the OCe11100-series adapter.

```
~# esxcfg-module -g elxnet
elxnet enabled = 1 options = ''
```

The server administrator must then run the following command to set the driver module parameters:

```
~ # esxcfg-module -s max_vfs=2,2 elxnet
```

And to verify that the value has been reprogrammed, type:

```
~ # esxcfg-module -g elxnet
elxnet enabled = 1 options = 'max_vfs=2,2'
```

After verification has completed, reboot the server to activate the command and burn the new driver parameters into the initial ram disk.

## 2.2 Installing the FC/FCoE Driver

This section provides installation information for the driver and the Emulex CIM Provider using the FC/FCoE interface on ESXi systems. Before using this product, you need a working knowledge of FC/FCoE and network storage devices.

### 2.2.1 Installing the FC/FCoE Driver and Management Software

The Broadcom Emulex driver is available through the VMware support site. Refer to the VMware support website for further details.

For VMware ESXi 5.1, 5.5, and 6.0 hosts, you can manage adapters using the Emulex OneCommand® Manager application on Windows® or the OneCommand Manager application for VMware vCenter application, but you must install and use the appropriate Emulex CIM Provider. See the *OneCommand Manager Application User Manual* and *OneCommand Manager Application for VMware vCenter User Manual* for instructions on installing the respective applications. The installation requires that the CIM Provider be installed. For more information on installing the CIM Provider, see the appropriate *CIM Provider Installation Guide*.

<b>NOTE</b>	Before installing the management application, you must install the FC/FCoE driver from the VMware website and then reboot the server.  Before installing the FC/FCoE drivers and CIM Provider, verify that the firmware version is correct. If it is, proceed with the installation. If it is not, update the firmware using the OneCommand Manager application from a Windows server or the OneCommand Manager application for VMware vCenter, and reboot the system before proceeding with the installation.
-------------	--

### 2.2.2 Uninstalling the FC/FCoE Driver

See the VMware Patch Download page for instructions.

### 2.2.3 Installing the Native mode FC/FCoE Driver esxcli plugin for ESXi5.5 and ESXi6.0

This section describes the installation and usage of the esxcli plug-in for the native mode FC/FCoE(lpfc) driver on ESXi5.5 and ESXi6.0 systems. This is a diagnostic tool similar to procfs for the FC/FCoE driver module.

<b>NOTE</b>	You can download the esxcli plug-in from the Broadcom website. Currently the vib requires a no-signature-check to install for ESXi6.0
-------------	--

To install the esxcli elxfc plug-in for ESXi5.5, perform the following steps:

1. Copy the following file to /var/log/vmware/ on the ESXi host:

```
vmware-esx-elxcli-fc-<version>.vib
```

2. On the ESXi host, install the vibs as follows:

```
esxcli software vib install -v <pathname>/ vmware-esx-elxcli-fc-<version>.vib
```

3. Restart the hostd using the following command:

```
/etc/init.d/hostd restart
```

Once the driver has restarted, the elxnet namespace appears under elxfc.

Supported commands:

```
esxcli elxfc
```

Usage:

```
esxcli elxfc <cmd> <cmd options>
```

**Table 1 Available Namespaces**

Namespace	Description
adapter	Broadcom Emulex FC/FCoE HBA instances adapter info
bginjerr	Broadcom Emulex FC/FCoE HBA instances blockguard error injection page info
event	Broadcom Emulex FC/FCoE HBA instances events info
eventsnap	Broadcom Emulex FC/FCoE HBA instances events info with snap after display
fcf	Broadcom Emulex FC/FCoE HBA instances fcf table info
heap	Broadcom Emulex FC/FCoE HBA instances heap usage stats and info
iostat	Broadcom Emulex FC/FCoE HBA instances i/o stats info
kv	Broadcom Emulex FC/FCoE HBA instances kv page info
lun	Broadcom Emulex FC/FCoE HBA instances SAN lun info
memlog	Broadcom Emulex FC/FCoE HBA instances memlog info
memory	Broadcom Emulex FC/FCoE HBA instances memory stats and info
mgmt	Broadcom Emulex FC/FCoE HBA instances mgmt stats and info
node	Broadcom Emulex FC/FCoE HBA instances nodes in SAN info
param	Broadcom Emulex FC/FCoE HBA instances dynamic config param info
port	Broadcom Emulex FC/FCoE HBA instances port info and stats
rsnn	Broadcom Emulex FC/FCoE HBA instances symbolic node name info
sli4q	Broadcom Emulex FC/FCoE HBA instances sli4q info
target	Broadcom Emulex FC/FCoE HBA instances SAN targets info
xlane	Broadcom Emulex FC/FCoE HBA instances xlane info

**Table 2 Available Commands**

Command	Description
list	Lists Broadcom Emulex FC/FCoE adapter instance names.
logmessage	Logs WARNING: lpfc: lpfc_mgmt_log_msg:xxx message to vmkernel.log
modinfo	Returns LPFC module information for list of FC/FCoE boards.

Broadcom recommend that you use the same or higher version of the plugin as the corresponding driver(lpfc) version installed.

For ESXi6.0 follow the same procedure as above, but use the following command in step2.

```
esxcli software vib install -v <pathname>/ EMULEX-ESX-elxlpfccli-<version>.vib  
-no-sig-check
```

## 2.3 Installing the NIC Driver

This section provides installation information for the NIC driver. Before using this product, you need a working knowledge of NIC and network-storage devices.

### 2.3.1 Installing the NIC Driver and Management Software

For ESXi 5.1, 5.5, and 6.0 systems, the latest Broadcom Emulex NIC driver is available out-of-box. Refer to the operating system guide for installation instructions.

<b>NOTE</b>	Before installing the OneCommand Manager application for VMware vCenter software plug-in, you must install the NIC driver from the VMware software website.
	Before installing the driver and CIM Provider, verify that the firmware version is correct. If it is, proceed with the installation. If it is not, you must update the firmware using the OneCommand Manager application for VMware vCenter software plug-in. See the <i>OneCommand Manager Application for VMware vCenter User Manual</i> for additional information. The installation requires that the CIM Provider be installed. For more information on installing the CIM Provider, see the appropriate <i>CIM Provider Installation Guide</i> .

### 2.3.2 Uninstalling the NIC Driver

See the VMware Patch Download page for instructions.

### 2.3.3 Installing the Native Mode Ethernet Driver esxcli Plug-in for ESXi 5.5 and 6.0

This section describes the installation and usage of the esxcli plug-in for the native mode Ethernet (elxnet) driver on ESXi 5.5 and 6.0 systems.

<b>NOTE</b>	You can download the esxcli plug-in from the Avago website.
-------------	---

To install the esxcli elxnet plug-in, perform the following steps:

1. Copy the following file to `/var/log/vmware/` on the ESXi host:  
`vmware-esx-elxcli-<version>.vib`
2. On the ESXi host, install the vibs as follows:  
`esxcli software vib install -v <pathname>/ vmware-esx-elxcli<version>.vib`
3. Restart `hostd`:  
`/etc/init.d/hostd restart`
4. Once the driver has restarted, the elxnet namespace appears under elxnet.

Supported commands:

<b>NOTE</b>	The following command can be used to get the <code>&lt;pci dev name&gt;</code> value: <code>#esxcli network nic list</code>
-------------	--

- **stats** – the stats command is invoked as follows:  
`esxcli elxnet stats get -p <pci dev name>`

For example:

```
elxnet stats get -p 0000:003:00.1
```

- **regdump (fatdump)** – the regdump command is invoked as follows:

```
esxcli elxnet regdump get -f <absolute file pathname> -p <pci device name>
```

For example:

```
esxcli elxnet regdump get -f /fat.txt -p 0000:002:00.0
```

- **debugMask (applicable for both ESXi 5.5 and 6.0)** – the debugMask command is invoked as follows:

- To retrieve the debugMask, type:

```
esxcli elxnet dbgmask get -p <pci dev name>
```

- To set the debugMask, type:

```
esxcli elxnet dbgmask set -p <pci dev name> -m <Debug Mask>
```

## 2.3.4 Virtualization Support

For the best performance, you must install VMware Tools in each guest operating system. For information on installing VMware Tools in a Linux or Windows guest operating system, refer to the appropriate VMware ESXi Server documentation.

## 2.3.5 Obtaining Information on the Installed NIC Driver

To get information on the installed NIC driver on ESXi 5.1 systems, enter:

```
esxcli software vib list | grep be2net
```

To get information on the install NIC driver on ESXi 5.5 and 6.0 systems, enter:

```
esxcli software vib list | grep elxnet
```

## 2.4 Installing the iSCSI Driver

This section provides installation information for the Broadcom Emulex iSCSI driver. Before using this product, you need a working knowledge of iSCSI and network-storage devices.

### 2.4.1 Installing the iSCSI Driver and Management Software

For ESXi 5.1, 5.5, and 6.0 systems, the Broadcom Emulex iSCSI driver is available out-of-box. A direct link to the iSCSI driver available for download from VMware is provided on the Downloads page on the Avago website.

#### NOTE

Before installing the OneCommand Manager for VMware vCenter Server® application, you must install the iSCSI driver from the VMware software website.

Before installing the driver and CIM Provider, verify that the firmware version is correct. If it is, proceed with the installation. If it is not, you must update the firmware using the OneCommand Manager for VMware vCenter Server application or the OneCommand Manager application on a Windows server. See the *OneCommand Manager for VMware vCenter User Manual* and the *OneCommand Manager Application User Manual* for additional information. The installation

---

requires that the CIM Provider be installed. For more information on installing the CIM Provider, see the appropriate *CIM Provider Installation Guide*.

The driver kit includes the IMA plug-in; see the documentation from VMware for information.

## 2.5 Updating Drivers with VMware Offline Bundle Media

VMware recommends using the offline bundle to update software on VMware ESXi 5.1, 5.5, and 6.0 platforms.

**NOTE** For more information about the ESXi Patch Management activities, refer to the VMware website.

To update a driver with the offline bundle media, perform the following steps:

1. Run the following command:

```
esxcli software vib install --maintenance-mode -d  
<pathname>/esx-5.5.0-<build_version>-be2iscsi-<driver-version>.zip
```

where *<build\_version>* represents the FC/FCoE, NIC, or iSCSI driver.

For example, to update the iSCSI driver, type the following command:

```
esxcli software vib install --maintenance-mode -d  
<pathname>/esx-5.5.0-472629-be2iscsi-11.0.190.0.zip
```

2. Reboot the VMware ESXi Server to activate the drivers.

# Chapter 3: Configuration

This section provides configuration information for the ESXi driver for FC/FCoE, NIC, and iSCSI protocols.

## 3.1 ESXi Command Line Tool Transition

For the ESXi 5.5 release and all subsequent operating system releases, VMware is starting the transition from esxcfg commands to esxcli commands. This section describes the transition to the esxcli commands.

**NOTE** Both sets of commands are supported in the ESXi 5.5 and 6.0 releases.

### 3.1.1 ESXi 5.1 Implementation

For ESXi 5.1, Broadcom used the esxcfg series of native command line applications to mine data, but the primary application used was the esxcfg-module to set and get driver parameters.

#### 3.1.1.1 FC/FCoE Driver Examples

For the FC/FCoE driver:

```
esxcfg-module -s "lpfc_log_verbose=0x10c3" lpfc820
```

This command could also be used to retrieve the driver parameters set.

For example:

```
esxcfg-module -g lpfc820
lpfc820 enabled = 1 options = 'lpfc_log_verbose=0x10c3'
```

#### 3.1.1.2 Ethernet Driver Examples

For the Ethernet driver, the following commands were also used to set and retrieve driver parameters:

■ To retrieve the parameters set:

```
~ # esxcfg-module -g be2net
be2net enabled = 1 options = 'max_vfs=2,2'
```

■ To set the module parameter that specifies the number of virtual functions:

```
~ # esxcfg-module -s max_vfs=2,2 be2net
```

■ To set the module parameter that enables or disables MSI-X:

```
~ # esxcfg-module -s msix=0 be2net
```

### 3.1.2 ESXi 5.5 and 6.0 Implementation

For ESXi 5.5 and 6.0, Broadcom uses the esxcli version of the command sequence, which has a different command syntax.

For example:

The command syntax to list the supported parameters by a driver:

```
~ # esxcli system module parameters list -m <driver binary name>
```

---

The command syntax for setting a parameter to a driver module:

```
~ # esxcli system module parameters set -p <parameter_string> -m <driver binary name>
```

### 3.1.2.1 FC/FCoE Driver Examples

To set extended logging for the Broadcom Emulex FC/FCoE driver:

```
~ # esxcli system module parameters set -p lpfc_log_verbose=0x10c3 -m lpfc
```

To list the parameter values assigned to a driver module:

```
~ # esxcli system module parameters list -m lpfc
```

This command is the same as listing parameters. The parameter set has the *value* column completed if applicable. This command lists all instance and global parameters, which makes it necessary to watch for any altered driver parameters.

### 3.1.2.2 Ethernet Driver Examples

To list the module parameters supported by the elxnet driver:

```
~ # esxcli system module parameters list -m elxnet
```

To set the number of virtual functions for the Broadcom Emulex elxnet driver:

```
~# esxcli system module parameters set -p max_vfs=2,2 -m elxnet
```

To verify the altered parameter value in the elxnet driver module:

```
~ # esxcli system module parameters list -m elxnet
```

To disable MSI-X support for the Broadcom Emulex elxnet driver:

```
~ # esxcli system module parameters set -p msix=0 -m elxnet
```

To verify the altered parameter value in the elxnet driver module:

```
~ # esxcli system module parameters list -m elxnet
```

To enable emi\_canceller for the Broadcom Emulex elxnet driver:

```
~ # esxcli system module parameters set -p emi_canceller=1 -m elxnet
```

To verify the altered parameter value in the elxnet driver module:

```
~ # esxcli system module parameters list -m elxnet
```

To disable vxlan\_offload for the Broadcom Emulex elxnet driver:

```
~ # esxcli system module parameters set -p vxlan_offload=0 -m elxnet
```

**NOTE** This parameter applies to OCe14000-series adapters only.

To enable RSS for the Broadcom Emulex elxnet driver:

```
# esxcli system module parameters set -p rss=1 -m elxnet
```

To change the global DebugMask value of the Broadcom Emulex elxnet driver:

```
# esxcli system module parameters set -p debugMask=0x0120 -m elxnet
```

Refer to [Section 4.2.3, NIC Informational Log Groups for ESXi 5.5 and 6.0](#), for more information.

To verify the altered parameter value in the elxnet driver module:

```
~ # esxcli system module parameters list -m elxnet
```

## 3.2 FC/FCoE Driver Configuration

You can configure driver parameters using native ESXi tools, the Emulex OneCommand Manager application (for use in non-lockdown mode only), or the OneCommand Manager for VMware vCenter Server application (for use in both lockdown and non-lockdown modes).

This document describes how to configure parameters using native ESXi tools. For a more comprehensive description of ESXi tools, refer to VMware's public website. If you have further questions, contact a VMware technical support representative.

See the *OneCommand Manager Application User Manual* and the *OneCommand Manager Command Line Interface User Manual* for information about the OneCommand Manager application.

See the *OneCommand Manager for VMware vCenter User Manual* for information about the OneCommand Manager for VMware vCenter Server application.

### 3.2.1 FC/FCoE Driver Parameters Configuration Methods

Configure the driver parameters using any of the following methods:

- Permanent (global)
- Permanent (per adapter)
- Temporary (global)
- Temporary (per adapter)

The OneCommand Manager application supports all four ways to configure driver parameters. This is the preferred method of setting configuration parameters. Refer to the *OneCommand Manager Application User Manual* or the *OneCommand Manager for VMware vCenter User Manual* for more information.

The native ESXi tools only support permanent configuration methods for the driver parameters. The following section provides further information on permanent configuration methods.

#### 3.2.1.1 Permanent FC/FCoE Configuration Methods Using Native ESXi Tools

Permanent configuration requires that the new values be saved in the ESXi environment. These changes are considered permanent because they persist across system reboots.

Refer to [Section 3.2.5, FC/FCoE Driver Configuration Parameters](#), for parameter names and values. Parameter values are in both hexadecimal and decimal.

**NOTE** For ESXi systems, the following steps must be executed from the Troubleshooting Administrative Shell environment. If your configuration does not provide access to this shell, refer to VMware's vSphere or VMware's vCenter server manual for enabling driver logging. Alternatively, refer to Broadcom's Emulex *CIM Provider Installation Guide* for driver logging.

To make changes that impact all adapters in the system (global changes):

1. From the Troubleshooting Administrative Shell environment's terminal window, use one of the following commands:

For ESXi 5.5 and 6.0 systems, type:

```
esxcli system module parameters set -p "param1=value param2=value ..." -m lpfc
```

**NOTE** Use quotation marks around the parameter values only when listing two or more.

For ESXi 5.1 systems, type

```
esxcfg-module -s "param=value param2=value..." <driver_name>
```

The <driver\_name> is obtained from the `vmkload_mod -l` call. Look for the `lpfc` prefix.

2. To reboot the server, type:

```
reboot
```

**NOTE**

VMware does not officially support unloading the driver via `vmkload_mod -u`. If you must unload the driver, contact VMware technical support.

NPIV port creation and deletion are performed by the VMware vSphere client or Virtual Center service. Refer to the VMware documentation for more information.

### 3.2.1.1.1 Example of Permanent Global Configuration

The following example sets `lun_queue_depth` (the maximum number of commands that can be sent to a single LUN) to 20 (the default is 30) for all Broadcom Emulex adapters in your system.

1. Locate the parameter `lpfc_lun_queue_depth` in [Table 1](#).
2. Set the permanent value using one of the following commands:

For ESXi 5.5 and 6.0 systems, type:

```
esxcli system module parameters set -p lpfc_lun_queue_depth=20 -m lpfc
```

For ESXi 5.1 systems, type:

```
esxcfg-module -s lpfc_lun_queue_depth=20 -m lpfc820
```

3. To reboot the server, type:

```
reboot
```

The new setting is used when the driver reloads.

To verify the setting, use one of the following commands:

For ESXi 5.5 and 6.0 systems, type:

```
esxcli system module parameters list -m lpfc
```

For ESXi 5.1 systems, type:

```
esxcfg-module -g lpfc820
```

### 3.2.1.1.2 Examples of Permanent Per-Adapter Configuration

The following example sets `lun_queue_depth` to 20 (the default is 30) for adapter #1.

1. Set the adapter-specific value using one of the following commands:

For ESXi 5.5 and 6.0 systems, type:

```
esxcli system module parameters set -p lpfc1_lun_queue_depth=20 -m lpfc
```

For ESXi 5.1 systems, type:

```
esxcfg-module -s lpfc1_lun_queue_depth=20 lpfc820
```

2. To reboot the server, type:

```
reboot
```

The new setting is used when the driver is reloaded.

---

To verify the setting, use one of the following commands:

For ESXi 5.5 and 6.0 systems, type:

```
esxcli system module parameters list -m lpfc
```

For ESXi 5.1 systems, type:

```
esxcfg-module -g lpfc820
```

The following example sets lun\_queue\_depth to 20 (the default is 30) for adapter #1 and lun\_queue\_depth to 10 (the default is 30) for adapter #2 on an ESXi 5.1 server.

1. Set the adapter-specific value using one of the following commands:

For ESXi 5.5 and 6.0 systems, type:

```
esxcli system module parameters set -p "lpfc1_lun_queue_depth=20, lpfc2_lun_queue_depth=10" -m lpfc
```

For ESXi 5.1 systems, type:

```
esxcfg-module -s "lpfc1_lun_queue_depth=20 lpfc2_lun_queue_depth=10" lpfc820
```

**NOTE** Type the command all on one line without a carriage return.

2. To reboot the server, type:

```
reboot
```

The new settings are used when the driver is reloaded.

To verify the settings, use one of the following commands:

For ESXi 5.5 and 6.0 systems, type:

```
esxcli system module parameters list -m lpfc
```

For ESXi 5.1 systems, type:

```
esxcfg-module -g lpfc820
```

### 3.2.2 Emulex ExpressLane Support

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.

Emulex ExpressLane™ enables special priority queuing for selected LUNs (ExpressLane LUNs). ExpressLane LUN performance is superior to that of regular LUNs. Mixed storage environments can use ExpressLane to alleviate congestion, reduce latency, and improve throughput, ensuring that key applications receive highest priority.

The following requirements must be met to use ExpressLane:

- ExpressLane is supported on LPe16000-series, LPe16202/OCe15100, LPe31000-series, and LPe32000-series adapters only.
- ExpressLane is only supported on ESXi 5.5 or later.
- An ExpressLane LUN must be enabled in the driver before it can be used by OneCommand Manager. Additionally, the priority value to mark each of the ExpressLane frames must be specified to the FC/FCoE driver through the appropriate driver parameters.
- ExpressLane is not supported for LUNs attached to virtual ports.

For additional information, see the *OneCommand Manager for VMware vCenter User Manual*.

### 3.2.3 Dynamically Adding LUNs

For instructions on dynamically adding LUNs, refer to the Using Rescan section of the VMware SAN Configuration documentation.

### 3.2.4 Dynamically Adding Targets

VMware does not provide a native mechanism for dynamically adding targets. After all target/LUN configuration steps have been successfully completed, add the target to the intended fabric zone.

To get the driver to log into the target, either the target or initiator link must be bounced. If the target is configured with security ACLs, the same link bounce requirement applies after the security ACLs are corrected.

To force the ESXi server to rescan all devices:

- Run the following command:

```
esxcfg-rescan vmhbaX
```

-Or-

- From the vSphere Client, click Configuration Tab > Storage Adapters, and then click Rescan All.

### 3.2.5 FC/FCoE Driver Configuration Parameters

[Table 1](#), lists the FC/FCoE driver module parameters, their descriptions, and their corresponding values in previous ESXi environments and in ESXi 5.5 and 6.0 native mode.

NOTE	For ESXi 5.1 systems, all adapter-specific parameters must have an <code>lpfcX_</code> prefix (where X is the driver instance number). For example, setting <code>lpfc0_lun_queue_depth=20</code> makes 20 the default maximum number of commands that can be sent to a single logical unit (disk) for <code>lpfc</code> instance 0.
------	--

Dynamic parameters do not require a system reboot for changes to take effect.

**Table 1 FC/FCoE Driver Parameters**

Module Parameter	Description	ESXi 5.1 Legacy Driver Model Values	ESXi 5.5 and 6.0 OP Native Mode Driver Model Values	Comments
throttle_log_cnt	Do not exceed this number of messages logged within 'throttle_log_time'	—	Default (Def) = 10 Minimum (Min) = 1 Maximum (Max) = 1000	For the native mode driver. Logging mechanism intended to speed up issue diagnosis by reducing the need to enable driver logging.
throttle_log_time	Do not exceed 'throttle_log_cnt' number of logs within this time limit (seconds)	—	Def = 1 Min = 1 Max = 60	For the native mode driver. Works with throttle_log_cnt.
compression_log	Define how often the compression logs are written (in seconds)	—	Def = 300 Min = 5 Max = 86400	For the native mode driver. The driver uses this parameter to periodically write status messages to the vmkernel log. The messages provide state analysis on the paths, targets, and adapter. It differs from throttle in that throttle stops the driver from spamming the logs on a very high frequency failure.
suppress_link_up	Suppress link up at initialization: 0x0 = bring up link 0x1 = do not bring up link 0x2 = never bring up link	—	Def = 0 Min = 0 Max = 2	For the native mode driver. Enable this parameter to assist with SAN issues during ESX boot.
max_targets	The maximum number of discovered targets allowed.	—	Def = 256 Min = 0 Max = 4096	For the native mode driver. Driver parameter to adjust the supported target count.
max_multiq	Sets how many completion queues the driver is requesting from ESXi for each HBA instance.  Each completion queue uses a MSI-X vector.  0 indicates disabled Multi-Queue.	—	Def = 4 Min = 0 Max = 8	Replaces use_mq.
lpfc_delay_discovery	Delay NPort discovery when the Clean Address bit is cleared.	Def = 0 Min = 0 Max = 1	Def = 0 Min = 0 Max = 1	No changes.
use_mq	Use the MultiQueue kernel API for ESXi 5.x operating systems.  0 = disable MultiQueue 1 = enable MultiQueue	Def = 1 Min = 0 Max = 1	—	Deprecated in the native mode driver. Replaced with max_multiq.
fcp_wq_count	Set the number of fast-path FCP work queues, if possible.	Def = 8 Min = 1 Max = 31	—	Deprecated in the native mode driver.
fcp_eq_count	Set the number of fast-path FCP event queues, if possible.	Def = 8 Min = 1 Max = 31	—	Deprecated in the native mode driver.

**Table 1 FC/FCoE Driver Parameters (Continued)**

Module Parameter	Description	ESXi 5.1 Legacy Driver Model Values	ESXi 5.5 and 6.0 OP Native Mode Driver Model Values	Comments
enable_fcp_priority	Enable (1) or disable (0) FCP priority.	Def = 0 Min = 0 Max = 1	Def = 0 Min = 0 Max = 1	No changes.
iocb_cnt	IOCBs allocated for extended link service, common transport, and abort sequence in 1024 increments.	Def = 1 Min = 1 Max = 5	Def = 1 Min = 1 Max = 5	No changes.
sli_mode	SLI mode selector: 0 - auto 2 - SLI-2 3 - SLI-3	Def = 0 Min = 0 Max = 3	—	Deprecated in the native mode driver.
devloss_tmo	The number of seconds the driver holds I/O waiting for a loss device to return.	Def = 10 Min = 1 Max = 255	Def = 10 Min = 1 Max = 255	No changes.
log_verbose	Verbose logging bit-mask.	Def = 0 Min = 0 Max = 0x7fffffff	Def = 0 Min = 0 Max = 0x7fffffff	No changes.
lun_queue_depth	The maximum number of FCP commands that can queue to a specific LUN.  <b>NOTE</b> The driver dynamically limits the runtime lun_queue_depth setting to 1/8th of the hba_queue_depth to prevent I/O starvation. An attempt to set the lun_queue_depth higher than the 1/8th setting results in a failure. The console logs and the adapter KeyVal page reflects the failure.	Def = 30 Min = 1 Max = 128	Def = 30 Min = 1 Max = 254	No changes.
tgt_queue_depth	The maximum number of FCP commands queued to a specific target port.	Def = 8192 Min = 10 Max = 8192	Def = 65535 Min = 10 Max = 65535	Default and maximum values increased.
hba_queue_depth	The maximum number of FCP commands queued to an FC/FCoE adapter.  The driver automatically adjusts the hba_queue_depth to match adapter capabilities. This setting may be overridden.	Def = 8192 Min = 32 Max = 8192	Def = 8192 Min = 32 Max = 8192	No changes.
enable_qfull	Enable the driver's lun_queue_depth ramp down/up functionality when the SCSI device status is Task Set Full (x28).  When enabled (value = 1), the Broadcom Emulex driver gradually reduces the LUN queue depth with each Task Set Full status completion, and then ramps back up as the I/O successfully completes.  When disabled (value = 0), the driver takes no action when an I/O completes with a Task Set Full status.	N/A	Def = 1 Min = 0 Max = 1	Disabling this parameter has a potential impact on the VM(s) overall performance. Consult your target vendor for guidance.

**Table 1 FC/FCoE Driver Parameters (Continued)**

Module Parameter	Description	ESXi 5.1 Legacy Driver Model Values	ESXi 5.5 and 6.0 OP Native Mode Driver Model Values	Comments
scan_down	Start scanning for devices from highest AL_PA to lowest.	Def = 1 Min = 0 Max = 1	Def = 1 Min = 0 Max = 1	No changes.
topology	Select FC topology. Valid values are: 0x0 = attempt loop mode then point-to-point 0x01 = internal loopback mode 0x02 = attempt point-to-point mode only 0x04 = attempt loop mode only 0x06 = attempt point-to-point mode then loop	Def = 0 Min = 0 Max = 6	Def = 0 Min = 0 Max = 6	No changes.
link_speed	Select link speed: 0 = auto select 1 = 1 Gigabaud 2 = 2 Gigabaud 4 = 4 Gigabaud 8 = 8 Gigabaud 10 = 10 Gigabaud 16 = 16 Gigabaud 32 = 32 Gigabaud	Def = 0 Min = 0 Max = 16	Def = 0 Min = 0 Max = 32	Updated for LPe32000-series adapters.  <b>NOTE</b> Setting this option incorrectly can cause the adapter to fail to initialize.  <b>NOTE</b> If you configure the link speed in a BIOS utility, the link speed may be overwritten by the operating system according to its own configuration settings. To avoid this issue, configure the link speed in both the operating system driver and the Boot BIOS or UEFI driver.
fcp_class	Select FC class of service for FCP sequences.	Def = 3 Min = 2 Max = 3	Def = 3 Min = 2 Max = 3	No changes.
use_adisc	Use address discovery on rediscovery, initiated by RSCN, to authenticate FCP devices instead of port login.	Def = 0 Min = 0 Max = 1	Def = 0 Min = 0 Max = 1	No changes.
first_burst_size	First burst size for targets that support first burst.	—	Def = 0 Min = 0 Max = 65536	For the native mode driver.
max_scsicmpl_time	Use SCSI command completion time to control queue depth to the device. 0 – SCSI command completion time is not used for controlling I/O queue depth. N – I/O queue depth is controlled to limit the I/O completion time to N ms.	Def = 0 Min = 0 Max = 60000	Def = 0 Min = 0 Max = 60000	No changes.
ack0	Enable ACK0 support. Use ACK0, instead of ACK1, for class 2 acknowledgement.	Def = 0 Min = 0 Max = 1		Deprecated in the native mode driver.

**Table 1 FC/FCoE Driver Parameters (Continued)**

Module Parameter	Description	ESXi 5.1 Legacy Driver Model Values	ESXi 5.5 and 6.0 OP Native Mode Driver Model Values	Comments
fdmi_on	<p>Enable FDMI support.</p> <p>For ESXi 5.1:</p> <p>0 = disable FDMI support (default setting)</p> <p>1 = enable only FDMI</p> <p>2 = enable FDMI</p> <p>For ESXi 5.5 and ESXi 6.0:</p> <p>0 = disable FDMI support (default setting)</p> <p>1 = enable FDMI support</p> <p>3 = enable FDMI with 60-second delay, use a subset of adapter and port attributes</p> <p>5 = enable FDMI without 60-second delay, use all adapter and port attributes</p> <p>7 = enable FDMI with 60-second delay, use all adapter and port attributes</p>	<p>Def = 0</p> <p>Min = 1</p> <p>Max = 2</p>	<p>Def = 0</p> <p>Min = 0</p> <p>Max = 7</p>	<p>FDMI support will default to FDMI-2. If that fails it will fallback to FDMI-1.</p> <p>Values 2, 4, and 6 are reserved.</p> <p>Values 3 and 7 are for compatibility with older switches. If issues arise using values 0, 1, or 5, contact Broadcom technical support.</p> <p>If FDMI is not operational, use a value of 7 to introduce a delay.</p> <p>A port reset or adapter reboot is required for the new setting to take effect.</p> <p>The 60-second delay starts after link up.</p> <p>Port attributes subset are attributes 1 to 6.</p> <p>Adapter attributes subset are attributes 1 to 0xB.</p> <p>All port attributes include attributes: 1–0xD, 0x101, 0x102, and 0x103.</p> <p>All adapter attributes include attributes: 1 to 0xC.</p> <p>See <a href="#">Table 2</a> for a list of the adapter port attributes.</p> <p>See <a href="#">Table 3</a> for a list of the adapter attributes.</p>
discovery_threads	The maximum number of extended link service commands that can be outstanding during discovery.	<p>Def = 32</p> <p>Min = 1</p> <p>Max = 64</p>	<p>Def = 32</p> <p>Min = 1</p> <p>Max = 64</p>	No changes.
max_luns	The maximum number of LUNs allowed.	<p>Def = 256</p> <p>Min = 1</p> <p>Max = 65535</p>	<p>Def = 256</p> <p>Min = 1</p> <p>Max = 256 (ESXi5.5)</p> <p>Max = 65535 (ESXi 6.0)</p>	No changes.
task_mgmt_tmo	The maximum time to wait for task management commands to complete.	—	<p>Def = 60</p> <p>Min = 5</p> <p>Max = 180</p>	For the native mode driver.
use_msi	<p>Use preferred MSI-X interrupt mode if possible.</p> <p>0 = MSI disabled (INTx mode)</p> <p>1 = MSI enabled</p> <p>2 = MSI-X enabled</p>	<p>Def = 2</p> <p>Min = 0</p> <p>Max = 2</p>	<p>Def = 2</p> <p>Min = 0</p> <p>Max = 2</p>	No changes.
fcf_failover_policy	FCF Fast failover = 1 Priority failover = 2	<p>Def = 1</p> <p>Min = 1</p> <p>Max = 2</p>	<p>Def = 1</p> <p>Min = 1</p> <p>Max = 2</p>	No changes.

**Table 1 FC/FCoE Driver Parameters (Continued)**

Module Parameter	Description	ESXi 5.1 Legacy Driver Model Values	ESXi 5.5 and 6.0 OP Native Mode Driver Model Values	Comments
enable_rrq	Enable Reinstate Recovery Qualifier functionality. 0x0: disabled, XRI/OXID use not tracked. 0x1: XRI/OXID reuse is timed with R_A_TOV, Reinstate Recovery Qualifier sent. 0x2: XRI/OXID reuse is timed with R_A_TOV, No Reinstate Recovery Qualifier sent.	Def = 2 Min = 0 Max = 2	Def = 2 Min = 0 Max = 2	No changes.
EnableXLane	Enable ExpressLane.	—	Def = 0 Min = 0 Max = 1	For the native mode driver.
XLanePriority	ExpressLane CS_CTL Priority. Sets the CS_CTL field in FC Header. See the switch vendor administration guide for additional information.	—	Def = 0x0 Min = 0x0 Max = 0x7F	For the native mode driver.
sg_seg_cnt	The maximum Scatter Gather Segment Count for DMA. The maximum data allowed in one SG element is 0x80000000.	Def = 64 Min = 64 Max = 256	Def = 64 Min = 64 Max = 4096	The maximum was increased to account for larger SCSI I/O sizes.
pci_max_read	The maximum DMA read byte count. Valid values are: 0 512 1024 2048 4096	Def = 0 Min = 0 Max = 4096	—	Deprecated in the native mode driver.
nlp_slab_cnt	NLP slab entries.	—	Def = 128 Min = 32 Max = 256	For the native mode driver. Controls the size of the driver's node table. This table in turn limits the driver's ability to discover remote ports, fabrics, initiators, and targets in a zone.
rb_slb_cnt	Receive Buffer slab entries.	—	Def = 256 Min = 32 Max = 256	For the native mode driver. Controls the maximum number of Receive Buffers that will be posted to the adapter.
lpfc_max_heap_size	Maximum allowable memory consumption per server for the LPFC module.	—	Def = 128 MB Min = 64 MB Max = 512 MB	For the native mode driver.
lpfc_mem_tracker	This turns memory tracking on (1) or off (0).	0 = disabled (default) 1 = enabled	0 = disabled (default) 1 = enabled	Keeps track of driver heap and slab alloc/free's. Only used when debugging driver memory issues.
lpfc_external_dif	Enable External DIF support on select devices	0 = disabled 1 = enabled (default)	0 = disabled 1 = enabled (default)	External DIF is only available on SLI4 / FC HBAs (LPe16000 / LPe32000-series adapters)

**NOTE** The values in [Table 2](#) and [Table 3](#) are taken from the FC-GS Fibre Channel Standard documents and are passed to the switch through FC-CT commands. Adapter port attributes are provided for each adapter port. Adapter attributes are provided once for each adapter, no matter the number of ports.

**Table 2 Adapter Port Attributes for fdmi\_on Parameter**

Hexadecimal Value	Information Type
0x001	Supported FC-4 Types
0x002	Supported Speed
0x003	Current Port Speed
0x004	Maximum Frame Size
0x005	Operating System Device Name
0x006	Host Name
0x007	Node Name
0x008	Port Name
0x009	Port Symbolic Name
0x00A	Port Type
0x00B	Supported Classes of Service
0x00C	Port Fabric Name
0x00D	Port Active FC-4 Types
0x101	Port State
0x102	Number of Discovered Ports
0x103	Port Identifier

**Table 3 Adapter Attributes for fdmi\_on Parameter**

Hexadecimal Value	Information Type
0x001	Node Name
0x002	Manufacturer
0x003	Serial Number
0x004	Model
0x005	Model Description
0x006	Hardware Version
0x007	Driver Version
0x008	Option ROM Version (boot code)
0x009	Firmware Version
0x00A	Operating System Name and Version
0x00B	Maximum CT Payload Length
0x00C	Node Symbolic Name

### 3.2.6 Creating an FC Remote Boot Disk

For instructions on creating an FC remote boot disk, refer to the VMware SAN Configuration documentation.

### 3.2.7 Managing Devices through the CIM Interface

VMware on the Visor-based ESXi platforms uses the CIM interface as the only standard management mechanism for device management.

#### 3.2.7.1 Using the OneCommand Manager GUI

For VMware ESXi 5.1, 5.5, and 6.0 hosts, you can manage adapters using the OneCommand Manager application on Windows, but you must install and use the appropriate Emulex CIM Provider.

**NOTE** For VMware ESXi 5.1, 5.5, and 6.0 hosts, if advanced adapter management capabilities are required (for example, iSCSI Management and port disable), use the OneCommand Manager application for VMware vCenter.

#### 3.2.7.2 Using the OneCommand Manager Application for VMware vCenter

The OneCommand Manager for VMware vCenter Server application uses the standard CIM interface to manage the adapters and supports CIM-based device and adapter management. The OneCommand Manager for VMware vCenter Server application also supports existing adapter management functionality based on its proprietary management stack and the standard HBA API interface. To manage LightPulse and OneConnect adapters (including updating the firmware) on an ESXi 5.1, 5.5, or 6.0 host using the OneCommand Manager for VMware vCenter Server application, you must install the out-of-box Emulex CIM Provider on the host.

For more information on installing the OneCommand Manager for VMware vCenter Server application and enabling the CIM Provider, see the *OneCommand Manager for VMware vCenter User Manual*.

### 3.2.8 Installing the Emulex CIM Provider

Refer to the *CIM Provider Installation Guide* for instructions on installing the Emulex CIM provider.

### 3.2.9 Creating, Deleting, and Displaying VPorts

The Broadcom Emulex driver for VMware supports NPIV by default. The only management API for creating and deleting a vPort and creating an NPIV-enabled virtual machine comes from ESXi. VPorts in the driver discover the fabric just like physical ports do, and are subject to the same SAN delays. As the number of VPorts increases, the amount of time it takes to complete remote port discovery increases. This is because the VPorts are created sequentially and each vPort executes discovery synchronously. If your NPIV-enabled virtual machines power-on automatically, powering on could take longer than usual. This is normal for NPIV virtual machines.

The following note applies to VPorts:

**NOTE** Ensure you are using the latest recommended firmware for vPort functionality. Check the Avago website for the latest firmware. Loop devices and NPIV are not supported on the same port at the same time. If you are running a loop topology and you create a vPort, the vPort's link state is offline. VMware ESX supports fabric mode only. You can create VPorts only on 4, 8, 16, and 32 GFC LightPulse adapters, and on OneConnect adapters. You cannot create VPorts on 1 GFC or 2 GFC adapters. The OneCommand Manager application sees all VPorts created by the driver, but the application has read-only access to them.

### 3.2.10 Configuring VVols in ESXi 6.0

The Broadcom Emulex native mode FC/FCoE driver supports the VVol feature released with ESXi 6.0 on OneConnect OCe11100-series, OCe14000-series, LPe12000-series, LPe16202/OCe15100, LPe16000-series, LPe31000-series, and LPe32000-series adapters. VMware's VVol feature allows for dynamic provisioning of storage, based upon the needs of a VM. VM disks, also called VVols, allow VMware administrators to manage storage arrays through the API. Arrays are logically partitioned into storage containers. VVols, are stored natively in the storage containers. I/O from ESX to the array is managed through an access point or Protocol Endpoint (PE) and the storage provider.

#### 3.2.10.1 Storage Containers

Storage containers are a logical abstraction and hold groups of VVols that are physically in the storage array. Storage containers are an alternative to traditional storage based upon LUNs or NFA shares. Storage containers are set up by a storage administrator. Storage container capacity is based on physical storage capacity. The minimum is one storage container per array and the maximum depends upon the array. One storage container can be simultaneously accessed via multiple PEs. Once the storage provider and PEs are in place, the storage container is visible to ESXi hosts.

#### 3.2.10.2 Protocol Endpoints

A PE is an access point that enables communication between an ESXi host and a storage array system. A PE is not a datastore; it is the I/O transport mechanism to access the storage container. A PE is part of the physical storage fabric. A PE is created by a storage administrator.

#### 3.2.10.3 Storage Providers

Storage providers are also referred to as VASA providers. Out-of-band communication between vCenter and the storage array is achieved through the storage provider. The storage provider creates the VVols.

For more information about VVols and instructions on configuring VVols, refer to the VMware and target vendor-supplied documentation.

## 3.3 NIC Driver Configuration

This section describes how to configure parameters for the ESXi NIC driver.

### 3.3.1 Configuring ESXi 5.5 and 6.0 NIC Driver Parameters

Table 4, NIC Ethernet Driver Parameters, lists the Ethernet driver module parameters, their descriptions, and their corresponding values in previous ESXi environments and in ESXi 5.5 and 6.0 native mode.

Table 4 NIC Ethernet Driver Parameters

Module Parameter	Description	ESXi 5.1 Legacy Driver Model Values	ESXi 5.5 and 6.0 OP Native Mode Driver Model Values	Comments
debugMask	The DebugMask is a bit-vector (uint32) and each bit represents a group. The DebugMask value is in eight-digit hexadecimal format (for example, 0x00000101).  DRIVER: 0x1 UPLINK: 0x2 QUEUE: 0x4 INTR: 0x8 MCC: 0x10 TX: 0x20 RX: 0x40 MGMT: 0x80 WORKER: 0x100 SRIOV: 0x200 EVENT: 0x400 VLAN: 0x800 VXLAN: 0x1000	—	Def = 0x1217 (includes DRIVER, UPLINK, MCC, QUEUE, SRIOV, and VXLAN groups)	For the native mode driver.
emi_canceller	Enable or disable the EMI Canceller.	Def = 0	Def = 0	No changes.
max_vfs	The number of PCI VFs to initialize. 0 = disabled 1–63 = enable this many VFs (depends on the ESXi version and adapter)	Def = 0 Min = 0 Max = 16	Def = 0 Min = 0 Max = 63 (depends on adapter)	The driver supports a maximum of 30 VFs per PF for the OCe11100-series adapters, 32 VFs per PF for the LPe16202 and OCe15100 adapters, 32 VFs per PF for the OCe14100-series adapters, and 63 VFs per PF for the OCe14400-series adapters.  <b>NOTE</b> HP Virtual Connect adapters support a maximum of 24 VFs per PF.
msix	Enable or disable MSI-X.	—	Def = 1	For the native mode driver.
rss	Enable or disable RSS.	—	Def = 0	For the native mode driver.
vxlan_offload	Enable or disable vxlan_offload.	—	Def = 1	For the native mode driver.
vlan_offload	Enable or disable VLAN filtering by the hardware.	—	—	Deprecated and not supported by the native mode driver.

#### 3.3.1.1 Deprecated Module Option

The `vlan_offload` module parameter of the previous Ethernet driver (be2net), which provided control for hardware VLAN filtering, is deprecated and is no longer supported in the elxnet driver.

### 3.3.1.2 New Module Options

In the elxnet driver, there are two new module parameters:

- **Enable or disable MSI-X support** – The driver default is to have MSI-X enabled.
- **Enable or disable vxlan\_offload support** – The driver default is to have vxlan\_offload enabled.

**NOTE** The following output is for illustrative purposes only. The actual output may vary depending on the adapter installed in the system.

The server administrator must run the following command to disable MSI-X:

```
~ # esxcli system module parameters set -p msix=0 -m elxnet
```

And to verify that the value has been reprogrammed:

```
~ # esxcli system module parameters list -m elxnet
```

The server administrator must run the following command to disable vxlan\_offload:

```
~ # esxcli system module parameters set -p vxlan_offload=0 -m elxnet
```

**NOTE** This parameter applies to OCe14000-series adapters only.

And to verify that the value has been reprogrammed:

```
~ # esxcli system module parameters list -m elxnet
```

### 3.3.2 Configuring ESXi 5.1 NIC Driver Parameters

The following table lists the user-configurable NIC driver parameters for ESXi 5.1. It includes a description of the parameter and its default value.

**Table 5 User-configurable NIC Driver Parameters**

Parameter	Default Value	Description
heap_initial	32 MB	The size of the memory heap, in bytes, that should be initially allocated for the driver.
heap_max	38 MB	The maximum possible size to which the driver heap is allowed to grow.

The following command line shows how to load the driver with the initial heap size set to 40 MB:

```
# vmkload_mod be2net heap_initial=41943040
```

To configure the NIC driver to load with this value as the initial heap size after each reboot, run the following command and reboot the system:

```
# esxcfg-module -s "heap_initial=41943040" be2net
```

### 3.3.3 NetQueue Support

- NetQueue enables support for multiple transmit and receive rings for improved performance. NetQueue is enabled by default.
- To determine whether NetQueue is enabled, run the following command from the ESXi Server console:

```
# esxcfg-advcfg -j netNetqueueEnabled
```

If it returns 'netNetqueueEnabled = FALSE', NetQueue is disabled.

- To enable NetQueue, run the following command:  
`# esxcfg-advcfg -k TRUE netNetqueueEnabled`  
Reboot the ESXi server for the change to take effect.

### 3.3.4 Receive-Side Scaling

RSS enables the ESXi 5.5 or 6.0 NIC driver to distribute incoming TCP traffic across multiple CPU cores for improved performance.

RSS is supported on OCe11100-series, OCe14000-series, LPe16202/OCe15100, LPe16000-series, LPe31000-series, and LPe32000-series adapters.

In ESXi 5.5 and 6.0, RSS can be selectively enabled for a particular vNIC in a VM. There can be VMs that require RSS and VMs that do not require RSS. This is configurable in the .vmx file. RSS is disabled by default.

To use RSS, the VM must be configured properly:

- The VM must have at least four cores and 4 to 8GB of RAM.
- The vNIC on which RSS is enabled must use VMXNET3 drivers.
- The VMXNET3 drivers must be updated to the latest version.
- The following entry must be added to the .vmx file in the datastore for each VM that requires RSS support for its traffic.

```
ethernet<x>.pnicFeatures="4"
```

Where <x> is the vNIC interface number on which the RSS feature is required.

The absence of the above line for the vNIC in the .vmx file will not force the VM or ESX to use the RSS feature, even though RSS is enabled and queues are created.

RSS limitations include:

- MSI-X must be enabled.
- When Multi-Channel is enabled on OCe11100-series adapters, RSS is supported on the first functions of port 0 and port 1 only.
- The number of regular NetQueues is reduced by 1 when RSS is enabled.
- RSS is not supported on one-port OCe11100-series adapters.
- RSS is supported when the ESXi host has at least 8 CPUs for UMC or 12 CPUs for a physical NIC (the number of CPUs must be greater than or equal to the number of NetQueues plus the number of RSS queues).

**Table 6 Available Queues for Each Adapter Model**

Adapter/Mode	Number of NetQueues (including the default queue)	Number of RSS Queues
OCe11100-series physical NIC	8	4
OCe11100-series UMC (function 0 or 1)	3	4
OCe11100-series UMC (functions 2 to 7)	3	0 (RSS is not supported)
OCe14000-series, LPe16202/OCe15100, and LPe16000-series adapters physical NIC	8	4
OCe14000-series, LPe16202/OCe15100, and LPe16000-series adapters UMC (functions 0 to 7)	4	4

**NOTE**

In cases where the ESXi host does not have the required number of CPUs, the number of NetQueues is reduced by the required number to support four RSS queues.

- RSS support in ESX 5.5 is limited and leaves the configuration of RSS parameters in the hardware completely to the driver. The driver configures the RSS policy in the hardware during driver initialization. The kernel will only request the driver to allocate an RSS queue when an RSS-enabled vNIC (VM) starts receiving packets.

To enable RSS:

1. Run the following command:

```
# esxcli system module parameters set -p "rss=1" -m elxnet
```

2. Reboot the ESXi server for the change to take effect.

To disable RSS:

1. Run the following command:

```
# esxcli system module parameters set -p "rss=0" -m elxnet
```

2. Reboot the ESXi server for the change to take effect.

For ESXi 6.0 systems only, `vsish` can be used to view the contents of the new node created for RSS in each of the uplinks (when RSS is enabled) once the MAC filter is applied. The following information is displayed:

- The number of RSS queues in the hardware
- The RSS hash key
- The RSS indirection table with load factor
- The RSS indirection table size

To view this information, run the following command:

```
vsish -e get /net/pNics/vmnicX/rxqueues/queues/<qid#>/rss
```

To view the CPU indirection table, run the following command:

```
vsish -e get /net/pNics/vmnicX/rxqueues/queues/<rss_qid>/rss/indTable
```

**NOTE**

This node is available only when there is at least one filter placed on the RSS queue (<qid#>). If there is no filter configured on <qid#>, the node <qid#> is not listed under `/net/pNics/<vmnic#>/rxqueues/queues/`.

The receive traffic on multiple RSS queues can be verified by viewing the driver statistics for receive traffic.

To view the driver statistics for receive traffic, run the following command:

```
esxcli elxnet stats get -p <><device_id>> | grep rx_compl >
```

where <device\_id> can be obtained from the `esxcli network nic list` command.

**NOTE**

In normal configurations, the RSS queues will be queues with ID - 8, 9, 10, and 11. When Multi-Channel is enabled, the RSS queues will be queues with ID - 4, 5, 6, and 7.

Multiple core usage can be seen by running the `esxtop` command in the ESX shell. When there is traffic, you can see that more than one CPU is being used even though there is only one active VM.

### 3.3.5

### How an ESXi Server Creates and Names Interfaces

The NIC driver supports a maximum of four adapters per system. For dual-channel adapters running in standard operating mode, the driver creates two interfaces (one for each physical port). The first and second interfaces are respectively named `vmnic0` and `vmnic1` (assuming there are no other network interfaces in your configuration). The same applies to vNIC-capable adapters if vNIC mode is disabled in the adapter BIOS (for those boards that support vNIC).

If vNIC mode is enabled in the adapter BIOS and the adapter is in NIC-only mode, the driver creates eight interfaces (four for each physical port). The interfaces are labeled vmnic0 through vmnic7 (assuming there are no other network interfaces in your configuration).

**NOTE** If vNIC mode is enabled in the adapter BIOS and the adapter is in FCoE or iSCSI mode, the driver creates only six NIC interfaces. The other two interfaces are reserved for FCoE and iSCSI.

All vNICs are fully functional and support the same capabilities as a standard NIC. The vNICs can also be linked to a virtual switch in the same way:

```
# esxcfg-nics -l //list recognized nics
# esxcfg-vswitch -l //list available vswitches
# esxcfg-vswitch -a vSwitch0 //create vSwitch0
# esxcfg-vswitch -A VMNet0 vSwitch0 //create virtual machine network, VMNet0 and
add it to vSwitch0
# esxcfg-vswitch -L vmnic0 vSwitch0 //link vmnic0 to vSwitch0
```

The only difference being that in vNIC mode, each of the four vNICs tied to a physical port shares the port's 10 GbE bandwidth.

### 3.3.6 Enabling SR-IOV

SR-IOV capability can be enabled for OneConnect OCe11100-series, OCe14000-series, and LPe16202/OCe15100 adapters if your system BIOS supports SR-IOV.

ESXi 5.5 and 6.0 fully support configuration networking options of a virtual function that is assigned to a virtual machine with compatibility for ESXi 5.1 or later, including default VLAN Tagging, VGT, and configuring a static MAC for a virtual function using the ESXi driver.

You can also perform these configuration tasks using the vSphere Client. For more information regarding the vSphere Client, see the support section of the VMware website.

ESXi 5.5 and 6.0 have limited support for extended configuration networking options of a virtual function that is assigned to a virtual machine with compatibility for ESXi 5.5 or later.

Supported extended SR-IOV options:

- Enabling VGT
- Enabling VST mode

Extended SR-IOV option limitations:

- Cannot change the size of the MTU (cannot enable jumbo frames)
- Cannot accept or drop incoming frames for a new address with the MAC address change option
- Cannot enable global promiscuous mode for virtual machine network adapters

For more information regarding supported networking configurations, see the support section of the VMware website.

**NOTE** SR-IOV cannot be enabled with other Multi-Channel technologies like UMC or vNIC.

To enable SR-IOV, perform the following steps:

1. Enable IOV capability in the system BIOS.
2. Press **<Ctrl>+p** to enter the BIOS.
3. Enable SR-IOV for each port of the adapter from the BIOS.

- 
4. If you want to use a newer driver, install the driver .vib file after booting the ESXi host.
  5. To load the Ethernet driver with options to enable SR-IOV, run one of the following commands:

ESXi 5.1

```
# esxcfg-module -s max_vfs=x,y be2net
```

ESXi 5.5 and 6.0

```
# esxcfg-module -s max_vfs=x,y elxnet
```

- Or -

```
# esxcli system module parameters set -p "max_vfs=x,y" -m elxnet
```

where *x* and *y* are the number of VFs to be enabled on each of the two NIC PFs. A value of 0 indicates that there are no VFs for the corresponding PF.

**NOTE**

The driver currently supports a maximum of 30 VFs per PF for the OCe11100-series adapters, 32 VFs per PF for the LPe16202/OCe15100 adapters, 32 VFs per PF for the OCe14100-series adapters, and 63 VFs per PF for the OCe14400-series adapters.

HP Virtual Connect adapters support a maximum of 24 VFs per PF.

For multiple adapters, use one of the following commands:

ESXi 5.1

```
# esxcfg-module -s max_vfs=x1,y1,x2,y2,x3,y3 be2net
```

ESXi 5.5 and 6.0

```
# esxcfg-module -s max_vfs=x1,y1,x2,y2,x3,y3 elxnet
```

- Or -

```
# esxcli system module parameters set -p "max_vfs=x1,y1,x2,y2,x3,y3" -m elxnet
```

where:

*x*1 is the number of VFs to be enabled on the first port of the first adapter

*y*1 is the number of VFs to be enabled on the second port of the first adapter

*x*2 is the number of VFs to be enabled on the first port of the second adapter

*y*2 is the number of VFs to be enabled on the second port of the second adapter

*x*3 is the number of VFs to be enabled on the first port of the third adapter

*y*3 is the number of VFs to be enabled on the second port of the third adapter

**NOTE**

The command can be extended to enable more adapters.

6. Reboot the host.
7. To confirm the number of configured VFs, run the following command:

```
# esxcfg-module -g be2net
```

8. To list the SR-IOV enabled pNICs, run the following command:

```
# esxcli network sriovnic list
```

9. To list the status of the VFs enabled on a PF, run the following command:

```
# esxcli network sriovnic vf list -n vmnic<X>
```

where *vmnic<X>* is the interface corresponding to the PF.

**NOTE**

The vSphere Client option Configuration > Advanced Settings also lists the VFs configured with the above command.

- 
10. To assign a VF to a VM, right-click the VM in the vSphere Client and select Edit Settings. Under the Hardware tab, click Add and select PCI Device to attach a VF to the VM.

**NOTE** A maximum of six VFs can be assigned to a VM.

### 3.3.7 Configuring Emulex Universal Multi-Channel

Emulex Universal Multi-Channel™, or UMC, enables the capability to divide a 10GbE port into multiple physical functions, with flexible bandwidth capacity allocation, that appear to the operating system and network as separate physical devices.

Multi-Channel can be configured on OCe14000-series adapters through the adapter BIOS or the OneCommand Manager for VMware vCenter application.

- To configure Multi-Channel using the adapter BIOS, see the *Boot for NIC, iSCSI, FCoE, and RoCE Protocols User Manual*.
- To configure Multi-Channel using the OneCommand Manager for VMware vCenter application, see the latest *OneCommand Manager for VMware vCenter User Manual*.

Refer to the *Emulex Universal Multi-Channel Reference Guide* for additional information on Multi-Channel.

#### 3.3.7.1 Using ARI

If Multi-Channel is enabled on an OCe14000-series network adapter, each port can be partitioned into isolated PFs (channels). You can configure up to 16 functions on a one-port OCe14400-series adapter, up to 8 functions per port on a one or two-port OCe14100-series adapter, and up to 4 functions per port on a four-port OCe14100-series adapter.

The maximum number of functions allowed on an adapter is controlled by the adapter's IPL file and the system's support for ARI.

ARI must be available to support more than eight functions on an adapter. OCe14000-series adapters automatically support ARI. However, the following requirements must be met to support more than eight functions on an adapter.

- The system hardware (the motherboard and BIOS) must support ARI.
- ARI must be enabled in the system BIOS.
- The host or guest operating system must support ARI:
  - Windows Server 2012 and newer versions
  - RHEL 5.9 and newer versions
  - RHEL 6.4 and newer versions
  - SLES 11 SP2 and newer versions
  - ESXi 5.1 and newer versions
- The application management tools, including the OneCommand Manager for VMware vCenter application, must support ARI.
- ARI must be enabled in the firmware using the OneCommand Manager for VMware vCenter application.

If these conditions are not met, you may be able to configure more than eight functions, but only up to eight functions will be running and discovered after a reboot.

### 3.3.8 Configuring VLANs

VLAN filtering is **not** supported in the current hardware. To configure VLANs [in legacy hardware \(OCe11000-series adapters\)](#), create the vSwitch with the required VLAN ID and use this interface as an adapter to this vSwitch. A native

---

VLAN can also be configured in the guest operating system in VGT mode. (For example, using vconfig in the Linux guest operating system.)

**NOTE** You cannot run LACP if UMC is enabled.

Refer to the *Emulex Universal Multi-Channel Reference Guide* for additional information on UMC.

### 3.3.8.1 Default VLAN Tagging

If default VLAN tagging is configured on a VF, a VLAN ID is automatically added to and removed from every packet sent and received over the VF.

To configure a default VLAN tag:

1. Power off the VM before proceeding to the next step.
2. Edit the `<vm-name>.vmx` file corresponding to the VM from the vmfs Datastore under `/vmfs/volumes/Datastore/<VM_DIR>/`.
3. Append the following line:  
`pciPassthru<X>.defaultVlan = "<vlan_id>"`  
where `<vlan_id>` is the required VLAN ID value for the pass-through VF `<X>` assigned to the VM.
4. Power on the VM.

### 3.3.8.2 Virtual Guest Tagging

If VGT is configured, a guest operating system can configure any VLAN ID using the native VLAN configuration method in the guest operating system.

To enable VGT:

1. Configure the special VLAN ID 4095 using the steps listed in [Section 3.3.8.1, Default VLAN Tagging](#).
2. Use the native method in the guest operating system to configure the required VLAN ID (for example, vconfig in the Linux guest operating system).

### 3.3.8.3 Configuring a Static MAC for a VF

A MAC address is automatically configured by ESXi for each VF. This default MAC address can be modified using the following steps:

1. Power off the VM before proceeding to the next step.
2. Edit the `<vm-name>.vmx` file corresponding to a VM from the vmfs Datastore under `/vmfs/volumes/Datastore/<VM_DIR>/`.
3. To configure the MAC address `00:50:56:00:00:02` on the first VF assigned to this VM, append the following lines:  
`pciPassthru0.MACAddressType = "static"`  
`pciPassthru0.MACAddress = "00:50:56:00:00:02"`
4. Power on the VM.

## 3.3.9 Configuring NPar (Dell Only)

**NOTE** This section is specific to Dell®.

NPar allows you to divide a 10Gb port into multiple physical functions, with flexible bandwidth capacity allocation, that appear to the operating system and network as separate physical devices.

### 3.3.9.1 Adapter Configuration

NPar can be configured on OCe14000-series adapters through the adapter BIOS or the OneCommand Manager for VMware vCenter Server application.

- To configure NPar using the adapter BIOS, see the latest *Boot for NIC, iSCSI, FCoE, and RoCE Protocols User Manual*.
- To configure NPar using the OneCommand Manager for VMware vCenter Server application, see the latest *OneCommand Manager for VMware vCenter User Manual*.

On the host operating system side, NPar provides up to 16 physical functions per device, if NParEP is enabled, using the standard PCI configuration space. The number of physical functions that can be mapped to a physical port depends on the adapter:

- Four-port adapter = four physical functions
- Two-port adapter = eight physical functions

Each physical function or partition is assigned a unique MAC address.

Partitions are available for virtual function assignment and for application segmentation via VLAN or IP subnets.

### 3.3.9.2 Adapter Requirements

- The partitions can be on separate subnets or VLANs
- Bandwidth allocation is flexible
- No operating system or BIOS changes are required
- No external switch changes are required
- Each partition should have standard NIC properties for stateless offload

The following items are supported on a per-partition basis:

- Per-partition statistics
- TSO required per partition
- MTU per partition
- Support for NetQueues

### 3.3.9.3 Using NParEP

- NParEP is available on OCe14000-series adapters only.
- NParEP support is available on Dell 13Gb or newer systems only.
- NParEP can be configured on OCe14000-series adapters using the adapter BIOS utility or the OneCommand Manager for VMware vCenter Server application.
- To configure NParEP using the adapter BIOS utility, see the *Boot for NIC, iSCSI, FCoE, and RoCE Protocols User Manual*.
- To configure NParEP using the OneCommand Manager for VMware vCenter application, see the *OneCommand Manager for VMware vCenter User Manual*.

## 3.3.10 Configuring Network Heap Size in ESXi Server 5.1

The amount of memory allocated by default for a network heap depends on the amount of memory configured in the system. The ESXi Server network stack allocates a minimum of 64 MB to the network heap to handle network data. Additional memory is allocated to the network heap if the system is configured with more memory. If the network load requires more than 64 MB of memory, the NIC driver cannot allocate it. When this happens, the driver logs messages in the file `/proc/vmware/log` indicating that the `alloc_skb()` call failed. This impacts network performance considerably.

---

To read the current size of the network heap, run the following command:

```
# esxcfg-advcfg -j netPktHeapMaxSize
netPktHeapMaxSize = 0
#
```

If the default value of 64 MB is in effect, this command shows the PktHeapMaxSize as 0. If any other value is in effect, the command returns that value.

You can allocate more memory for the network heap using the esxcfg-advcfg command. For example, to set the heap size to 128 MB, run the following commands:

```
# esxcfg-advcfg -k 128 netPktHeapMaxSize
# esxcfg-advcfg -j netPktHeapMaxSize
netPktHeapMaxSize = 128
#
```

The new value takes effect after a reboot.

### 3.3.11 Performance Tuning

#### 3.3.11.1 Using vmxnet Emulation

Using vmxnet3 or vmxnet2 as the emulation driver in guest operating systems is crucial for optimal network performance. To configure vmxnet3 as the emulation driver in guest operating systems, you must install VMware Tools in the guest operating systems. For information on installing VMware Tools in a guest operating system, refer to the appropriate VMware ESXi Server documentation.

After VMware Tools are installed, if you add a network adapter for a guest operating system, select vmxnet3 or vmxnet2 as the adapter type.

#### 3.3.11.2 Enabling TSO

Some Broadcom Emulex adapters support TSO, which is necessary to achieve optimal transmit throughput performance with low CPU utilization with the adapter. TSO is enabled by default in ESXi 5.1, 5.5, and 6.0 servers.

To view the current TSO configuration in the vSphere client, perform the following:

1. Select the Configuration tab.
2. Under Software, click Advanced Settings.
3. Under Net Features, view the current value of Net.UseHwTSO.
  - If the value is 1, TSO is enabled.
  - If the value is 0, TSO is disabled.

#### 3.3.11.3 Enabling Jumbo Frames

Enabling jumbo frames reduces CPU utilization and is a recommended practice. To use jumbo frames, you must increase the MTU size in the vSwitch and also in the guest operating system. Broadcom Emulex adapters support MTU sizes between 64 bytes and 9000 bytes. For optimal performance, set the MTU size to the maximum value supported by your network environment. The desired MTU size must be configured in the vSwitch as well as the guest operating systems. Jumbo frames are not enabled by default in ESXi servers.

To configure the MTU size of a vSwitch using the vSphere client, perform the following:

1. Select the Configuration tab.
2. Under Hardware, click Networking.
3. Under Networking, select Properties.

4. Select the vSwitch you want to edit and click Edit.
5. Under vSwitch Properties on the General tab, set the MTU value to 9000.

To configure the MTU size of a vSwitch using the CLI, enter the following command:

```
esxcfg-vswitch -m 9000 vSwitch1
```

### **3.3.11.3.1 Setting the MTU Size for a Linux Guest Operating System**

To set the MTU size in each Linux guest operating system to 9000, run the following command:

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

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

### **3.3.11.3.2 Setting the MTU Size for a Windows Guest Operating System**

To set the MTU size in each Windows guest operating system, perform the following:

1. Go to the Start menu and select Control Panel > System.
2. Select the Hardware tab and open Device Manager.
3. Expand the Network Adapters heading.
4. Right-click the appropriate NIC, and select Properties.
5. Select the Advanced tab and set the MTU value.

## **3.3.11.4 Using the Port Statistics Counters**

You can view all of the port statistics counters maintained by the adapter's Ethernet driver for potential performance issues. Excessive drop or error counters are an indication of a bad link or defective hardware.

To view the statistics of the vmnics on the ESXi host, run the following command:

```
# vsish -e get /net/pNics/<vmnicx>/stats
```

where <vmnicx> is the corresponding interface.

To view the port statistics counters on Linux guest VMs on ESXi 5.1, run the following command:

```
# ethtool -S eth<N>
```

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

[Table 7](#) contains a list of ethtool -S option port statistics counters and their descriptions.

To view the port statistics counters on ESXi 5.5 and 6.0, run the following command:

```
# esxcli network nic stats get -n vmnic0
```

Refer to [Table 24](#) for the equivalent command to get driver priv stats.

**Table 7 Port Statistics Counters**

Ethtool Statistics Counter Name	Description
rx_packets	The number of packets received.
rx_unicast	The number of unicast packets received.
rx广播	The number of broadcast packets received.
rx_multicast	The number of multicast packets received.
rx_bytes	The total number of bytes received.
rx_errors	The total number of error packets received.

**Table 7 Port Statistics Counters (Continued)**

Ethtool Statistics Counter Name	Description
rx_length_errors	The number of packets received with length errors. A length error occurs if an incoming packet is undersized (less than 65 bytes) or oversized (greater than 1522/9216 bytes).
rx_over_errors	Receiver ring buffer overflow errors.
rx_crc_errors	The number of packets received with CRC errors.
rx_frame_errors	The number of packets received with a frame alignment error.
rx_fifo_errors	The number of received packets dropped when FIFO entering the packet demux block overflows, or an internal FIFO going into the main packet buffer tank (PMEM) overflows.
rx_buf_post_fail	The number of times the driver could not allocate a buffer to replenish the receive ring. If this error is frequently observed, you can increase the network heap size using the following command:  esxcfg-advcfg -j netPktHeapMaxSize Reboot the host.
tx_packets	The number of packets transmitted.
tx_bytes	The total number of bytes transmitted.
tx_errors	The total number of error packets transmitted.
tx_timeout_counter	The number of times that a transmit timeout occurred.
tx_rate	The total throughput transmitted from a vmnic interface in Mb/s.
rx_rate	The total throughput received by a vmnic interface in Mb/s.
on_die_temperature	The current temperature of the chip die in degrees Celsius.
link_down_reason	A value indicating the reason for a link being down: 0 – The link is down for an unknown reason. 1 – The link is down due to DCC. 2 – The link is down due to FIP. 3 – The link is down due to vNIC configuration. 4 – The link is down due to a receive PAUSE flood.
ints-netq-<N>	The number of interrupts on Netqueue <N>.
rx_events-netq-<N>	The number of events received on Netqueue <N>.
rx_packets-netq-<N>	The number of packets received on Netqueue <N>.
rx_bytes-netq-<N>	The number of bytes received on Netqueue <N>.
tx_events-netq-<N>	The number of events transmitted on Netqueue <N>.
tx_packets-netq-<N>	The number of packets transmitted on Netqueue <N>.
tx_bytes-netq-<N>	The number of bytes transmitted on Netqueue <N>.

### 3.3.11.5 VXLAN Configuration

The elxnet driver supports hardware offloads for VXLAN on the OCe14000-series adapters. These offloads are enabled by default and can be disabled using the module parameter `vxlan_offload`.

To view the VXLAN statistics, run the following command:

```
# esxcli elxnet stats get -p <pcidevname> | grep vxlan
```

For more information on configuring VXLAN through VMware vCloud Networking and Security Manager and VMware vSphere Distributed Switch (vDS), refer to the *VMware VXLAN Deployment Guide* on the VMware website.

## 3.4 Broadcom Emulex Drivers for VMware NSX 6.2

This section describes the new functionality associated with the Broadcom Emulex drivers for VMware NSX 6.2.

### 3.4.1 Module Parameters

rss	Creates $x$ RSS (receive-side scaling) queues on one of the Net queues ( $x \leq 8$ )
drss	Creates $y$ RSS queues on the default queue ( $y \leq 8$ , excluding the default queue)

### 3.4.2      **Functionality**

- RSS and DRSS (default queue receive-side scaling) are enabled by default, unless explicitly disabled. This is an NSX requirement.
  - rss=1 – creates four Net queue RSS rings.
  - drss=1 – creates four default queue RSS rings.
  - If system or adapter resources are:
    - If DRSS is not specified
      - and-
      - RSS is not specified or is  $\leq 4$
    - The Net queues are given priority and the RSS rings are reduced as required. The RSS rings are used only if at least two RSS rings can be created with the available resources.
    - If RSS is specified and the value is  $> 4$ 
      - or-
      - DRSS is specified
    - The RSS queues are given priority and the maximum possible RSS queues are created. The Net queue count is reduced to a minimum of four.

If both Net queue and default queue based RSS rings are to be created, DRSS rings are given priority.

The following table shows the number of RSS and DRSS rings created, corresponding to the value of RSS and DRSS module parameters.

**Table 8 RSS and DRSS Rings**

RSS	DRSS	No. of RSS Rings Created <sup>a</sup>	No. of DRSS Rings Created <sup>b</sup>
$U^c   1$	$U^c   1$	4	4
0	0	0	0
1	0	4	0
0	1	0	4
$x \geq 2$	0	$x, \max 8$	0
$x \geq 2$	$U^c   1$	$x, \max 8$	4
$U^c   1$	$y \geq 2$	4	$y, \max 8$
0	$y \geq 2$	0	$y, \max 8$
$x \geq 2$	1	$x, \max 8$	4
$x \geq 2$	$y \geq 2$	$x, \max 8$	$y, \max 8$

- a. Net queue RSS queues include leading Net queue.

- 
- b. DRSS rings exclude default Q.
  - c. Unspecified.

### 3.4.3 Usage

If there are two Broadcom Emulex adapters installed in the host, then the commands to enable or disable RSS and DRSS queues are as follows:

<b>NOTE</b>	If RSS or DRSS is to be disabled for a function, it must be specified as zero for that function.  If the value is unspecified for an adapter, then by default RSS and DRSS are enabled with four rings.
-------------	---

To enable RSS with four rings on each function and disable DRSS, type the following command and press Enter:

```
esxcfg-module elxnet -s "rss=4,4,4,4 drss=0,0,0,0"
```

To enable DRSS with four rings on each function and disable RSS, type the following command and press Enter:

```
esxcfg-module elxnet -s "rss=0,0,0,0 drss=4,4,4,4"
```

To enable DRSS and RSS with four rings each, type the following command and press Enter:

```
esxcfg-module elxnet -s "rss=4,4,4,4 drss=4,4,4,4"
```

To enable DRSS and RSS with the default number of rings (four rings), type the following command and press Enter:

```
esxcfg-module elxnet -s "rss=1,1,1,1 drss=1,1,1,1"
```

To enable DRSS with four rings and RSS with eight rings, type the following command and press Enter:

```
esxcfg-module elxnet -s "rss=8,8,8,8 drss=4,4,4,4"
```

<b>NOTE</b>	In case of insufficient resources, the maximum possible RSS queues are created after reducing Net queues to a minimum of four.
-------------	--

To disable both DRSS and RSS, type the following command and press Enter:

```
esxcfg-module elxnet -s "rss=0,0,0,0 drss=0,0,0,0"
```

### 3.4.4 Limitations

RSS with dynamic Net queue is not supported.

#### 3.4.5 Limitations Exclusive to OCe11000-Series Adapters

- RSS and DRSS features do not support UDP traffic.
- Only one RSS pool is supported per port. Therefore, either DRSS or RSS is supported, but not both.
- The leading queue does not participate in RSS traffic.
- Only the first physical function on the port supports RSS in Multi-Channel mode, due to limited resources.
- RSS is not supported when SR-IOV is enabled on the adapter.

## 3.5 iSCSI Driver Configuration

This section describes how to configure parameters for the ESXi iSCSI driver.

### 3.5.1 Configuring iSCSI Driver Options

The following table lists the user configurable iSCSI driver options. It includes a description of the parameters, default values, and the limits within which they can be configured.

**NOTE** If the value given for a parameter is outside the supported range (Minimum and Maximum values), the driver will log an error in the Event Log and continue to load using the default value of the parameter.

**Table 9 User-configurable iSCSI Driver Options**

Parameter	Default Value	Minimum Value	Maximum Value	Description
LDTO	20 seconds	0 seconds	3600 seconds	Link Down Timeout (in seconds) This determines the amount of time the initiator driver waits for the controller's physical link to become available before reporting that the LUNs are unavailable to the operating system.
ETO	30 seconds	0 seconds	3600 seconds	Extended Timeout (in seconds) This determines the amount of time the initiator driver waits for the target to become available after it has lost connection to the target during an I/O operation. <b>NOTE</b> If the minimum value is set between 0 and 19, the driver assumes a value of 20 internally. No modifications are seen in the esxcfg-module -g be2iscsi command output.
im_policy	2	0	4	Controls the rate of interrupts for the adapter. For more information, see <a href="#">Section 3.5.2, Interrupt Moderation Policy Settings</a> .
large I/O	128	128	512	Maximum transfer size in a single I/O request in kilobytes. By default, the iSCSI driver supports a maximum of 128 KB of data in a single I/O request. This option can be used to enable support for 512 KB of data in a single I/O request.
log_level	0x00	—	0xff	Enables the logging of debug information in system logs for a specific path, or all paths, in the driver.

The following command line shows how to configure the driver with an LDTO value of 25 seconds:

```
# vmkload_mod be2iscsi ldto=25
```

To configure the ESXi Server to load the iSCSI driver with this value after each reboot, run the following commands and reboot the system:

```
# esxcfg-module -s "ldto=25" be2iscsi
# /usr/sbin/esxcfg-boot -r
# reboot
```

The im\_policy (Interrupt Moderation policy) parameter configures the Broadcom Emulex iSCSI driver to use different settings for Interrupt Moderation. An im\_policy value of 1 achieves the highest interrupt rate, whereas the value 4 provides the lowest interrupt rate. The default value is 2. An im\_policy of 0 turns off the Interrupt Moderation algorithm in the driver.

The large\_io option can be used to modify the maximum transfer size in a single SCSI command. By default, the Broadcom Emulex iSCSI driver in ESXi supports up to 128 KB and 32 Scatter Gather entries in a single SCSI command. If applications issue I/O requests larger than 128 KB or need more than 32 Scatter Gather entries, the request will be split into multiple requests by the driver. By specifying large\_io=512; the iSCSI driver can support up to 512 KB of data and a total of 128 Scatter Gather entries in a single SCSI command.

**NOTE**

By setting the option to 512, the amount of physical memory consumed by the driver increases. Also, though intermediate values between 64 and 512 will be accepted, the memory used by the driver will be the same as what will be used for large\_io=512.

To set the large\_io parameter in ESXi, type the following commands:

```
esxcfg-module -s large_io=512 be2iscsi --> to set the parameter  
esxcfg-module -g be2iscsi --> to view the parameter value  
esxcfg-module -i be2iscsi --> to view the options
```

Even though the VMware operating system can be tuned to accept larger I/O sizes from guest operating systems, the guest operating systems will sometimes need to be tuned to create those larger I/Os. For example, a Windows 32-bit guest operating system will by default have a maximum transfer size of 64 KB, even if the VMware kernel is tuned to allow 512-KB I/O transfers. You must modify the guest registry to achieve the maximum transfer rate set in the VMware kernel.

To do this, use the registry editor to add or modify the following entry to change the maximum transfer size for the Windows guest:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Symmpi\Parameters\Device\MaximumSGList
```

The formula used to calculate the proper value for MaximumSGList is:

- For a 32-bit Windows guest use the following formula:

```
MaximumSGList = ((Maximum Transfer Size) / 4) + 1
```

- For a 64-bit Windows guest use the following formula:

```
MaximumSGList = ((Maximum Transfer Size) / 8) + 1
```

For example, to allow a 256 KB transfer size on a 32-bit guest, this would be the formula to use:  $256 / 4 = 64 + 1 = 65$  (decimal) or 0x41(hexadecimal).

The maximum value allowed for MaximumSGList is 255 or 0xFF. For the particular value of 0xFF, the internal value passed to Windows is increased to 0x101, allowing support for a full 1-MB transfer (2 MB for 64-bit).

The MaximumSGList is 129 for a 512-KB transfer.

Other guest operating systems may also need to be tuned similarly. Consult the tuning guides that accompany those operating systems.

### 3.5.2

### Interrupt Moderation Policy Settings

Interrupt Moderation Policy settings control the rate of interrupts for the adapter. By default, the driver implements an Interrupt Moderation scheme based on the I/O load and the interrupt rate. The default setting for the Interrupt

Moderation Policy attempts to vary the interrupt rate between 3500 to 10000 interrupts per second. In addition, the Broadcom Emulex iSCSI driver allows other configuration settings which are listed in the following table.

**Table 10 Interrupt Moderation Policy Settings**

Setting	Parameter	Description
Disabled	im_policy=0	Interrupt Moderation algorithm is turned off in the driver.
Aggressive	im_policy=1	Achieves the highest interrupt rate among all available settings.
Moderate	im_policy=2	This is the default value.
Conservative	im_policy=3	Achieves a lower interrupt rate than Moderate.
Very Conservative	im_policy=4	Achieves the minimum interrupt rate among all available settings.

While the default setting of Moderate may work for most configurations, there are instances when the setting may need to be changed. Changing the Interrupt Moderation Policy setting should be based on the initiator system configuration, the number of iSCSI targets that will be connected, the I/O load, and the throughput and latency offered by these iSCSI targets.

On systems capable of sustaining a higher interrupt rate and if the number of targets that will be connected is less (up to 8), the more Aggressive setting will result in lower latency and higher values of I/O operations per second (IOPs). But the higher interrupt rate could also result in system stalls and freezes, especially during higher queue depth values and smaller sized I/O requests.

On a configuration that involves a large number of iSCSI targets (more than 32 or 64) and higher values of queue depth, the default setting may prove to be too aggressive and the Interrupt Moderation setting may need to be changed to Conservative or Very Conservative. Though this will increase latency of an I/O request, the lower interrupt rate may allow the system to function under a high load.

### 3.5.2.1 iSCSI Error Handling

The goal of iSCSI error handling is to tolerate link level and target level failures up to configured timeout values so I/O errors are not seen by the application or operating system. Error handling is triggered under the following conditions:

- Loss of an immediate link to the initiator (for example, a cable disconnect or a port failure).  
The firmware detects and notifies the driver of a loss of the link. When this happens, the driver will queue the I/O requests internally up to a configured timeout period so the operating system does not see I/O errors. This timeout is known as LDTO.
- Loss of connection to the target due to target or network disconnection at the target.  
If the driver has I/O requests pending with the target and the target becomes unavailable (due to target going down or failing over, or network issues at the target), the driver queues up the I/O request internally up to a configured timeout period. This timeout is known as ETO.

When the configured threshold for LDTO and ETO is reached and the initiator is still unable to connect to the target, the driver fails all I/O requests. At this point, I/O errors will be seen by the application and operating system.

#### NOTE

Following a link up, switch ports can take a long time to initialize and go to a forwarding state. Because of this, additional time should be added to the ETO and LDTO settings to eliminate I/O disruption and/or target unavailability. If the switch port is connected to a single host, then PortFast mode can be enabled on the switch port to eliminate delays in transitioning to a forwarding state.

### 3.5.3 Configuring LDTO and ETO on ESXi Server

The following table lists the default values of LDTO and ETO on ESXi Server and the limits within which they can be configured.

**Table 11 LDTO and ETO Default Values on ESX Server**

Parameter	Default Value	Minimum Value	Maximum Value
LDTO	20 seconds	0 seconds	3600 seconds
ETO	30 seconds	0 seconds	3600 seconds

**NOTE**

If the value of ETO is set to a number between 0 and 19, the driver assumes a value of 20 seconds internally. You will not see any modification to the esxcfg-module -g be2iscsi command output.

LDTO and ETO values are configurable during insmod time. The ETO value specified during insmod is the default ETO value applied to all targets.

The following command line shows how to configure the driver with an LDTO value of 25 seconds:

```
# vmkload_mod be2iscsi ldto=25
```

To configure ESXi Server to load the iSCSI driver with this value after each reboot, run the following commands and reboot the system:

```
# esxcfg-module -s "ldto=25" be2iscsi
# /usr/sbin/esxcfg-boot -r
# reboot
```

### 3.5.4 Multipath I/O Support

This section describes the installation and login processes for multipath I/O support.

#### 3.5.4.1 Configuring and Enabling Support for ESXi MPIO on Non-boot Targets

To configure and enable support for ESXi MPIO on non-boot targets:

1. Connect your configuration for multipath.
2. Log into your targets with all paths using vSphere Client or iSCSISelect. For more information about using vSphere Client, refer to [Section 3.5.4.3, Logging into Targets Using vSphere Client](#). For information about using iSCSISelect, see the *Boot Manual for Emulex OneConnect Adapters* or the *Boot Manual for Emulex LightPulse Adapters*.

#### 3.5.4.2 Configuring and Enabling Support for ESXi MPIO on Boot Targets

To configure and enable support for ESXi MPIO on boot targets:

1. Connect your configuration for multipath. Multipath can be configured before or after installation.
2. Log into your targets with all paths using vSphere Client or iSCSISelect. For more information about using vSphere Client, refer to [Section 3.5.4.3, Logging into Targets Using vSphere Client](#). For information about using iSCSISelect, see the *Boot Manual for Emulex OneConnect Adapters* or the *Boot Manual for Emulex LightPulse Adapters*.

### 3.5.4.3 Logging into Targets Using vSphere Client

To log into targets using the vSphere Client:

1. Log into the server you would like to configure from the vSphere Client and select the Configuration tab.
2. In the Hardware section, select the Storage Adapters link.
3. Select the OneConnect host adapter in the Storage Adapters list.
4. Select the initiator port, and then click Properties.
5. From the iSCSI Initiator Properties screen, click the Dynamic Discovery tab.
6. Click Add to display the Add Send Target Server menu.
7. Type in the IP address of the first target portal and click OK.
8. The initiator finds the target machines.
9. Click the Static Discovery tab to find all target portals.
10. Click Close. A dialog box indicates that a rescan is needed.
11. Click Yes. The configured LUNs are listed.
12. Repeat the steps to log into the other target portal to set up MPIO.
13. To check that both paths are connected to the same LUN, follow these steps:
  - a. Select the LUN and right-click.
  - b. Click Manage Paths. Multipaths on the LUN are shown.

### 3.5.4.4 Error Handling Under MPIO and Cluster Configurations

In an MPIO or cluster configuration, fault tolerant software is present on the system that makes the iSCSI driver error handling redundant. These configurations also require that I/O errors be reported as soon as they are detected, so the software can failover to an alternate path or an alternative node as quickly as possible.

When the iSCSI driver is run under these configurations, the error handling implemented in the iSCSI driver must be turned off by setting the default value of LDTO and ETO to 0. The changes will take effect during the next driver load.

## 3.5.5 Reading the Driver Statistics for a Specified Port

You can read the iSCSI driver's statistics for a specific port on the adapter to diagnose potential performance issues.

To read the driver statistics for a particular port, use the following command:

```
cat /proc/be2iscsi/be2iscsi<XX>/driver_stats
```

where XX is the specific port number.

## 3.5.6 Setting the Log Level

The log level setting enables the logging of debug information in system logs for a specific path, or all paths, in the driver.

Use one of the following methods to set the log level:

1. User configurable driver parameters – use either of the following commands to set the log level.

**NOTE** The new log level value will take effect after the next reboot.

```
# esxcfg-module -s "log_level=0xff" be2iscsi
-Or-
# vmkload_mod be2iscsi "log_level=0xff"
```

- 
2. 'log\_level' proc entry setting – use the following commands to read and write values to the log\_level proc entry.
    - cat /proc/be2iscsi/log\_level – This command displays the current log level and help.
    - echo 0xff > /proc/be2iscsi/log\_level – This command sets the log\_level value in the driver to enable the desired logging.

See [Table 12](#) for a list of available log level values.

**Table 12 Log Level Values**

log_level Value	Description
0x01	Enables logging in the iSCSI driver initialization and unload path.
0x02	Enables logging in the iSCSI driver I/O path.
0x04	Enables logging in the iSCSI driver error handling path.
0x08	Enables logging in the iSCSI driver configuration path, such as adding targets, deleting targets, and so on.
0x10	Enables logging in the iSCSI driver IOCTL path.

## Chapter 4: Troubleshooting

Your system may operate in an unexpected manner in certain circumstances. This section explains many of these circumstances and offers one or more workarounds for each situation.

### 4.1 Troubleshooting the FC/FCoE Driver

This section provides troubleshooting information for the FC/FCoE driver.

[Table 13](#), identifies some of the common situations and their potential resolutions.

**Table 13 Troubleshooting the FC/FCoE Driver**

Situation	Resolution
Port link fails to come up.	If an FC link fails to come up, verify that an 8GFC adapter is not trying to connect to a 1GFC device. Only 2, 4, and 8GFC devices are supported on 8 GFC adapters. 8 and 16 GFC devices are supported on 16 GFC adapters. 4 GFC devices are conditionally supported on 16 GFC adapters. If you want 4 GFC support on a 16 GFC adapter, contact Broadcom customer support.  For FCoE adapters, you make sure the vfc and vlan are correctly configured for all FCoE links. If the adapter firmware and FCoE switch do not complete FIP correctly, an FCF is not presented to the driver and the virtual link for FCoE will not come up.
The Broadcom Emulex driver is not loaded and all paths are down.	Use the lspci utility to determine if the Broadcom Emulex ports are being properly identified. If not, find out if the driver ISO was correctly installed. You must have the correct driver for the installed adapter because the device PCI IDs are installed with the driver package.  Examine the /var/log/vmkernel.log file for lpfc or lpfc820 log messages indicating an error. If you specified driver logging (refer to <a href="#">Section 3.2, FC/FCoE Driver Configuration</a> ) make sure you spelled the driver parameters correctly. The ESX module subsystem will not load the driver on reboot if the parameters are not spelled correctly. In this case, contact Broadcom support.
All paths are down.	Use the driver's KV pages to get critical information.  First check the link state. The KV command shown in the footnote shows the driver's current link, whether or not it has found a fabric and the link speed. <sup>a</sup>  If the data shows Link Up -Ready and Mode Online, check the discovered nodes. Fabric, Initiator and Target types show the SAN as it was presented and discovered by the driver. Additionally, if you are experiencing periodic path outage, the command shown in the second footnote in a script loop will show if the node status is changing or if the node is going offline. <sup>b</sup>  If your target or initiator is not in the driver's discovered list, check your zone membership and the state of all zone members.  Contact Broadcom customer support if you are unable to resolve missing zone members.

**Table 13 Troubleshooting the FC/FCoE Driver**

Situation	Resolution
The FC/FCoE driver fails to recognize an adapter and logs “ <i>unknown IOCB</i> ” messages in the system log during driver load. The adapter is running outdated firmware.	Download and install the adapter firmware that complies with the minimum supported version (or later) listed on the Avago website.
System panics when booted with a failed adapter installed.	Remove the failed adapter and reboot.
The FC/FCoE driver does not discover all remote ports in the configuration switch zone. Some initiators or targets may appear to be missing.	Evaluate your switch zone. Count how many entries there are and add at least seven more (to account for fabric logins). If the sum exceeds 128, you must increase the driver's node table size. The following commands increase it to 200 entries. Refer to <a href="#">Section 3.2, FC/FCoE Driver Configuration</a> , for more information on this driver parameter.  Globally: <code>esxcli system module parameters set -p lpfc_nlp_slab_cnt=200 -m lpfc</code> Per instance: <code>esxcli system module parameters set -p lpfc0_nlp_slab_cnt=200 -m lpfc</code> A reboot is required.

a. Data from KV command 1:

```
[root@chara:~] /usr/lib/vmware/vmkmgmt_keyval/vmkmgmt_keyval -i vmhba3/Emulex -k adapter -g
Key 'adapter':
lpfc Adapter Page
```

```
Emulex LightPulse FC SCSI 11.0.206.6000
Emulex LPe12002-M8 8Gb 2-port PCIe Fibre Channel Adapter on PCI bus 0000:04 device 00 fn 1 port 1 Link
Speed: 8 Gb
```

```
BoardNum: 3
FW Version: 2.02X11
HW Version: 31004549
ROM Version: 5.12a5
SerialNum: VM21932214
Vendor Id: f10010df
```

```
SLI Rev: 3
MQ: Unavailable
NPIV Supported: VPIs max 255 VPIs used 0
RPIs max 4096 RPIs used 18 IOCBs inuse 0 IOCB max 16 txq cnt 0 txq max 0 txcmplq 0
XRIIs max 4096 FCP 320
FCP BDEs max 66 DMA buf size 1008
```

```
Queue Depth
LUN 30
HBA FCP 2048
```

```
PCI read error: 0 retry attempts: 0
```

```
Link Up - Ready:
EDTOV 2000 ms RATOV 10 sec
PortID 0x20700
Fabric
Current speed 8G
```

```
WWPN 10:00:00:00:c9:f4:48:af WWNN 20:00:00:00:c9:f4:48:af
```

```
Mode: Online
```

b. Data from KV command 2:

```
[root@chara:~] /usr/lib/vmware/vmkmgmt_keyval/vmkmgmt_keyval -i vmhba11/Emulex -k node -g
Key 'node':
 1pfc Node page:
```

WWNN	WWPN	ScsiID	
DID	Type	Status	
10:00:00:05:33:a6:a8:bb	20:06:00:05:33:a6:a8:bb	xffffffe	Fabric
10:00:00:05:33:a6:a8:bb	21:fc:00:05:33:a6:a8:bb	xffffffc	Fabric
00:00:00:00:00:00:00:00	00:00:00:00:00:00:00:00	xffffffd	Fabric
20:00:00:00:c9:f3:f9:b6	10:00:00:00:c9:f3:f9:b6	x021600	Initiator
20:00:00:00:c9:f4:48:af	10:00:00:00:c9:f4:48:af	x020700	Initiator
20:00:00:90:fa:02:19:16	10:00:00:90:fa:02:19:16	x021500	Initiator
20:00:00:90:fa:02:19:17	10:00:00:90:fa:02:19:17	x020e00	Initiator
20:00:00:90:fa:5d:2f:29	10:00:00:90:fa:5d:2f:29	x020902	Initiator
20:00:00:90:fa:5d:2f:31	10:00:00:90:fa:5d:2f:31	x020903	Initiator
20:00:00:90:fa:5d:92:2f	10:00:00:90:fa:5d:92:2f	x02094a	Initiator
20:00:00:11:0d:e0:54:00	20:00:00:11:0d:e0:54:00	0 x02090e	Target
20:06:00:11:0d:19:9a:00	20:06:00:11:0d:19:9a:00	1 x021000	Target
20:07:00:11:0d:19:9b:00	20:07:00:11:0d:19:9b:00	2 x021100	Target

#### 4.1.1 FC/FCoE Driver Log Messages

Log messages have traditionally been organized into logical groups based on code functionality in the FC driver. With the introduction of the latest Broadcom Emulex adapters, that grouping is modified to account for additional behaviors. The traditional grouping is maintained, but recently added messages are no longer grouped together.

The messages provided in this section are unmaskable error conditions. They are automatically added to the system console log.

You can examine the /var/log/vmkernel file.log to see any of these messages. If you have concerns, the best policy is to run a vm-support dump and contact VMware or Broadcom support staff.

Log messages are organized into logical groups based on code functionality within the driver. Each group consists of a block of 100 log message numbers. Most groups require a single block of 100 message numbers; however, some groups (INIT, FCP) require two blocks.

[Table 14](#) lists the groups and defines the associated number ranges for ESXi 5.1 systems.

[Table 15](#) lists the groups and defines the associated number ranges for ESXi 5.5 and 6.0 systems.

**Table 14 Message Log Table for ESXi 5.1 Systems**

LOG Message Verbose Mask Definition	Verbose Bit	Verbose Description
LOG_ELS	0x1	Extended link service events
LOG_DISCOVERY	0x2	Link discovery events
LOG_MBOX	0x4	Mailbox events
LOG_INIT	0x8	Initialization events
LOG_LINK_EVENT	0x10	Link events
LOG_FCP	0x40	FCP traffic history
LOG_NODE	0x80	Node table events
LOG_TEMP	0x100	Temperature sensor events
LOG_MISC	0x400	Miscellaneous and FCoE events

**Table 14 Message Log Table for ESXi 5.1 Systems (Continued)**

LOG Message Verbose Mask Definition	Verbose Bit	Verbose Description
LOG_SLI	0x800	SLI events
LOG_FCP_ERROR	0x1000	Selective FCP events
LOG_LIBDFC	0x2000	IOCTL events
LOG_VPORT	0x4000	NPIV events
LOG_EVENT	0x10000	IOCTL event
LOG_DAEMON	0x20000	IOCTL Daemon events
LOG_FIP	0x40000	FIP event
LOG_PROC	0x80000	Procfs events
LOG_FCP_UNDERRUN	0x100000	FCP underruns
LOG_TRC_NODE	0x00200000	Node Trace Events
LOG_ALL_MSG	0x7fffffff	Log all messages

**Table 15 Message Log Table for ESXi 5.5 and 6.0 Systems**

LOG Message Verbose Mask Definition	Verbose Bit	Verbose Description
LOG_ELS	0x1	Extended link service events
LOG_DISCOVERY	0x2	Link discovery events
LOG_MBOX	0x4	Mailbox events
LOG_INIT	0x8	Initialization events
LOG_LINK_EVENT	0x10	Link events
LOG_MGMT_ERROR	0x20	IODM error logging
LOG_FCP	0x40	FCP traffic history
LOG_NODE	0x80	Node table events
LOG_TEMP	0x100	Temperature sensor events
LOG_BG	0x200	BlockGuard events
LOG_MISC	0x400	Miscellaneous and FCoE events
LOG_SLI	0x800	SLI events
LOG_FCP_ERROR	0x1000	Selective FCP events
LOG_LIBDFC	0x2000	IOCTL events
LOG_VPORT	0x4000	NPIV events
LOG_SECURITY	0x8000	Security events
LOG_EVENT	0x10000	IOCTL event
LOG_FIP	0x20000	FIP event
LOG_FCP_UNDER	0x40000	FCP underrun errors
LOG_KVPAGE	0x80000	KV page verbose
LOG_TASKMGMT	0x100000	Task management events
LOG_MGMT_TRACE	0x200000	IODM management trace logging
LOG_ALL_MSG	0x7fffffff	Log all messages

#### 4.1.1.1 ESXi 5.5 and 6.0 Message Log Example

The following is an example of a LOG message on ESXi 5.5 and 6.0 systems:

```
#define LOG_ELS          0x00000001      /* ELS events */  
#define LOG_DISCOVERY    0x00000002      /* Link discovery events */  
#define LOG_MBOX          0x00000004      /* Mailbox events */  
#define LOG_INIT          0x00000008      /* Initialization events */  
#define LOG_LINK_EVENT    0x00000010      /* Link events */  
#define LOG_MGMT_ERROR    0x00000020      /* IODM management error logging */  
#define LOG_FCP           0x00000040      /* FCP traffic history */  
#define LOG_NODE          0x00000080      /* Node table events */  
#define LOG_TEMP          0x00000100      /* Temperature sensor events */  
#define LOG_BG            0x00000200      /* BlockGuard events */  
#define LOG_MEM_HEAP      0x00000400      /* Mem tracker heap logging */  
#define LOG_SLI           0x00000800      /* SLI events */  
#define LOG_FCP_ERROR     0x00001000      /* log errors, not underruns */  
#define LOG_LIBDFC        0x00002000      /* Libdfc events */  
#define LOG_VPORT         0x00004000      /* NPIV events */  
#define LOG_MEM_SLAB      0x00008000      /* Mem tracker slab logging */  
#define LOG_EVENT         0x00010000      /* CT,TEMP,DUMP, logging */  
#define LOG_FIP           0x00020000      /* FIP events */  
#define LOG_FCP_UNDER     0x00040000      /* FCP underruns errors */  
#define LOG_KVPAGE        0x00080000      /* KV page verbose */  
#define LOG_TASKMGMT      0x00100000      /* Task Management events. */  
#define LOG_MGMT_TRACE    0x00200000      /* IODM management trace logging */  
#define LOG_SCSI_CMD      0x00400000      /* ALL SCSI commands */  
#define LOG_EDIF          0x00800000      /* External DIF events */  
#define LOG_KV_ERROR      0x01000000      /* Key-value mgmt error logging */  
#define LOG_ESXCLI        0x02000000      /* Esxcli mgmt interface logging */  
#define LOG_ALL_MSG       0x7fffffff      /* LOG all messages */
```

In the above LOG message:

- lpfc - driver binary
- lpfc\_mbx\_cmpl\_read\_topology - function generating the log
- 1: identifies Broadcom Emulex HBA1.
- 1305 identifies the LOG message number.

**NOTE**

If the word 'Data:' is present in a LOG message, any information to the right of 'Data:' is intended for Broadcom technical support/engineering use only.

Unless otherwise noted in the ACTION: attribute, report these errors to Broadcom technical support. Broadcom requests that when reporting occurrences of these error messages, you provide a tarball of all vmkernel files in /var/log.

#### 4.1.1.2 ESXi 5.1 Message Log Example

The following is an example of a LOG message on ESXi 5.1 systems:

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

In the above LOG message:

- lpfc 0000:03:06.0: identifies the PCI location of the particular FC/FCoE HW port.
- 0: identifies Broadcom Emulex HBA0.
- 1305 identifies the LOG message number.

**NOTE** If the word 'Data:' is present in a LOG message, any information to the right of 'Data:' is intended for Broadcom technical support/engineering use only.  
Unless otherwise noted in the ACTION: attribute, report these errors to Broadcom technical support. Broadcom requests that when reporting occurrences of these error messages, you provide a tarball of all vmkernel files in /var/log.

## 4.1.2 Extended Link Service Events (0100 to 0199)

### 4.1.2.1 **elx\_mes0100: FLOGI failure Status:<status>/<extended\_status> TMO:<timeout>**

DESCRIPTION: An extended link service FLOGI command that was sent to the fabric failed.

DATA: (1) ulpStatus, (2) ulpWord[4], (3) ulpTimeout

ACTION: This error could indicate a fabric configuration error or an internal driver issue. If this issue persists, report the error to Broadcom technical support.

### 4.1.2.2 **elx\_mes0111: Dropping received ELS cmd**

DESCRIPTION: The driver dropped an extended link service response ring entry.

DATA: (1) ulpStatus, (2) ulpWord[4], (3) ulpTimeout

ACTION: This error could indicate a software driver or a firmware issue. If this issue persists, report the error to Broadcom technical support.

### 4.1.2.3 **elx\_mes0113: A FLOGI ELS command <elsCmd> was received from DID <did> in Loop Mode**

DESCRIPTION: While in Loop Mode an unknown or unsupported extended link service command was received.

DATA: None

ACTION: Check the device DID.

### 4.1.2.4 **elx\_mes0115: Unknown ELS command <elsCmd> received from N\_Port <did>**

DESCRIPTION: Received an unsupported extended link service command from a remote N\_Port.

DATA: None

ACTION: Check the remote N\_Port for a potential issue.

### 4.1.2.5 **elx\_mes0122 FDISC Failed (value). Fabric Detected Bad WWN**

DESCRIPTION: The driver's F\_Port discovery failed. The switch reported a bad WWN in the FLOGI request.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

---

**4.1.2.6 elx\_mes0124 Retry illegal cmd <value> retry:<value> delay:<value>**

DESCRIPTION: The port rejected an extended link service command as illegal. The driver is retrying.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.2.7 elx\_mes0125: FDISC Failed (value). Fabric out of resources**

DESCRIPTION: The fabric rejected an F\_Port discovery because the switch cannot support any more virtual ports.

DATA: None

ACTION: Reconfigure the switch to support more NPIV logins. If this issue persists, contact Broadcom technical support.

**4.1.2.8 elx\_mes0126: FDISC failed (ulpStatus/ulpWord[4])\n**

DESCRIPTION: The ELS F\_Port discovery command has failed.

DATA: None

ACTION: Check the port and switch configuration.

**4.1.2.9 elx\_mes0127: ELS timeout**

DESCRIPTION: 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) ulploTag

ACTION: If the extended link service command is not going through the adapter, reboot the system. If this issue persists, report the error to Broadcom technical support.

**4.1.2.10 elx\_mes0133: PLOGI: no memory for reg\_login**

DESCRIPTION: Memory allocation error.

DATA: (1) nlp\_DID, (2) nlp\_state, (3) nlp\_flag, (4) nlp\_rpi

ACTION: This is a memory allocation error. Check the system resources and unload any unused modules.

**4.1.2.11 elx\_mes0134: PLOGI: cannot issue reg\_login**

DESCRIPTION: The ELS port login mailbox command has failed.

DATA: (1) nlp\_DID, (2) nlp\_state, (3) nlp\_flag, (4) nlp\_rpi

ACTION: Check the port and switch configuration.

**4.1.2.12 elx\_mes0135: cannot format reg\_login**

DESCRIPTION: The system 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

ACTION: None required.

**4.1.2.13 elx\_mes0136: PLOGI completes to N\_Port < DID > completion**

DESCRIPTION: A port login has completed for which there is no NDLP.

DATA: (1) ulpStatus, (2) ulpWord[4]

ACTION: None required.

---

**4.1.2.14 elx\_mes0137: No retry ELS command <ELS\_CMD> to remote**

DESCRIPTION:

DATA: (1) ulpStatus, (2) ulpWord[4]

ACTION: None required.

**4.1.2.15 elx\_mes0138: ELS rsp: Cannot issue reg\_login for <DID>**

DESCRIPTION: The REG\_LOGIN mailbox command failed.

DATA: (1) nlp\_DID, (2) nlp\_state, (3) nlp\_flag, (4) nlp\_rpi

ACTION: None required.

**4.1.2.16 elx\_mes0140: PLOGI Reject: invalid nname**

DESCRIPTION: An invalid node WWN was provided.

DATA: None

ACTION: None required.

**4.1.2.17 elx\_mes0141: PLOGI Reject: invalid pname**

DESCRIPTION: An invalid port WWN was provided.

DATA: None

ACTION: None required.

**4.1.2.18 elx\_mes0142: PLOGI RSP: Invalid WWN**

DESCRIPTION: The port login sent to the port by a remote port had an invalid WWN.

DATA: None

ACTION: None required.

**4.1.2.19 elx\_mes0144: Not a valid WCQE code: <Completion Code>**

DESCRIPTION: The completion queue handler detected an invalid type.

DATA: None

ACTION: None required.

**4.1.2.20 elx\_mes0147: Failed to allocate memory for an RSCN event**

DESCRIPTION: Memory could not be allocated to send the RSCN event to the management application.

DATA: None

ACTION: None required.

**4.1.2.21 elx\_mes0148: Failed to allocate memory for a LOGO event**

DESCRIPTION: Memory could not be allocated to send the N\_Port logout event to the FC transport.

DATA: None

ACTION: None required.

---

#### 4.1.2.22 **elx\_mes0154: Authentication not complete**

DESCRIPTION: Memory could not be allocated to send the N\_Port logout event to the FC transport.

DATA: None

ACTION: None required.

### 4.1.3 **Link Discovery Events (0200 to 0299)**

#### 4.1.3.1 **elx\_mes0200: CONFIG\_LINK bad hba state <hba\_state>**

DESCRIPTION: A CONFIG\_LINK mailbox command completed and the driver was not in the right state.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

#### 4.1.3.2 **elx\_mes0203: Devloss timeout on WWPN <address> N\_Port <nlp\_DID>**

DESCRIPTION: 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

ACTION: If the device generating this message is not a target to which the adapter is connected, this error does not affect the data integrity of the I/O between the adapter and the attached storage and can be ignored.

#### 4.1.3.3 **elx\_mes0206: Device discovery completion error**

DESCRIPTION: This indicates that an uncorrectable error was encountered during device (re)discovery after a link up. FC devices are not accessible if this message is displayed.

DATA: None

ACTION: Reboot the system. If the issue persists, report the error to Broadcom technical support. Run with verbose mode on for more details.

#### 4.1.3.4 **elx\_mes0207: Device <DID> (<WWN>) sent invalid service parameters. Ignoring device.**

DESCRIPTION: Invalid service parameters were received from the DID. Ignoring this remote port.

DATA: DID, WWN

ACTION: Verify the remote port's configuration. If the issue persists, report the error to Broadcom technical support. Run with verbose mode on for more details.

#### 4.1.3.5 **elx\_mes0222: Initial FLOGI/FDISK timeout**

DESCRIPTION: The driver sent the initial FLOGI or FDISK to the fabric and never got a response back.

DATA: None

ACTION: Check the fabric configuration. The driver recovers from this and continues with device discovery.

#### 4.1.3.6 **elx\_mes0223: Timeout while waiting for NameServer login**

DESCRIPTION: The login request to the NameServer was not acknowledged within R\_A\_TOV.

DATA: None

ACTION: Check the fabric configuration. The driver recovers from this and continues with device discovery.

---

#### 4.1.3.7 **elx\_mes0224: NameServer Query timeout**

DESCRIPTION: 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

ACTION: Check the fabric configuration. The driver recovers from this and continues with device discovery.

#### 4.1.3.8 **elx\_mes0227: Node Authentication timeout**

DESCRIPTION: The driver has lost track of what N\_Ports are being authenticated.

DATA: None

ACTION: None required. The driver should recover from this event.

#### 4.1.3.9 **elx\_mes0228: CLEAR LA timeout**

DESCRIPTION: The driver issued a CLEAR\_LA that never completed.

DATA: None

ACTION: None required. The driver should recover from this event.

#### 4.1.3.10 **elx\_mes0230: Unexpected timeout, hba linkstate <link\_state>**

DESCRIPTION: Discovery has timed out and the adapter state is not ready.

DATA: None

ACTION: None required.

#### 4.1.3.11 **elx\_mes0231: RSCN timeout**

DESCRIPTION: The driver has lost track of what N\_Ports have RSCNs pending.

DATA: (1) fc\_ns\_retry, (2) lpfc\_max\_ns\_retry

ACTION: None required. The driver should recover from this event.

#### 4.1.3.12 **elx\_mes0233: Nodelist not empty**

DESCRIPTION: Driver unloaded or hotplug detected a node still in use.

DATA: None

ACTION: None required.

#### 4.1.3.13 **elx\_mes0237: Pending Link Event during Discovery: State <state>**

DESCRIPTION:

DATA: None

ACTION: None required.

#### 4.1.3.14 **elx\_mes0241: NameServer Rsp Error Data: <data>**

DESCRIPTION:

DATA: None

ACTION: None required.

---

#### 4.1.3.15 **elx\_mes0246: RegLogin failed**

DESCRIPTION: The firmware returned a failure for the specified RegLogin.

DATA: (1) Did, (2) mbxStatus, (3) hbaState

ACTION: This message indicates that the firmware could not perform a RegLogin for the specified DID. A limitation on how many nodes an adapter can view may exist.

#### 4.1.3.16 **elx\_mes0249: Cannot issue Register Fabric login: Err <err>**

DESCRIPTION: The firmware could not issue the fabric reg login; the err value is unique for each possible failure.

DATA: None

ACTION: None required.

#### 4.1.3.17 **elx\_mes0251: NameServer login: no memory**

DESCRIPTION: The firmware could not allocate memory for the NDLP structure.

DATA: None

ACTION: None required.

#### 4.1.3.18 **elx\_mes0252: Cannot issue NameServer login**

DESCRIPTION: The firmware could not issue an extended link service port login to the nameserver DID.

DATA: None

ACTION: Check the port connection and the switch configuration.

#### 4.1.3.19 **elx\_mes0253: Register VPI: Can't send mbox**

DESCRIPTION: The system could not issue the REG\_LOGIN command for this vPort.

DATA: None

ACTION: None required.

#### 4.1.3.20 **elx\_mes0254: Register VPI: no memory goto mbox\_err\_exit**

DESCRIPTION: The system could not allocate memory for the REG\_LOGIN mailbox command.

DATA: None

ACTION: None required.

#### 4.1.3.21 **elx\_mes0255: Issue FDISC: no IOCB**

DESCRIPTION: All of the pre-allocated IOCBs are in use.

DATA: None

ACTION: None required.

#### 4.1.3.22 **elx\_mes0256: Issue FDISC: Cannot send IOCB**

DESCRIPTION: The system is unable to send the fabric IOCB.

DATA: None

ACTION: Check the switch configuration.

---

**4.1.3.23 elx\_mes0257: GID\_FT Query error: <ulpStatus> <fc\_ns\_retry>**

DESCRIPTION: The GID\_FT common transport request for the nameserver has failed.

DATA: None

ACTION: Check the switch configuration.

**4.1.3.24 elx\_mes0258: Register Fabric login error: <mbxStatus>**

DESCRIPTION: The REG\_LOGIN for the fabric has failed.

DATA: None

ACTION: Check the port configuration and the switch configuration.

**4.1.3.25 elx\_mes0259: No NPIVFabric support**

DESCRIPTION: The switch to which the port is connected does not support NPIV.

DATA: None

ACTION: Check the switch configuration.

**4.1.3.26 elx\_mes0260: Register NameServer error: <mbxStatus>**

DESCRIPTION: The REG\_LOGIN mailbox command has failed for the nameserver.

DATA: None

ACTION: Check the switch configuration.

**4.1.3.27 elx\_mes0261: Cannot Register NameServer login**

DESCRIPTION: Either a memory allocation issue or an invalid parameter was sent to the REG\_LOGIN.

DATA: None

ACTION: At least one message (0142, 0121, 0133, 0134, or 0135) should precede this message.

**4.1.3.28 elx\_mes0262: No NPIV Fabric support**

DESCRIPTION: The switch to which the port is connected does not support NPIV.

DATA: None

ACTION: Check the switch configuration.

**4.1.3.29 elx\_mes0263: Discovery Mailbox error: state: <port\_state>: <sparam\_mbox> <cfglink\_mbox>**

DESCRIPTION: 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.

ACTION: Try to unload and reload the driver when it is convenient.

**4.1.3.30 elx\_mes0264: No NPIV Fabric support**

DESCRIPTION: The switch to which the port is connected does not support NPIV.

DATA: None

ACTION: Check the switch configuration.

---

**4.1.3.31 elx\_mes0266: Issue NameServer Req <cmdcode> err <rc> Data: <fc\_flag> <fc\_rscn\_id\_cnt>**

DESCRIPTION: The driver was unable to send the nameserver common transport command.

DATA: (1) VPorts fc\_flag, (2) VPorts fc\_rscn\_id\_cnt

ACTION: Check the switch and port configurations.

**4.1.3.32 elx\_mes0267: NameServer GFF Rsp <did> Error (<ulpStatus> <un.ulpWord[4]>) Data: <fc\_flag> <fc\_rscn\_id\_cnt>**

DESCRIPTION: The nameServer GFF common transport request failed.

DATA: (1) VPorts fc\_flag, (2) VPorts fc\_rscn\_id\_cnt

ACTION: Check the switch and port configurations.

**4.1.3.33 elx\_mes0268: NS cmd <cmdcode> Error (<ulpStatus> <un.ulpWord[4]>)**

DESCRIPTION: The nameServer common transport request failed.

DATA: None.

ACTION: Check the switch and port configurations.

**4.1.3.34 elx\_mes0271: Illegal State Transition: node <nlp\_DID> event <evt>, state <nlp\_state> Data:<nlp\_rpi> <nlp\_flag>**

DESCRIPTION: The current node state does not have a handler for this event.

DATA: (1) nlp\_rpi, (2) nlp\_flag

ACTION: Verify that all targets are still visible to the SCSI mid-layer.

**4.1.3.35 elx\_mes0272: Illegal State Transition: node <nlp\_DID> event <evt>, state <nlp\_state> Data: <nlp\_rpi> <nlp\_flag>**

DESCRIPTION: The driver is completing a port login but does not have the rcv\_plogi flag set.

DATA: (1) nlp\_rpi, (2) nlp\_flag

ACTION: Verify that all targets are still visible to the SCSI mid-layer.

**4.1.3.36 elx\_mes0273: Unexpected discovery timeout,vport State <port\_state>**

DESCRIPTION: The discovery process has timed out.

DATA: None

ACTION: Verify that all targets are visible.

**4.1.3.37 elx\_mes0282: did:<value> ndlp:<value> pusgmap:<value> refcnt<value>, ndlp->nlp\_DID, (void \*)ndlp, lpfc\_init.c-ndlp->nlp\_usg\_map**

DESCRIPTION: Driver clean-up has found a node that is still on the node list during driver unload or PCI hotplug removal.

DATA: None.

ACTION: None required.

**4.1.3.38 elx\_mes0283: Failed to allocate mbox cmd memory**

DESCRIPTION: Mailbox allocation error.

DATA: None

ACTION: None required.

---

**4.1.3.39 elx\_mes0285: Allocated DMA memory size <alloclen> is less than the requested DMA memory size <reqlen>**

DESCRIPTION: Memory allocation was truncated.

DATA: None

ACTION: None required.

**4.1.3.40 elx\_mes0286: lpfc\_nlp\_state\_cleanup failed to allocate statistical data buffer <nlp\_DID>**

DESCRIPTION: Memory allocation failed for the node's statistical data.

DATA: None

ACTION: None required.

**4.1.3.41 elx\_mes0287: lpfc\_alloc\_bucket failed to allocate statistical data buffer <nlp\_DID>**

DESCRIPTION: Memory allocation failed for the node's statistical data.

DATA: None

ACTION: None required.

**4.1.3.42 elx\_mes0288: Unknown FCoE event type <event\_type> event tag <event\_tag>**

DESCRIPTION: The firmware has detected an unknown FCoE event.

DATA: None

ACTION: Check the FCoE switch configuration and the adapter DCBX mode.

**4.1.3.43 elx\_mes0289: Issue register VFI failed: Err <rc>**

DESCRIPTION: The driver could not register the virtual fabric index for the FCFI.

DATA: None

ACTION: Check the switch and port configurations.

**4.1.3.44 elx\_mes0290: The SLI4 DCBX asynchronous event is not handled yet**

DESCRIPTION: The SLI-4 DCBX asynchronous event is not handled yet.

DATA: None

ACTION: None required.

**4.1.3.45 elx\_mes0291: Allocated DMA memory size <alloc\_len> is less than the requested DMA memory size <req\_len>**

DESCRIPTION: The asynchronous DCBX events are not handled in the driver.

DATA: None

ACTION: Check the switch configuration.

**4.1.3.46 elx\_mes0293: PM resume failed to start workerthread: error=<error>**

DESCRIPTION: The PCI resume (hotplug) could not start the worker thread for the driver.

DATA: None

ACTION: Unload and reload the driver.

---

#### 4.1.3.47 **elx\_mes0294: PM resume failed to enable interrupt**

DESCRIPTION: The PCI resume (hotplug) could not get an interrupt vector.

DATA: None

ACTION: Unload and reload the driver.

#### 4.1.3.48 **elx\_mes0297:invalid device group <pci\_dev\_grp>**

DESCRIPTION: While unloading the driver, the driver detect a PCI device that it should not have claimed.

DATA: None

ACTION: None required.

#### 4.1.3.49 **elx\_mes0299: Invalid SLI revision <sli\_rev>**

DESCRIPTION: While processing a host attention error or an unrecoverable error, the driver detected an invalid SLI revision.

DATA: None

ACTION: None required.

### 4.1.4 **Mailbox Events (0300 to 0339)**

#### 4.1.4.1 **elx\_mes0300: LATT: Cannot issue READ\_LA: Data: <rc>**

DESCRIPTION: The link attention handler could not issue a READ\_LA mailbox command.

DATA: None

ACTION: None required.

#### 4.1.4.2 **elx\_mes0303: Ring <ringno> handler: portRspPut <portRspPut> is bigger then rsp ring <portRspMax>**

DESCRIPTION: The port rsp ring put index is larger than the size of the rsp ring.

DATA: None

ACTION: This error could indicate a software driver, firmware, or hardware issue. Report the error to Broadcom technical support.

#### 4.1.4.3 **elx\_mes0304: Stray mailbox interrupt, mbxCommand <mbxcommand> mbxStatus <mbxstatus>**

DESCRIPTION: Received a mailbox completion interrupt and there are no outstanding mailbox commands.

DATA: None

ACTION: This error could indicate a hardware or firmware issue. If the issue persists, report the error to Broadcom technical support.

#### 4.1.4.4 **elx\_mes0306: CONFIG\_LINK mbxStatus error <mbxStatus> HBA state <hba\_state>**

DESCRIPTION: The driver issued a CONFIG\_LINK mbox command to the adapter that failed.

DATA: None

ACTION: This error could indicate a firmware or hardware issue. Report the error to Broadcom technical support.

---

#### 4.1.4.5 **elx\_mes0310: Mailbox command <mbxcommand> timeout**

DESCRIPTION: 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

ACTION: This error could indicate a software driver or firmware issue. If no I/O is going through the adapter, reboot the system. If the issue persists, report the error to Broadcom technical support.

#### 4.1.4.6 **elx\_mes0311 Mailbox command <value> cannot issue data: <value> <value>**

DESCRIPTION: The driver detected an HBA error and can't issue the mailbox.

DATA: (1) sli flags (2) hba flags

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

#### 4.1.4.7 **elx\_mes0312: Ring <ringno> handler: portRspPut <rspPutIdx> is bigger then rsp ring <numRiocb>**

DESCRIPTION: The IOCB command ring put pointer is ahead of the get pointer.

DATA: None

ACTION: None required.

#### 4.1.4.8 **elx\_mes0315: Ring <ringno> issue: portCmdGet <local\_getidx> is bigger then cmd ring <max\_cmd\_idx>**

DESCRIPTION: The port cmd ring get index is greater than the size of cmd ring.

DATA: None

ACTION: This error could indicate a software driver, firmware, or hardware issue. Report the error to Broadcom technical support.

#### 4.1.4.9 **elx\_mes0317: iotag <ulp\_IoTag> is out of range: max iotag <max\_iotag> wd0 <wd0>**

DESCRIPTION: The IoTag in the completed IOCB is out of range.

DATA: None

ACTION: This error could indicate a software driver, firmware, or hardware issue. Report the error to Broadcom technical support.

#### 4.1.4.10 **elx\_mes0319: READ\_SPARAM mbxStatus error <mbxStatus> hba state <hba\_state>**

DESCRIPTION: The driver issued a READ\_SPARAM mbox command to the adapter that failed.

DATA: None

ACTION: This error could indicate a firmware or hardware issue. Report the error to Broadcom technical support.

#### 4.1.4.11 **elx\_mes0320: CLEAR\_LA mbxStatus error <mbxStatus> hba state <hba\_state>**

DESCRIPTION: The driver issued a CLEAR\_LA mbox command to the adapter that failed.

DATA: None

ACTION: This error could indicate a firmware or hardware issue. Report the error to Broadcom technical support.

#### 4.1.4.12 **elx\_mes0323: Unknown Mailbox command <mbxCommand> Cmpl**

DESCRIPTION: A unknown mailbox command completed.

DATA: None

ACTION: This error could indicate a software driver, firmware, or hardware issue. Report the error to Broadcom technical support.

---

**4.1.4.13 elx\_mes0324: Config port initialization error, mbxCmd <mbxCommand> READ\_NVPARM, mbxStatus <mbxStatus>**

DESCRIPTION: A read nparams mailbox command failed during port configuration.

DATA: None

ACTION: This error could indicate a software driver, firmware, or hardware issue. Report the error to Broadcom technical support.

**4.1.4.14 elx\_mes0330: IOCB wake NOT set**

DESCRIPTION: The completion handler associated with the IOCB was never called.

DATA:(1) timeout, (2) timeleft/jiffies

ACTION: This error could indicate a software driver, firmware, or hardware issue. If the issue persists, report the error to Broadcom technical support.

**4.1.4.15 elx\_mes0332: IOCB wait issue failed, Data <value>**

DESCRIPTION: A driver-issued I/O failed to complete in polling mode.

DATA: (1) error value.

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.4.16 elx\_mes0334: Unknown IOCB command**

DESCRIPTION: Received an unknown IOCB command completion.

DATA: (1) type, (2) ulpCommand, (3) ulpStatus, (4) ulploTag, (5) ulpContext

ACTION: This error could indicate a software driver or firmware issue. If this issue persists, report the error to Broadcom technical support.

**4.1.4.17 elx\_mes0335: Unknown IOCB command**

DESCRIPTION: Received an unknown IOCB command completion.

DATA: (1) ulpCommand, (2) ulpStatus, (3) ulploTag, (4) ulpContext

ACTION: This error could indicate a software driver or firmware issue. If this issue persists, report the error to Broadcom technical support.

**4.1.4.18 elx\_mes0338: IOCB wait timeout error - no wake response Data <value> <value>**

DESCRIPTION: Driver issued I/O did not get a wake signal in polling mode.

DATA: (1) wait time (2) wake value

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.4.19 elx\_mes0340: Adapter temperature is OK now**

DESCRIPTION: The adapter temperature has reverted to normal range.

DATA: Temperature in Celsius

ACTION: No action needed, informational.

---

**4.1.4.20 elx\_mes0341: Ring <ringno> Cannot find buffer for an unsolicited iocb tag <un.ulpWord[3]>**

DESCRIPTION: No additional pre-allocated buffers are available to handle unsolicited buffers.

DATA: None

ACTION: Verify that this port is not being managed by multiple ports.

**4.1.4.21 elx\_mes0342: Ring <ringno> Cannot find buffer for an unsolicited iocb tag <unsli3.sli3Words>**

DESCRIPTION: There was a multiple IOCB unsolicited command and sufficient buffer space cannot be allocated for it.

DATA: None

ACTION: None required.

**4.1.4.22 elx\_mes0343: Ring <ringno> Cannot find buffer for an unsolicited iocb tag <un.ulpWord[3]>**

DESCRIPTION: No additional pre-allocated buffers are available to handle unsolicited buffers.

DATA: None

ACTION: None required.

**4.1.4.23 elx\_mes0344: Ring <ringno> Cannot find buffer for an unsolicited iocb tag <unsli3.sli3Words[7]>**

DESCRIPTION: No additional pre-allocated buffers are available to handle unsolicited buffers.

DATA: None

ACTION: None required.

**4.1.4.24 elx\_mes0345: Resetting board due to mailbox timeout**

DESCRIPTION: A mailbox command failed to complete. The driver is resetting the port.

DATA: None

ACTION: If the mailbox command fails again, set the lpfc\_log\_verbose to LOG\_MBOX and retry.

**4.1.4.25 elx\_mes0346: Ring <ring number> handler: unexpected ASYNC\_STATUS evt\_code <evtcode>**

DESCRIPTION: The adapter received an asynchronous event that was not a temperature event.

DATA: None

ACTION: None required.

**4.1.4.26 elx\_mes0347: Adapter is very hot, please take corrective action. Temperature: <value> Celsius**

DESCRIPTION: The adapter temperature is above normal range

DATA: Temperature in Celsius

ACTION: Shutdown and remove the adapter. Contact customer support.

**4.1.4.27 elx\_mes0348: NameServer login: node freed**

DESCRIPTION: The enable mode failed to free up the nameserver login.

DATA: None

ACTION: None required.

---

**4.1.4.28 elx\_mes0349: rc should be MBX\_SUCCESS**

DESCRIPTION: The next mailbox command on the mailbox queue has failed.

DATA: None

ACTION: None required.

**4.1.4.29 elx\_mes0350: rc should have been MBX\_BUSY**

DESCRIPTION: Attempting to unregister a default RPI from an interrupt context and the mailbox state is not busy.

DATA: None

ACTION: None required.

**4.1.4.30 elx\_mes0352: Config MSI mailbox command failed, mbxCmd <u.mb.mbxCommand>, mbxStatus <u.mb.mbxStatus>**

DESCRIPTION: The mailbox command sent to the firmware to configure the adapter to use MSI-X has failed.

DATA: None

ACTION: Ensure the hardware platform supports MSI-X.

**4.1.4.31 elx\_mes0359: Not a valid slow-path completion event: majorcode=<value>, minorcode=<value>**

DESCRIPTION: SLI-4: The EQE is not valid.

DATA: None

ACTION: None required.

**4.1.4.32 elx\_mes0360: Unsupported EQ count. <entry\_count>**

DESCRIPTION: The firmware cannot create an event queue of this size.

DATA: None

ACTION: None required.

**4.1.4.33 elx\_mes0361: Unsupported CQ count. <entry\_count>**

DESCRIPTION: The firmware cannot create a completion queue of this size.

DATA: None

ACTION: None required.

**4.1.4.34 elx\_mes0362: Unsupported MQ count. <entry\_count>**

DESCRIPTION: The firmware cannot create MQ count of this size.

DATA: None

ACTION: None required.

**4.1.4.35 elx\_mes0364: Invalid param**

DESCRIPTION: SLI-4: The post SGL function was passed an invalid XRI.

DATA: None

ACTION: None required.

---

**4.1.4.36 elx\_mes0365: Slow-path CQ identifier <cqid> does not exist**

DESCRIPTION: The Completion Queue ID passed in the Event Queue entry does not reference a valid completion queue.

DATA: None

ACTION: None required.

**4.1.4.37 elx\_mes0366: Not a valid fast-path completion event: majorcode=<major code hex>, minorcode=<minor code hex>**

DESCRIPTION: The major or minor code in the Event Queue field is not valid.

DATA: None

ACTION: None required.

**4.1.4.38 elx\_mes0367: Fast-path completion queue does not exist**

DESCRIPTION: The fast path completion queue referenced by the CQID does not exist.

DATA: None

ACTION: None required.

**4.1.4.39 elx\_mes0368: Miss-matched fast-path completion queue identifier: eqcqid=<cqid>, fcpcqid=<queue\_id>**

DESCRIPTION: The CQID in the event queue entry does not match the fcpcqid that was passed into the routine.

DATA: None

ACTION: None required.

**4.1.4.40 elx\_mes0369: No entry from fast-path completion queue fcpcqid=<queue\_id>**

DESCRIPTION: No completions exist in the completion queue referenced by fcpcqid.

DATA: None

ACTION: None required.

**4.1.4.41 elx\_mes0370: Invalid completion queue type <type>**

DESCRIPTION: The event queue entry is not for a mailbox or a work queue entry.

DATA: None

ACTION: None required.

**4.1.4.42 elx\_mes0371: No entry from the CQ: identifier <queue\_id>, type <type>**

DESCRIPTION: No completion queue event exists for this event queue entry.

DATA: None

ACTION: None required.

**4.1.4.43 elx\_mes0372: iotag <iotag> is out of range: max iotag (<sli.last\_iotag>)**

DESCRIPTION: The IOCB lookup cannot be performed because the iocb\_tag is out of range.

DATA: None

ACTION: None required.

---

**4.1.4.44 elx\_mes0376: READ\_REV Error. SLI Level <sli\_rev> FCoE enabled <hba\_flag & HBA\_FCOE\_SUPPORT>**

DESCRIPTION: This SLI-4 only adapter setup function was called for a non-SLI-4 device.

DATA: None

ACTION: None required.

**4.1.4.45 elx\_mes0377: Error <rc> parsing vpd. Using defaults.**

DESCRIPTION: Could not parse the VPD data, so the driver is using the default values.

DATA: None

ACTION: None required.

**4.1.4.46 elx\_mes0381: Error <rc> during queue setup.**

DESCRIPTION: Could not set up all the queues that the driver requires to exchange I/Os with the adapter.

DATA: None

ACTION: Reload the driver.

**4.1.4.47 elx\_mes0382: READ\_SPARAM command failed status <issue status>, mbxStatus <mailbox status>**

DESCRIPTION: The READ\_SPARAM mailbox command has failed during initialization. The adapter has been set to error state.

DATA: None

ACTION: Perform a dump with hbacmd and then try reloading the driver.

**4.1.4.48 elx\_mes0383: Error <error> during scsi sgl post operation**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.4.49 elx\_mes0384: There is pending active mailbox cmd**

DESCRIPTION: The mailbox commands have overlapped. This command should have been added to the mailbox queue.

DATA: None

ACTION: None required.

**4.1.4.50 elx\_mes0385: rc should have been MBX\_BUSY**

DESCRIPTION: 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

ACTION: None required.

**4.1.4.51 elx\_mes0387: Failed to allocate an iocbq**

DESCRIPTION: Failed to get an IOCBQ from the list of available IOCBQs.

DATA: None

ACTION: None required.

---

**4.1.4.52 elx\_mes0388: Not a valid WCQE code: <hex cqe\_code>**

DESCRIPTION: The event code is invalid. This event is dropped.

DATA: None

ACTION: Verify that the adapter's firmware is current.

**4.1.4.53 elx\_mes0391: Error during rpi post operation**

DESCRIPTION: The driver was trying to post pages to the firmware to keep target login information and encountered a failure.

DATA: None

ACTION: Unload and reload the driver.

**4.1.4.54 elx\_mes0393: Error <rc> during rpi post operation**

DESCRIPTION: The driver was trying to post pages to the firmware to keep target login information and encountered a failure.

DATA: None

ACTION: Unload and reload the driver.

**4.1.4.55 elx\_mes0394: Failed to allocate CQ\_EVENT entry**

DESCRIPTION: The asynchronous event handler was unable to allocate an event queue entry to which to transfer the asynchronous event.

DATA: None

ACTION: This could be a V-LINK clear from the switch or a fatal error from the firmware. Perform a dump from the OneCommand Manager application.

**4.1.4.56 elx\_mes0395: The mboxq allocation failed**

DESCRIPTION: The asynchronous link event handler could not allocate a mailbox command to issue the READ\_LA (read link attention) mailbox command.

DATA: None

ACTION: None required.

**4.1.4.57 elx\_mes0396: The lpfc\_dmabuf allocation failed**

DESCRIPTION: The asynchronous link event handler could not allocate a DMA buffer for the mailbox command to issue the READ\_LA (read link attention) mailbox command.

DATA: None

ACTION: None required.

**4.1.4.58 elx\_mes0397: The mbuf allocation failed**

DESCRIPTION: The asynchronous link event handler could not allocate DMA-able memory for the READ\_LA mailbox command.

DATA: None

ACTION: None required.

---

**4.1.4.59 elx\_mes0398: Invalid link fault code: < hex link\_fault >**

DESCRIPTION: The attempt to read the link attention register has returned an unknown value.

DATA: None

ACTION: None required.

**4.1.4.60 elx\_mes0399: Invalid link attention type: <hex link\_type >**

DESCRIPTION: The READ\_LA mailbox command has returned an invalid link type.

DATA: None

ACTION: None required.

## 4.1.5 Initialization Events (0400 to 0599)

**4.1.5.1 elx\_mes0400: Phys Attribute Count Exceeded, Max <value>, Actual <value>**

DESCRIPTION: Too many driver configuration parameters have been set. The limit is given as Max.

DATA: (1) Maximum number (2) Actual number

ACTION: Reduce the number of actual parameters.

**4.1.5.2 elx\_mes0402: Cannot find virtual addr for buffer tag on ring <ringno>**

DESCRIPTION: A DMA buffer is not available for this unsolicited command.

DATA: (1) tag, (2) next, (3) prev, (4) postbufq\_cnt

ACTION: None required.

**4.1.5.3 elx\_mes0403: lpfc\_nodev\_tmo attribute cannot be set to <val>, allowed range is [<>LPFC\_MIN\_DEVLOSS\_TMO>, <LPFC\_MAX\_DEVLOSS\_TMO>]**

DESCRIPTION: Attempt to set the nodev timeout value is outside the range of the devloss timeout range.

DATA: None

ACTION: Set the nodev timeout between the minimum and maximum of the devloss timeout range.

**4.1.5.4 elx\_mes0404: Config Param <value> set to <value>**

DESCRIPTION: The driver is setting a persistent vPort parameter to a different value.

DATA: (1) New value

ACTION: None. This message is notification only.

**4.1.5.5 elx\_mes0405: Config Param <value> set to <value>**

DESCRIPTION: The driver is setting a persistent vPort parameter to a different value.

DATA: (1) New value

ACTION: None. This message is notification only.

**4.1.5.6 elx\_mes0406: Adapter maximum temperature exceeded (<temperature>), taking this port offline**

DESCRIPTION: The driver has received an error from the adapter indicating that the maximum allowable temperature has been exceeded.

DATA: (1) work\_hs, (2) work\_status[0], (3) work\_status[1]

ACTION: Make sure that the server fans are not blocked. Shut down the server if the airflow is restricted.

---

**4.1.5.7 elx\_mes0408: Cannot create debugfs root**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.5.8 elx\_mes0409: Cannot create debugfs nodelist**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.5.9 elx\_mes0410: Cannot find virtual addr for mapped buf on ring <ringno>**

DESCRIPTION: The driver cannot find the specified buffer in its mapping table. Thus it cannot find the virtual address needed to access the data.

DATA: (1) phys, (2) next, (3) prev, (4) postbufq\_cnt

ACTION: This error could indicate a software driver or firmware issue. If the issue persists report the error to Broadcom technical support.

**4.1.5.10 elx\_mes0411: Cannot create debugfs hbqinfo**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.5.11 elx\_mes0412: Cannot create debugfs hba**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.5.12 elx\_mes0413: Cannot create debugfs dumpHBASlim**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.5.13 elx\_mes0414: Cannot create debugfs dumpHostSlim**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.5.14 elx\_mes0415: Cannot create debugfs slow\_ring trace**

DESCRIPTION:

DATA: None

ACTION: None required.

---

**4.1.5.15 elx\_mes0416: Cannot create debugfs slow\_ring buffer**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.5.16 elx\_mes0417: Cannot create debugfs**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.5.17 elx\_mes0418: Cannot create debugfs disc trace buffer**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.5.18 elx\_mes0419: Cannot create debugfs discovery trace**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.5.19 elx\_mes0423: Vport Attribute Instance Error. Defaulting lpfc\_#attr to <value>, error value <value>, allowed range is [min, max]**

DESCRIPTION: A vPort attribute was set out of range. The driver reset the parameter to its default.

DATA: None

ACTION: Set the module parameter between the minimum and maximum values.

**4.1.5.20 elx\_mes0424: Vport Attribute Count Exceeded, Max <value>, Actual <value>**

DESCRIPTION: The total number of vPort attributes set exceeded the max allowed.

DATA: None

ACTION: Reduce the number of set attributes to below the maximum.

**4.1.5.21 elx\_mes0425: lpfc\_#attr" attribute cannot be set to <value>, allowed range is [min, max]**

DESCRIPTION: Driver attribute lpfc\_#attr was defined with an out-of-range value.

DATA: None

ACTION: Set the parameter between the minimum and maximum values.

**4.1.5.22 elx\_mes0426: lpfc\_enable\_auth attribute cannot be set to <value>, allowed range is [<min>, <max>]**

DESCRIPTION:

DATA: None

ACTION: None required.

---

**4.1.5.23 elx\_mes0427: Cannot re-enable interrupt after slot reset.**

DESCRIPTION: The driver was not able to enable the interrupt after an adapter reset.

DATA: None

ACTION: Unload and reload the driver.

**4.1.5.24 elx\_mes0430: PM resume Failed to enable interrupt**

DESCRIPTION: The driver's power management resume function could not enable the interrupt.

DATA: None

ACTION: Perform another PM suspend and resume or adapter reset.

**4.1.5.25 elx\_mes0431: Failed to enable interrupt.**

DESCRIPTION: The driver failed to start the interrupt.

DATA: None

ACTION: Unload and reload the driver.

**4.1.5.26 elx\_mes0433: Wakeup on signal: rc=<rc>**

DESCRIPTION: A signal other than the LPFC\_DATA\_READY was received on the worker thread.

DATA: None

ACTION: Unload and reload the driver.

**4.1.5.27 elx\_mes0434: PM resume failed to start worker thread: error=<error>.**

DESCRIPTION: The driver's power management resume function could not start the worker thread.

DATA: None

ACTION: Unload and reload the driver.

**4.1.5.28 elx\_mes0435: Adapter failed to get Option ROM version status <rc>.**

DESCRIPTION: The driver could not read the adapter's option ROM.

DATA: None

ACTION: Reset the adapter. Make sure that the adapter's firmware is current.

**4.1.5.29 elx\_mes0436: Adapter failed to init, timeout, status reg <status>**

DESCRIPTION: The adapter failed during power-up diagnostics after it was reset.

DATA: None

ACTION: This error could indicate a hardware or firmware issue. If the issue persists, report the error to Broadcom technical support.

**4.1.5.30 elx\_mes0437: Adapter failed to init, chipset, status reg <status>**

DESCRIPTION: The adapter failed during power-up diagnostics after it was reset.

DATA: None

ACTION: This error could indicate a hardware or firmware issue. If the issue persists, report the error to Broadcom technical support.

---

**4.1.5.31 elx\_mes0438: Adapter failed to init, chipset, status reg <status>**

DESCRIPTION: The adapter failed during power-up diagnostics after it was reset.

DATA: None

ACTION: This error could indicate a hardware or firmware issue. If the issue persists, report the error to Broadcom technical support.

**4.1.5.32 elx\_mes0439: Adapter failed to init, mbxCmd <mbxCommand> READ\_REV, mbxStatus <mbxStatus>**

DESCRIPTION: Adapter initialization failed when issuing a READ\_REV mailbox command.

DATA: None

ACTION: This error could indicate a hardware or firmware issue. If the issue persists, report the error to Broadcom technical support.

**4.1.5.33 elx\_mes0440: Adapter failed to init, READ\_REV has missing revision information**

DESCRIPTION: A firmware revision initialization error was detected.

DATA: None

ACTION: This error could indicate a hardware or firmware issue. Update the firmware. If the issue persists, report the error to Broadcom technical support.

**4.1.5.34 elx\_mes0442: Adapter failed to init, mbxCmd <mbxCommand> CONFIG\_PORT, mbxStatus <mbxStatus>**

DESCRIPTION: Adapter initialization failed when issuing a CONFIG\_PORT mailbox command.

DATA: (1) hbainit

ACTION: This error could indicate a hardware or firmware issue. If the issue persists, report the error to Broadcom technical support.

**4.1.5.35 elx\_mes0443: Adapter failed to set maximum DMA length mbxStatus <u.mb.mbxStatus>.**

DESCRIPTION: Cannot set the maximum DMA length to reflect cfg\_pci\_max\_read.

DATA: None

ACTION: Set module parameter lpfc\_pci\_max\_read to 512, 1024, 2048, or 4096.

**4.1.5.36 elx\_mes0445: Firmware initialization failed.**

DESCRIPTION: The driver was unable to initialize the hardware.

DATA: None

ACTION: This error could indicate a hardware or firmware issue. If the issue persists, report the error to Broadcom technical support.

**4.1.5.37 elx\_mes0446: Adapter failed to init, mbxCmd <mbxCommand> CFG\_RING, mbxStatus <mbxStatus>, ring <num>**

DESCRIPTION: Adapter initialization failed when issuing a CFG\_RING mailbox command.

DATA: None

ACTION: This error could indicate a hardware or firmware issue. If the issue persists, report the error to Broadcom technical support.

---

**4.1.5.38 elx\_mes0448: Adapter failed to init, mbxCmd <mbxCommand> READ\_SPARM, mbxStatus <mbxStatus>**

DESCRIPTION: Adapter initialization failed when issuing a READ\_SPARM mailbox command.

DATA: None

ACTION: This error could indicate a hardware or firmware issue. If the issue persists, report the error to Broadcom technical support.

**4.1.5.39 elx\_mes0449: Phys attribute Instance Error. Defaulting to lpfc\_#attr to <value>. Allowed range is [min, max]**

DESCRIPTION: A physical device attribute has an out-of-range value. The driver is correcting it.

DATA: (1) value written, (2) minimum value, (3) maximum value

ACTION: Write the default value.

**4.1.5.40 elx\_mes0450: lpfc\_%attr attribute cannot be set to <value>, allowed range is [%min, %max]**

DESCRIPTION: Sysfs attribute value written exceeds attribute range.

DATA: (1) attribute name, (2) value written, (3) minimum value, (3) maximum value

ACTION: Write a value within the supported range.

**4.1.5.41 elx\_mes0451: Failed to enable interrupt**

DESCRIPTION:

DATA: None.

ACTION: None required.

**4.1.5.42 elx\_mes0453: Adapter failed to init, mbxCmd <mbxCommand> READ\_CONFIG, mbxStatus <mbxStatus>**

DESCRIPTION: Adapter initialization failed when issuing a READ\_CONFIG mailbox command.

DATA: None

ACTION: This error could indicate a hardware or firmware issue. If the issue persists, report the error to Broadcom technical support.

**4.1.5.43 elx\_mes0456: Adapter failed to issue ASYNCEVT\_ENABLE mbox status <rc>.**

DESCRIPTION: The mailbox command to enable an asynchronous event notification failed.

DATA: None

ACTION: Make sure the adapter firmware is current. Reload the driver.

**4.1.5.44 elx\_mes0457: Adapter Hardware Error**

DESCRIPTION: The driver received an interrupt indicating a possible hardware issue.

DATA: (1) status, (2) status1, (3) status2

ACTION: This error could indicate a hardware or firmware issue. If the issue persists, report the error to Broadcom technical support.

**4.1.5.45 elx\_mes0462: Too many cmd / rsp ring entries in SLI2 SLIM Data: <values> /home/pely/svn\_linux\_lpfc\_upstream/branches/vmware-8.2.1-ucna/kernel/lpfc/lpfc\_sli.c:**

DESCRIPTION:

DATA: None

ACTION: None required.

---

**4.1.5.46 elx\_mes0472: Unknown PCI error state: <value>**

DESCRIPTION: The PCI bus has detected an error.

DATA: (1) state value

ACTION: The driver resets the adapter and attempts recovery. If the issue persists, contact Broadcom technical support.

**4.1.5.47 elx\_mes0474: Unable to allocate memory for issuing "MBOX\_CONFIG\_MSI command"**

DESCRIPTION: Mailbox memory pool allocation error.

DATA: None

ACTION: None required.

**4.1.5.48 elx\_mes0475: Not configured for supporting MSI-X cfg\_use\_msi: <cfg\_use\_msi>.**

DESCRIPTION: The lpfc\_use\_msi module parameter should have been set to 2.

DATA: None

ACTION: Set module parameter lpfc\_use\_msi=2.

**4.1.5.49 elx\_mes0476: HBA not supporting SLI-3 or later SLI Revision: <sli\_rev>.**

DESCRIPTION: The adapter does not support SLI-3 or SLI-4.

DATA: None

ACTION: This adapter does not support MSI. Set lpfc\_use\_msi=0.

**4.1.5.50 elx\_mes0479: Deferred Adapter Hardware Error**

DESCRIPTION: An adapter hardware error was sent to the driver.

DATA: (1) work\_hs, (2) work\_status[0], (3) work\_status[1]

ACTION: Perform a dump using hbacmd.

**4.1.5.51 elx\_mes0482: Illegal interrupt mode**

DESCRIPTION: The driver could not set MSI-X, MSI or INTx interrupt modes.

DATA: None

ACTION: This could be a server issue. Reboot. If this issue persists, report the error to Broadcom technical support.

**4.1.5.52 elx\_mes0483: Invalid link-attention link speed: <value>, bf\_get(lpfc\_acqe\_link\_speed, acqe\_link).**

DESCRIPTION: The link speed reported in the link attention interrupt is invalid.

DATA: None

ACTION: Check the switch configuration.

**4.1.5.53 elx\_mes0492: Unable to allocate memory for issuing SLI\_CONFIG\_SPECIAL mailbox command**

DESCRIPTION: A memory allocation fault occurred when issuing a mailbox.

DATA: None

ACTION: This could be a transient error. If this issue persists, report the error to Broadcom technical support.

---

**4.1.5.54 elx\_mes0493: SLI\_CONFIG\_SPECIAL mailbox failed with status <rc>.**

DESCRIPTION: The mailbox command failed.

DATA: None

ACTION: Make sure the adapter's firmware is current. Unload and reload the driver.

**4.1.5.55 elx\_mes0494: Unable to allocate memory for issuing "SLI\_FUNCTION\_RESET mailbox command"**

DESCRIPTION: Mailbox memory pool allocation error.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

**4.1.5.56 elx\_mes0495: SLI\_FUNCTION\_RESET mailbox failed with status <shdr\_status> add\_status <shdr\_add\_status>, mbx status <rc>.**

DESCRIPTION: Mailbox command failed.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

**4.1.5.57 elx\_mes0496: Failed allocate slow-path EQ**

DESCRIPTION: The event queue for the slow path was not allocated.

DATA: None

ACTION: Unload and reload the driver.

**4.1.5.58 elx\_mes0497: Failed allocate fast-path EQ**

DESCRIPTION: The event queue for the fast path was not allocated.

DATA: None

ACTION: Unload and reload the driver.

**4.1.5.59 elx\_mes0498: Adapter failed to init, mbxCmd <cmd> INIT\_LINK, mbxStatus <status>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.5.60 elx\_mes0499: Failed allocate fast-path FCP CQ (<fcq\_cqidx>).**

DESCRIPTION: The completion queue event for the fast path could not be allocated.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

**4.1.5.61 elx\_mes0500: Failed allocate slow-path mailbox CQ**

DESCRIPTION: The driver failed to allocate the slow-path mailbox completion queue.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

---

**4.1.5.62 elx\_mes0501: Failed allocate slow-path ELS CQ**

DESCRIPTION: The driver failed to allocate the slow-path extended link service completion queue.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

**4.1.5.63 elx\_mes0503: Failed allocate fast-path FCP**

DESCRIPTION: The driver failed to allocate the fast-path FCP.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

**4.1.5.64 elx\_mes0504: Failed allocate slow-path ELS WQ**

DESCRIPTION: The driver failed to allocate the slow-path extended link service WQ.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

**4.1.5.65 elx\_mes0505: Failed allocate slow-path MQ**

DESCRIPTION: The driver failed to allocate the slow-path MQ.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

**4.1.5.66 elx\_mes0506: Failed allocate receive HRQ**

DESCRIPTION: The driver failed to allocate the receive HRQ.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

**4.1.5.67 elx\_mes0507: Failed allocate receive DRQ**

DESCRIPTION: The driver failed to allocate the receive DRQ.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

**4.1.5.68 elx\_mes0520: Slow-path EQ not allocated**

DESCRIPTION: The slow-path event queue is not allocated.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

**4.1.5.69 elx\_mes0521: Failed setup of slow-path EQ rc = <value>**

DESCRIPTION: The slow-path event queue setup failed with status rc.

DATA: (1) status code

ACTION: The driver fails to load. Contact Broadcom technical support.

---

**4.1.5.70 elx\_mes0522: Fast-path EQ <fcp\_eqidx> not allocated**

DESCRIPTION: The fast-path event queue is not allocated.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

**4.1.5.71 elx\_mes0523: Failed setup of fast-path EQ <fcp\_eqidx>, rc = <rc>**

DESCRIPTION: The fast-path event queue setup failed.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

**4.1.5.72 elx\_mes0526: Fast-path FCP CQ <fcp\_cqidx> not allocated**

DESCRIPTION: The fast-path FCP is not allocated.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

**4.1.5.73 elx\_mes0527: Failed setup of fast-path FCP CQ <fcp\_cqidx>, rc = <rc>**

DESCRIPTION: The fast-path FCP completion queue setup failed.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

**4.1.5.74 elx\_mes0528: Mailbox CQ not allocated**

DESCRIPTION: The mailbox completion queue is not allocated.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

**4.1.5.75 elx\_mes0529: Failed setup of slow-path mailbox CQ: rc = <value>**

DESCRIPTION: The driver failed to setup the Completion Queue. Failure code is reported.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.5.76 elx\_mes0530: ELS CQ not allocated**

DESCRIPTION: The extended link service completion queue is not allocated.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

**4.1.5.77 elx\_mes0531: Failed setup of slow-path ELS CQ: rc = <value>**

DESCRIPTION: The extended link service completion queue is allocated, but failed initial setup.

DATA: (1) status

ACTION: The driver fails to load. Contact Broadcom technical support.

---

**4.1.5.78 elx\_mes0534: Fast-path FCP WQ <fcp\_eqidx> not allocated**

DESCRIPTION: The fast-path FCP WQ is not allocated.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

**4.1.5.79 elx\_mes0535: Failed setup of fast-path FCP WQ <fcp\_wqidx>, rc = <rc>**

DESCRIPTION: The fast-path FCP WQ setup failed.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

**4.1.5.80 elx\_mes0536: Slow-path ELS WQ not allocated**

DESCRIPTION: The slow-path extended link service WQ is not allocated.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

**4.1.5.81 elx\_mes0537: Failed setup of slow-path ELS WQ: rc = <value>**

DESCRIPTION: The driver failed to setup the Work Queue. A failure code is reported.

DATA: (1) (2) (3)

ACTION: Software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.5.82 elx\_mes0538: Slow-path MQ not allocated**

DESCRIPTION: The slow-path MQ is not allocated.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

**4.1.5.83 elx\_mes0539: Failed setup of slow-path MQ: rc = <value>**

DESCRIPTION: The slow-path MQ is allocated, but failed initial setup.

DATA: (1) status

ACTION: The driver fails to load. Contact Broadcom technical support.

**4.1.5.84 elx\_mes0540: Receive Queue not allocated**

DESCRIPTION: The Receive Queue is not allocated.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

**4.1.5.85 elx\_mes0541: Failed setup of Receive Queue: rc = <value>**

DESCRIPTION: The Receive Queue is allocated, but failed setup.

DATA: (1) status

ACTION: The driver fails to load. Contact Broadcom technical support.

---

**4.1.5.86 elx\_mes0542: lpfc\_create\_static\_vport failed to allocate mailbox memory**

DESCRIPTION: The driver failed to allocate the mailbox memory for vPort creation.

DATA: None

ACTION: Static VPorts does not load. Contact Broadcom technical support.

**4.1.5.87 elx\_mes0543: lpfc\_create\_static\_vport failed to allocate vport\_info**

DESCRIPTION: The driver failed to allocate the VPort\_info.

DATA: None

ACTION: Static VPorts does not load. Contact Broadcom technical support.

**4.1.5.88 elx\_mes0545: lpfc\_create\_static\_vport bad information header <value> <value>, le32\_to\_cpu(vport\_info->signature), le32\_to\_cpu(vport\_info->rev) & VPORT\_INFO\_REV\_MASK;**

DESCRIPTION: Invalid information header; the signature or revision is invalid.

DATA: None

ACTION: Static VPorts does not load. Contact Broadcom technical support.

**4.1.5.89 elx\_mes0582: Error <rc> during sgl post operation**

DESCRIPTION: The SGL post operation failed.

DATA: None

ACTION: None required.

**4.1.5.90 elx\_mes0602: Failed to allocate CQ\_EVENT entry**

DESCRIPTION: The driver failed to allocate a CQ\_EVENT entry.

DATA: None

ACTION: None required.

**4.1.5.91 elx\_mes0603: Invalid work queue CQE subtype <subtype>**

DESCRIPTION: Invalid work queue CQE.

DATA: None

ACTION: None required.

## 4.1.6 FCP Traffic History (0700 to 0799)

**4.1.6.1 elx\_mes0700: Bus Reset on target <i> failed**

DESCRIPTION: The bus reset for the specified target failed.

DATA: None

ACTION: None required.

**4.1.6.2 elx\_mes0706: IOCB Abort failed - outstanding <value> failed <value>**

DESCRIPTION: The driver did not recover all I/O following a reset task management command.

DATA: (1) outstanding I/O count (2) number of unrecovered I/O

ACTION: The reset call fails. ESXi tries to recover.

---

**4.1.6.3 elx\_mes0713: SCSI layer issued Device Reset (<value>, <value>) reset status <value> flush status <value>**

DESCRIPTION: A device reset has completed on (tgt, lun). Status values are displayed.

DATA: (1) tgt (2) lun (3) task mgmt status (4) flush status

ACTION: None required.

**4.1.6.4 elx\_mes0714: SCSI layer issued bus reset**

DESCRIPTION: The SCSI layer is requesting the driver to abort all I/Os to all targets on this adapter.

DATA: (1) ret

ACTION: Check the state of the targets in question.

**4.1.6.5 elx\_mes0717: FCP command <value> residual underrun converted to error Data: <values>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.6.6 elx\_mes0718: Unable to dma\_map single request\_buffer: <value>**

DESCRIPTION: The driver could not map a single virtual address to a DMA address.

DATA: (1) DMA mapping error

ACTION: None. The driver fails the I/O back to ESXi.

**4.1.6.7 elx\_mes0720: FCP command <value> residual overrun error. Data: <values>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.6.8 elx\_mes0721: Device Reset rport failure: rdata <rdata>**

DESCRIPTION: The reset of the Rport failed.

DATA: None

ACTION: None required.

**4.1.6.9 elx\_mes0724: I/O flush failure for context <ctx> on <tgt:lun> cnt <value>**

DESCRIPTION: The I/O flush to the {LUN, TARGET, or HOST} has failed.

DATA: (1) count of unrecovered I/O

ACTION: None required. The reset is retried.

**4.1.6.10 elx\_mes0727: TMF <cmd> to TGT <TGT#> LUN <LUN#> failed (<ulpStatus>, <ulpWord[4]>)**

DESCRIPTION: The task management function command failed.

DATA: None

ACTION: None required. The task management function command gets retried.

---

**4.1.6.11 elx\_mes0748: Abort handler timed out waiting for abort to complete:ret <status> ID <target id> LUN <lun id> snum <serial number>**

DESCRIPTION: The abort handler timed out waiting for abort to complete.

DATA: None

ACTION: None required.

**4.1.6.12 elx\_mes0798: Device Reset rport failure: rdata <value>**

DESCRIPTION: Driver failed a device reset - no rdata buffer.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

## **4.1.7 Node Table Events (0900 to 0999)**

**4.1.7.1 elx\_mes0915: Register VPI failed: <mbxStatus>**

DESCRIPTION: Could not register the VPI.

DATA: None

ACTION: None required.

## **4.1.8 Security Events (1000 to 1099)**

**4.1.8.1 elx\_mes1000: Authentication is enabled but authentication service is not running**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.8.2 elx\_mes1005: AUTHENTICATION\_FAILURE Nport: <port>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.8.3 elx\_mes1006: Bad Name tag in auth message**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.8.4 elx\_mes1007: Bad Name length in auth message**

DESCRIPTION:

DATA: None

ACTION: None required.

---

**4.1.8.5 elx\_mes1008: Bad Number of Protocols**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.8.6 elx\_mes1009: Bad param type**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.8.7 elx\_mes1010: Bad Tag 1**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.8.8 elx\_mes1011: Auth\_neg no has function chosen.**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.8.9 elx\_mes1012: Auth\_negotiate Bad Tag 2**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.8.10 elx\_mes1013: Auth\_negotiate no DH\_group found.**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.8.11 elx\_mes1014: dhchap challenge bad name tag.**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.8.12 elx\_mes1015: dhchap challenge bad name length.**

DESCRIPTION:

DATA: None

ACTION: None required.

---

**4.1.8.13 elx\_mes1016: dhchap challenge Hash ID not Supported.**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.8.14 elx\_mes1017: dhchap challenge could not find DH Group.**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.8.15 elx\_mes1018: dhchap challenge No Public key for non-NULL DH Group.**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.8.16 elx\_mes1021: ERROR: attempted to queue security work, when no workqueue created.**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.8.17 elx\_mes1028: Start Authentication: No buffers**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.8.18 elx\_mes1029: Reauthentication Failure**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.8.19 elx\_mes1031: Start Authentication: Get config failed.**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.8.20 elx\_mes1032: Start Authentication: get config timed out.**

DESCRIPTION:

DATA: None

ACTION: None required.

---

**4.1.8.21 elx\_mes1034: Not Expecting Challenge - Rejecting Challenge.**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.8.22 elx\_mes1035: Transport ID does not math - Rejecting Challenge.**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.8.23 elx\_mes1036: Authentication transaction reject - re-auth request reason <value> exp <value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.8.24 elx\_mes1037: Authentication transaction reject - restarting authentication. reason <value> exp <value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.8.25 elx\_mes1039: Not Expecting Reply - rejecting. State <value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.8.26 elx\_mes1040: Bad Reply trans\_id - rejecting. Trans\_id: <value> Expecting: <value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.8.27 elx\_mes1043: Authentication LS-RJT**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.8.28 elx\_mes1045: Issue AUTH\_NEG failed. Status: <value>**

DESCRIPTION:

DATA: None

ACTION: None required.

---

**4.1.8.29 elx\_mes1048: Issue AUTH\_REJECT failed.**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.8.30 elx\_mes1049: Authentication is enabled but authentication service is not running**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.8.31 elx\_mes1050: Authentication mode is disabled, but is required by the fabric.**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.8.32 elx\_mes1053: Start Authentication: Security service offline.**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.8.33 elx\_mes1055: Authentication parameter is disabled, but is required by the fabric.**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.8.34 elx\_mes1056: Authentication mode is disabled, but is required by the fabric.**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.8.35 elx\_mes1057: Authentication transaction reject. reason <value> exp <value>**

DESCRIPTION:

DATA: None

ACTION: None required.

## **4.1.9 Miscellaneous and FCoE Events (1200 to 1299)**

**4.1.9.1 elx\_mes1201: Failed to allocate dfc\_host**

DESCRIPTION: The driver failed to allocate a DFC host and bind it to the management stack.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

---

**4.1.9.2 elx\_mes1209: C\_CT Request error Data: <value> <value>**

DESCRIPTION: IOCTL common transport response error - the driver is failing the IOCTL request.

DATA: (1) response buffer flag (2) Data Size

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.9.3 elx\_mes1210: Invalid cmd size: <cmd value> <cmdsz value> <rspsz value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.9.4 elx\_mes1211: genreq alloc failed: <value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.9.5 elx\_mes1213: FCoE cmd overflow: <off value> + <cnt value> > <cmdsz value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.9.6 elx\_mes1214: Cannot issue FCoE cmd, SLI not active: <off value> rc = -EACCESS**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.9.7 elx\_mes1215: Cannot issue FCoE cmd: not ready or not in maint mode <off value> <rc value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.9.8 elx\_mes1216: FCoE IOCB failed: <off value> <rc value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.9.9 elx\_mes1223: menlo\_write: couldn't alloc genreq <value>**

DESCRIPTION:

DATA: None

ACTION: None required.

---

**4.1.9.10 elx\_mes1230: Could not find buffer for FCoE cmd: <off value> <indmp value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.9.11 elx\_mes1231: bad bpl**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.9.12 elx\_mes1235: Could not find buffer for FCoE cmd: <off value> poff: <value> cnt: <value> mlastcnt: <value> addl: <value> addh: <value> mdsz: <value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.9.13 elx\_mes1238: FCoE IOCB failed: <off value> <rc value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.9.14 elx\_mes1240: Unable to allocate command buffer memory.**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.9.15 elx\_mes1243: Menlo command error. code=<value>.**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.9.16 elx\_mes1244: Unable to allocate response buffer memory.**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.9.17 elx\_mes1246: FCoE chip is running golden firmware. Update FCoE chip firmware immediately**

DESCRIPTION:

DATA: None

ACTION: None required.

---

**4.1.9.18 elx\_mes1247: FCoE chip is running diagnostic firmware. Operational use suspended.**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.9.19 elx\_mes1248: FCoE chip is running unknown firmware.**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.9.20 elx\_mes1249: Invalid FRU data found on adapter. Return adapter to Emulex for repair**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.9.21 elx\_mes1250: Menlo command error. code=<value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.9.22 elx\_mes1251: Menlo command error. code=<value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.9.23 elx\_mes1252: Menlo command error. code=<value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.9.24 elx\_mes1262: Failed to allocate dfc\_host**

DESCRIPTION: The driver could not allocate memory to the dfc\_host\_struct.

DATA: None

ACTION: None required.

**4.1.10 Link Events (1300 to 1399)**

**4.1.10.1 elx\_mes1300: Link Down Event in loop back mode**

DESCRIPTION: The driver received a link down event while in loopback mode - unexpected event.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

---

#### 4.1.10.2 **elx\_mes1302: Invalid speed for this board: Reset link speed to auto: <value>**

DESCRIPTION: The driver detected an invalid link speed. Resetting the Link to Auto mode.

DATA: (1) Invalid speed detected

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

#### 4.1.10.3 **elx\_mes1303: Link Up Event <eventTag> received Data: <value> <value> <value> <value> <value> <value> <value>**

DESCRIPTION: 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]

ACTION: If numerous link events are occurring, check the physical connections to the Fibre Channel network.

#### 4.1.10.4 **elx\_mes1305: Link Down Event <eventTag> received Data: <value> <value> <value>**

DESCRIPTION: A link down event was received.

DATA: (1) fc\_eventTag, (2) hba\_state, (3) fc\_flag

ACTION: If numerous link events are occurring, check the physical connections to the Fibre Channel network.

#### 4.1.10.5 **elx\_mes1306: Link Up Event in loop back mode <eventTag> received Data: <value> <value> <value> <value>**

DESCRIPTION: Link up notification; configured for loopback.

DATA: (1) fc\_eventTag, (2) granted\_AL\_PA, (3) UlnkSpeed, (4) alpa\_map[0]

ACTION: None required.

#### 4.1.10.6 **elx\_mes1308: Menlo Maint Mode Link up Event <value> rcvd Data: <value> <value> <value>**

DESCRIPTION: Link up notification in Menlo maintenance mode.

DATA: (1) fc\_eventTag, (2) port\_state, (3) vPort fc\_flag

ACTION: None required.

#### 4.1.10.7 **elx\_mes1309: Link Down Event <value> received Data <value> <value> <value>**

DESCRIPTION: The port generated a link down event to the host.

DATA: (1) fc\_eventTag (2) port\_state (3) vPort fc\_flag

ACTION: None required.

#### 4.1.10.8 **elx\_mes1310: Link Up Event npiv not supported in loop topology**

DESCRIPTION: Loop topologies are not supported when NPIV is enabled.

DATA: None

ACTION: Put the link into Fabric mode.

### 4.1.11 **Port Setup Events (1400 to 1499)**

#### 4.1.11.1 **elx\_mes1400: Failed to initialize sgl list.**

DESCRIPTION: Failed to initialize the SGL list during initialization.

DATA: None

ACTION: Reboot the server. If the issue persists, contact Broadcom technical support.

---

**4.1.11.2 elx\_mes1401: Failed to enable pci device.**

DESCRIPTION: Failed to enable a PCI device during initialization.

DATA: None

ACTION: Reboot the server. If the issue persists, contact Broadcom technical support.

**4.1.11.3 elx\_mes1402: Failed to set up pci memory space.**

DESCRIPTION: PCI initialization failed.

DATA: None

ACTION: Reboot the server. If the issue persists, contact Broadcom technical support.

**4.1.11.4 elx\_mes1403: Failed to set up driver resource.**

DESCRIPTION: Driver resource initialization failed.

DATA: None

ACTION: None required.

**4.1.11.5 elx\_mes1404: Failed to set up driver resource.**

DESCRIPTION: Driver resource initialization failed.

DATA: None

ACTION: None required.

**4.1.11.6 elx\_mes1405: Failed to initialize iocb list.**

DESCRIPTION: IOCB initialization failed.

DATA: None

ACTION: None required.

**4.1.11.7 elx\_mes1406: Failed to set up driver resource.**

DESCRIPTION: Initialization failed to set up a driver resource.

DATA: None

ACTION: None required.

**4.1.11.8 elx\_mes1407: Failed to create scsi host.**

DESCRIPTION: Initialization failed to create a SCSI host.

DATA: None

ACTION: None required.

**4.1.11.9 elx\_mes1408: Port Failed POST – portsmpmr=<value>, perr=<port error>, sfi=<sfi reg>, nip=<nip reg>, ipc=<ipc reg>, scr1=<value>, scr2=<value>, hscratch=<value>, pstatus=<port status>**

DESCRIPTION: The adapter's power-on self-test has failed.

DATA: None

ACTION: Make sure the adapter firmware is up to date. Contact Broadcom technical support if the issue persists after a system reboot.

---

**4.1.11.10 elx\_mes1410: Failed to set up pci memory space.**

DESCRIPTION: Initialization failed to set up the PCI memory space.

DATA: None

ACTION: None required.

**4.1.11.11 elx\_mes1411: Failed to set up driver resource.**

DESCRIPTION: Initialization failed to set up a driver resource.

DATA: None

ACTION: None required.

**4.1.11.12 elx\_mes1412: Failed to set up driver resource.**

DESCRIPTION: Initialization failed to set up a driver resource.

DATA: None

ACTION: None required.

**4.1.11.13 elx\_mes1413: Failed to initialize iocb list.**

DESCRIPTION: Initialization failed to initialize the IOCB list.

DATA: None

ACTION: None required.

**4.1.11.14 elx\_mes1414: Failed to set up driver resource.**

DESCRIPTION: Initialization failed to set up a driver resource.

DATA: None

ACTION: None required.

**4.1.11.15 elx\_mes1415: Failed to create scsi host.**

DESCRIPTION: Initialization failed to create a SCSI host.

DATA: None

ACTION: None required.

**4.1.11.16 elx\_mes1416: Failed to allocate sysfs attr**

DESCRIPTION: Initialization failed to allocate a sysfs attribute.

DATA: None

ACTION: None required.

**4.1.11.17 elx\_mes1418: Invalid HBA PCI-device group: <dev\_grp>**

DESCRIPTION: An invalid adapter PCI-device group was detected.

DATA: None

ACTION: None required.

---

**4.1.11.18 elx\_mes1419: Invalid HBA PCI-device group: <dev\_grp>**

DESCRIPTION: An invalid adapter PCI-device group was detected.

DATA: None

ACTION: None required.

**4.1.11.19 elx\_mes1420: Invalid HBA PCI-device group: <dev\_grp>**

DESCRIPTION: An invalid adapter PCI-device group was detected.

DATA: None

ACTION: None required.

**4.1.11.20 elx\_mes1421: Failed to set up hba**

DESCRIPTION: Initialization failed to set up the adapter.

DATA: None

ACTION: None required.

**4.1.11.21 elx\_mes1422: Unrecoverable Error Detected during POST uerr\_lo\_reg=<ue lo>, uerr\_hi\_reg=<ue hi>**

DESCRIPTION: The adapter has notified the driver that it has encountered an unrecoverable error.

DATA: None

ACTION: Perform a dump from the OneCommand Manager application. Then, unload and reload the driver.

**4.1.11.22 elx\_mes1423: HBA Unrecoverable error: uerr\_lo\_reg=<ue lo>, uerr\_hi\_reg=<ue hi>, ue\_mask\_lo\_reg=<ue mask lo>, ue\_mask\_hi\_reg=<ue mask hi>**

DESCRIPTION: The adapter has notified the driver that it has encountered an unrecoverable error.

DATA: None

ACTION: Perform a dump from the OneCommand Manager application. Then, unload and reload the driver.

**4.1.11.23 elx\_mes1424: Invalid PCI device group: <pci\_dev\_grp>**

DESCRIPTION: An invalid adapter PCI-device group was detected.

DATA: None

ACTION: None required.

**4.1.11.24 elx\_mes1425: Invalid PCI device group: <pci\_dev\_grp>**

DESCRIPTION: An invalid adapter PCI-device group was detected.

DATA: None

ACTION: None required.

**4.1.11.25 elx\_mes1426: Invalid PCI device group: <pci\_dev\_grp>**

DESCRIPTION: An invalid adapter PCI-device group was detected.

DATA: None

ACTION: None required.

---

**4.1.11.26 elx\_mes1427: Invalid PCI device group: <pci\_dev\_grp>**

DESCRIPTION: An invalid adapter PCI-device group was detected.

DATA: None

ACTION: None required.

**4.1.11.27 elx\_mes1428: Invalid PCI device group: <pci\_dev\_grp>**

DESCRIPTION: An invalid adapter PCI-device group was detected.

DATA: None

ACTION: None required.

**4.1.11.28 elx\_mes1429: Invalid PCI device group: <pci\_dev\_grp>**

DESCRIPTION: An invalid adapter PCI-device group was detected.

DATA: None

ACTION: None required.

**4.1.11.29 elx\_mes1430: Failed to initialize sgl list.**

DESCRIPTION: The driver failed to initialize the SGL list.

DATA: None

ACTION: None required.

**4.1.11.30 elx\_mes1431: Invalid HBA PCI-device group: <dev\_grp>**

DESCRIPTION: An invalid adapter PCI-device group was detected.

DATA: None

ACTION: None required.

**4.1.11.31 elx\_mes1432: Failed to initialize rpi headers.**

DESCRIPTION: The RPI headers required by the firmware failed to initialize.

DATA: None

ACTION: None required.

**4.1.11.32 elx\_mes1476: Failed to allocate sysfs attr.**

DESCRIPTION: The driver failed to allocate a sysfs attribute.

DATA: None

ACTION: None required.

**4.1.11.33 elx\_mes1477: Failed to set up hba**

DESCRIPTION: The driver failed to set up an adapter.

DATA: None

ACTION: None required.

## 4.1.12 IOCTL Events (1600 to 1699)

None.

## 4.1.13 VPort Events (1800 to 1832)

### 4.1.13.1 **elx\_mes1800: Could not issue unreg\_vpi**

DESCRIPTION: The driver's attempt to unregister a VPI failed.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

### 4.1.13.2 **elx\_mes1801: Create vport work array FAILED: cannot do scsi\_host\_get**

DESCRIPTION: The driver failed to create a working list of VPorts.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

### 4.1.13.3 **elx\_mes1802: HBQ <index>: local\_hbqGetIdx <index> is > than hbqp->entry\_count <count>**

DESCRIPTION: An error occurred when processing a queue related to an adapter in a particular slot.

DATA: (1) hbqno, (2) local\_hbqGetIdx, (3) entry\_count

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

### 4.1.13.4 **elx\_mes1803: Bad hbq tag. Data: <tag> <count>**

DESCRIPTION: An error occurred when processing queue-related tags for an adapter in a particular slot.

DATA: (1) tag, (2) buffer\_count

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

### 4.1.13.5 **elx\_mes1804: Invalid asynchronous event code: <evt code>**

DESCRIPTION: The asynchronous event code that the firmware passed to the driver is invalid.

DATA: None

ACTION: None required.

### 4.1.13.6 **elx\_mes1805: Adapter failed to init. Data: <command> <status> <queue num>**

DESCRIPTION: An error occurred when processing queue-related tags for an adapter in a particular slot.

DATA: (1) mbxCommand, (2) mbxStatus, (3) hbaqno

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

### 4.1.13.7 **elx\_mes1806: Mbox <command> failed. No vport.**

DESCRIPTION: A mailbox command could not be communicated because there was no vPort associated with the mailbox command.

DATA: (1) mbxCommand

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

---

**4.1.13.8 elx\_mes1807: IOCB <value> failed. No vport**

DESCRIPTION: An IOCB command could not be communicated because there was no vPort associated with the mailbox command.

DATA: (1) ulpCommand

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.13.9 elx\_mes1808: Create VPORT failed: NPIV is not enabled: SLI mode <mode>**

DESCRIPTION: The driver failed to create a port because the adapter was in the wrong mode or was not capable of NPIV.

DATA: (1) sli\_rev

ACTION: Load the driver with NPIV enabled on an adapter that supports SLI-3.

**4.1.13.10 elx\_mes1809: Create VPORT failed: Max VPORTs (<vpi>) exceeded.**

DESCRIPTION: The driver failed to create a port because the maximum number of ports supported by the driver is exceeded.

DATA: (1) max\_vpi

ACTION: No Action. The driver cannot create any more VPorts.

**4.1.13.11 elx\_mes1810: Create VPORT failed: Cannot get instance number.**

DESCRIPTION: The driver failed to allocate resources for an adapter and could not assign an instance number.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.13.12 elx\_mes1811: Create VPORT failed: vpi x<vpi>**

DESCRIPTION: The driver failed to create a port and had to eliminate all its resources.

DATA: (1) vpi

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.13.13 elx\_mes1812: vport\_delete failed: Cannot delete physical host**

DESCRIPTION: An attempt to delete a port failed because it was to delete a physical port and not a virtual port. Only VPorts on physical ports can be deleted on an NPIV system.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.13.14 elx\_mes1813: Create VPORT failed. Cannot get sparam.**

DESCRIPTION: The port could not be created because it could not be initialized, possibly due to unavailable resources.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.13.15 elx\_mes1814: Mbox <u.mb.mbxCommand> failed, no vport**

DESCRIPTION: The vPort field of this mailbox command was not completed.

DATA: None

ACTION: None required.

---

**4.1.13.16 elx\_mes1815 Could not issue unreg\_did (default rpis)**

DESCRIPTION: An attempt to unregister RPI failed.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.13.17 elx\_mes1818: VPort failed init, mbxCmd <mailbox command> READ\_SPARM mbxStatus <mailbox status>, rc = <status>**

DESCRIPTION: A pending mailbox command issued to initialize a port failed.

DATA: (1) mbxCommand, (2) mbxStatus, (3) rc

ACTION: This is a software driver error. If this issue persists, report these issues to Broadcom technical support.

**4.1.13.18 elx\_mes1820: Unable to select SLI-3. Not supported by adapter.**

DESCRIPTION: The adapter is not capable of operating in a given mode.

DATA: None

ACTION: This is an informational message. SLI-3 mode is only available on some adapters. Do not attempt to force an adapter to run in SLI-3 mode if that adapter does not support SLI-3 mode. Adapters that do not support SLI-3 are configured to run in SLI-2 mode. Nevertheless, Broadcom recommends you use the auto setting (0).

**4.1.13.19 elx\_mes1821: Create VPORT failed. Invalid WWN format**

DESCRIPTION: The port could not be created due to an invalid WWNN or WWPN format.

DATA: None

ACTION: Provide a valid WWN when creating VPorts.

**4.1.13.20 elx\_mes1822: Invalid <name>: <xx: xx: xx: xx: xx: xx: xx: xx>**

DESCRIPTION: 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]

ACTION: When creating a vPort you must furnish a valid WWN.

**4.1.13.21 elx\_mes1823: Create VPORT failed. Duplicate WWN on HBA.**

DESCRIPTION: The port could not be created because it would duplicate an existing WWNN adapter address. The resources for the port had to be discarded.

DATA: None

ACTION: Provide a unique WWN.

**4.1.13.22 elx\_mes1825: Vport Created.**

DESCRIPTION: This message indicates that a port was created in the system. It is provided at this level to ensure that it always appears at all log levels.

DATA: None

ACTION: No action, informational.

---

**4.1.13.23 elx\_mes1826: Vport Disabled.**

DESCRIPTION: The port had to be disabled in the system.

DATA: None

ACTION: No action, informational.

**4.1.13.24 elx\_mes1827: Vport Enabled.**

DESCRIPTION: The port was enabled after possible recovery from some errors.

DATA: None

ACTION: No action, informational.

**4.1.13.25 elx\_mes1828: Vport Deleted.**

DESCRIPTION: A vPort was deleted.

DATA: None

ACTION: No action, informational.

**4.1.13.26 elx\_mes1830: Signal aborted mbxCmd <command>**

DESCRIPTION: A pending mailbox command was aborted because the thread received a signal.

DATA: None

ACTION: The command is retried.

**4.1.13.27 elx\_mes1831: Create VPORT Interrupted.**

DESCRIPTION: The port creation process was unexpectedly interrupted at a critical time and the operation was unsuccessful.

DATA: None

ACTION: The process was interrupted while creating a vPort. Retry the command.

**4.1.13.28 elx\_mes1832: No pending MBOX command to handle.**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.14 Extended Link Service Events (1833 to 2800)**

**4.1.14.1 elx\_mes1835: Vport discovery quiesce failed: state <port\_state> fc\_flags <fc\_flag> wait msecs <jiffies\_to\_msecs(jiffies - start\_time)>**

DESCRIPTION: Could not pause discovery on this vPort.

DATA: None

ACTION: None required.

**4.1.14.2 elx\_mes1836: Could not issue unreg\_login(all\_rpis) status <rc>**

DESCRIPTION: The unreg\_login cannot be issued.

DATA: None

ACTION: None required.

---

**4.1.14.3 elx\_mes1837: vport\_delete failed: Cannot delete static vport.**

DESCRIPTION: Static VPorts cannot be deleted.

DATA: None

ACTION: None required.

**4.1.14.4 elx\_mes1838: Failed to INIT\_VPI on vpi <vpi> status <rc>**

DESCRIPTION: Failed to INIT\_VPI.

DATA: None

ACTION: None required.

**4.1.14.5 elx\_mes1839: Create VPORT failed. vname allocation failed.**

DESCRIPTION: The driver failed to allocate a buffer for Virtual Machine name.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.14.6 elx\_mes1840: Delete VPORT cannot proceed at this time due to SCSI layer busy.**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.14.7 elx\_mes1920: Exec format error, Dropping Link state event**

DESCRIPTION: There was no dfchba instance available for a Link State event - dropping.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.14.8 elx\_mes1923: Exec format error, Dropping rscn event**

DESCRIPTION: There was no dfchba instance available for a RSCN event - dropping.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.14.9 elx\_mes1926: Exec format error**

DESCRIPTION: There was no dfchba instance available for an IOCTL loopback test - dropping.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.14.10 elx\_mes1927: Exec format error, Dropping temp event**

DESCRIPTION: There was no dfchba instance available for a Temperature event - dropping.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

---

**4.1.14.11 elx\_mes1928: Exec format error, Dropping dump event**

DESCRIPTION: There was no dfchba instance available for a Dump event - dropping.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.14.12 elx\_mes1929: Exec format error**

DESCRIPTION: There was no dfchba instance available for an IOCTL loopback XRI read - dropping.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.14.13 elx\_mes1934: ENOMEM DMA coherent resource unavailable**

DESCRIPTION: The driver failed to allocate a DMA buffer for an IOCTL request.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.14.14 elx\_mes1935: Loopback test did not receive any data**

DESCRIPTION: The driver ran a loopback test, but did not receive a response.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.14.15 elx\_mes1936: ENOMEM Kernel resource unavailable**

DESCRIPTION: The driver failed to allocate a DMA buffer during a loopback test.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.14.16 elx\_mes1944: ENOMEM kernel memory resource unavailable**

DESCRIPTION: The driver failed to allocate a kernel buffer for a timed out I/O request.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.14.17 elx\_mes1949: ENOEXEC NULL parameter passed to function**

DESCRIPTION: The driver tried to post a receive buffer, but no receive buffers are available.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.14.18 elx\_mes1950: ENOMEM IOCB resource not available**

DESCRIPTION: The driver could not allocate the IOCBs needed to post loopback receive buffers.

DATA: (1) (2) (3)

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

---

**4.1.14.19 elx\_mes1951: ENOMEM MBUF resource not available**

DESCRIPTION: The driver failed to get the memory buffer needed for a loopback test.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.14.20 elx\_mes1952: ENOMEM DMA resource not available**

DESCRIPTION: The driver failed to get the DMA buffers needed for a loopback test

DATA: (1) (2) (3)

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.14.21 elx\_mes1957: EPERM Illegal BDE count [<value>]**

DESCRIPTION: The driver received too many receive buffers for a loopback operation.

DATA: (1) receive buffer count

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.14.22 elx\_mes2000: Failed to allocate mbox for READ\_FCF cmd**

DESCRIPTION: The driver failed to allocate a mailbox for the READ\_FCF command.

DATA: None

ACTION: None required.

**4.1.14.23 elx\_mes2001: Unable to allocate memory for issuing SLI\_CONFIG\_SPECIAL mailbox command**

DESCRIPTION: Unable to allocate memory for issuing the SLI\_CONFIG\_SPECIAL mailbox command.

DATA: None

ACTION: None required.

**4.1.14.24 elx\_mes2002: Error Could not grow rpi count**

DESCRIPTION: An error occurred because the RPI count could not be increased.

DATA: None

ACTION: None required.

**4.1.14.25 elx\_mes2007: Only Limited Edition cmd Format supported <iocb.ulpCommand>**

DESCRIPTION: SLI-4 only supports the Limited Edition command format.

DATA: None

ACTION: None required.

**4.1.14.26 elx\_mes2008: Error <rc> posting all rpi headers**

DESCRIPTION: The RPI headers could not be posted to the firmware.

DATA: None

ACTION: None required.

---

**4.1.14.27 elx\_mes2009: Failed to allocate mbox for ADD\_FCF cmd**

DESCRIPTION: The driver failed to allocate a mailbox for the ADD\_FCF command.

DATA: None

ACTION: None required.

**4.1.14.28 elx\_mes2010: Resume RPI Mailbox failed status <status>, mbxStatus <mbx status>.**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.14.29 elx\_mes2011: Unable to allocate memory for issuing SLI\_CONFIG\_SPECIAL mailbox command**

DESCRIPTION: Unable to allocate memory for issuing SLI\_CONFIG\_SPECIAL mailbox command.

DATA: None

ACTION: None required.

**4.1.14.30 elx\_mes2012: Mailbox failed, mbxCmd <mbx\_cmd> READ\_CONFIG, mbxStatus <mbx status>.**

DESCRIPTION: The READ\_CONFIG mailbox command failed.

DATA: None

ACTION: None required.

**4.1.14.31 elx\_mes2013: Could not manually add FCF record 0, status <rc>**

DESCRIPTION: Could not add an FCF record to the FCF list.

DATA: None

ACTION: None required.

**4.1.14.32 elx\_mes2014: Invalid command <iocb.ulpCommand>**

DESCRIPTION: The IOCB command is invalid.

DATA: None

ACTION: None required.

**4.1.14.33 elx\_mes2015: Invalid CT command <iocb.ulpCommand>**

DESCRIPTION: An Invalid Command-Type in the IOCB is not supported.

DATA: None

ACTION: None required.

**4.1.14.34 elx\_mes2017: REG\_FCFI mbxStatus error <mbx status> HBA state <port\_state>.**

DESCRIPTION: The REG\_FCFI mailbox command has failed.

DATA: None

ACTION: None required.

---

**4.1.14.35 elx\_mes2018: REG\_VFI mbxStatus error <mbx status> HBA state <port\_state>.**

DESCRIPTION: The REG\_VFI mailbox command has failed.

DATA: None

ACTION: None required.

**4.1.14.36 elx\_mes2022: INIT VPI Mailbox failed status <status>, mbxStatus <mbxStatus>**

DESCRIPTION: The INIT VPI mailbox command has failed.

DATA: None

ACTION: None required.

**4.1.14.37 elx\_mes2400: Failed to allocate xri for ELS sgl**

DESCRIPTION: Initialization failed to allocate XRI for the extended link service SGL.

DATA: None

ACTION: None required.

**4.1.14.38 elx\_mes2500: EQ\_CREATE mailbox failed with status <shdr\_status> add\_status <shdr\_add\_status>, mbx status <rc>**

DESCRIPTION: The mailbox command sent to create the event queue has failed.

DATA: None

ACTION: None required.

**4.1.14.39 elx\_mes2501: CQ\_CREATE mailbox failed with status <shdr\_status> add\_status <shdr\_add\_status>, mbx status <rc>**

DESCRIPTION: The mailbox command sent to create the completion queue has failed.

DATA: None

ACTION: None required.

**4.1.14.40 elx\_mes2502: MQ\_CREATE mailbox failed with status <shdr\_status> add\_status <shdr\_add\_status>, mbx status <rc>**

DESCRIPTION: The mailbox command sent to create the mailbox queue has failed.

DATA: None

ACTION: None required.

**4.1.14.41 elx\_mes2503: WQ\_CREATE mailbox failed with status <shdr\_status> add\_status <shdr\_add\_status>, mbx status <rc>**

DESCRIPTION: The mailbox command sent to create the work queue has failed.

DATA: None

ACTION: None required.

**4.1.14.42 elx\_mes2504: RQ\_CREATE mailbox failed with status <shdr\_status> add\_status <shdr\_add\_status>, mbx status <rc>**

DESCRIPTION: The mailbox command sent to create the receive queue has failed.

DATA: None

ACTION: None required.

---

**4.1.14.43 elx\_mes2505: EQ\_DESTROY mailbox failed with status <shdr\_status> add\_status <shdr\_add\_status>, mbx status <rc>**

DESCRIPTION: The mailbox command sent to delete the event queue has failed.

DATA: None

ACTION: None required.

**4.1.14.44 elx\_mes2506: CQ\_DESTROY mailbox failed with status <shdr\_status> add\_status <shdr\_add\_status>, mbx status <rc>**

DESCRIPTION: The mailbox command sent to delete the completion queue has failed.

DATA: None

ACTION: None required.

**4.1.14.45 elx\_mes2507: MQ\_DESTROY mailbox failed with status <shdr\_status> add\_status <shdr\_add\_status>, mbx status <rc>**

DESCRIPTION: The mailbox command sent to delete the mailbox queue has failed.

DATA: None

ACTION: None required.

**4.1.14.46 elx\_mes2508: WQ\_DESTROY mailbox failed with status <shdr\_status> add\_status <shdr\_add\_status>, mbx status <rc>**

DESCRIPTION: The mailbox command sent to delete the work queue has failed.

DATA: None

ACTION: None required.

**4.1.14.47 elx\_mes2509: RQ\_DESTROY mailbox failed with status <shdr\_status> add\_status <shdr\_add\_status>, mbx status <rc>**

DESCRIPTION: The mailbox command sent to delete the receive queue has failed.

DATA: None

ACTION: None required.

**4.1.14.48 elx\_mes2510: RQ\_DESTROY mailbox failed with status <shdr\_status> add\_status <shdr\_add\_status>, mbx status <rc>**

DESCRIPTION: The mailbox command sent to delete the receive queue has failed.

DATA: None

ACTION: None required.

**4.1.14.49 elx\_mes2511: POST\_SGL mailbox failed with status <shdr\_status> add\_status <shdr\_add\_status>, mbx status <rc>**

DESCRIPTION: The mailbox command sent to post the SGL pages to the firmware has failed.

DATA: None

ACTION: None required.

---

**4.1.14.50 elx\_mes2512: MODIFY\_EQ\_DELAY mailbox failed with status <shdr\_status> add\_status <shdr\_add\_status>, mbx status <rc>**

DESCRIPTION: The mailbox command sent to modify the event queue delay in the firmware has failed.

DATA: None

ACTION: None required.

**4.1.14.51 elx\_mes2513: POST\_SGL\_BLOCK mailbox command failed status <shdr\_status> add\_status <shdr\_add\_status> mbx status <rc>**

DESCRIPTION: The mailbox command sent to post the SGL pages to the firmware has failed.

DATA: None

ACTION: None required.

**4.1.14.52 elx\_mes2514: POST\_RPI\_HDR mailbox failed with status <shdr\_status> add\_status <shdr\_add\_status>, mbx status <rc>**

DESCRIPTION: The mailbox command sent to post the RPUI header pages to the firmware has failed.

DATA: None

ACTION: None required.

**4.1.14.53 elx\_mes2515: ADD\_FCF\_RECORD mailbox failed with status <rc>**

DESCRIPTION: The mailbox command to add the FCF record has failed.

DATA: None

ACTION: None required.

**4.1.14.54 elx\_mes2521: READ\_FCF\_RECORD mailbox failed with status <shdr\_status> add\_status <shdr\_add\_status>, mbx**

DESCRIPTION: The READ\_FCF\_RECORD mailbox command has failed.

DATA: None

ACTION: None required.

**4.1.14.55 elx\_mes2522: Synchronous READ\_FCF\_RECORD mailbox failed with status <value> add\_status <value>**

DESCRIPTION: Driver failed to read the active FCF Record on an FCoE link - FCF may not be available.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.14.56 elx\_mes2523: Allocated DMA memory size (<alloc\_len>) is less than the requested DMA memory size (<req\_len>)**

DESCRIPTION: The ADD\_FCF\_RECORD mailbox command failed to retrieve the length required from the firmware.

DATA: None

ACTION: None required.

---

**4.1.14.57 elx\_mes2524: Failed to get the non-embedded SGE virtual address**

DESCRIPTION: The READ\_FCF\_RECORD mailbox command could not retrieve the scatter/gather entry that was requested.

DATA: None

ACTION: None required.

**4.1.14.58 elx\_mes2527: Failed to allocate non-embedded SGE array.**

DESCRIPTION: The driver failed to allocate the non-embedded scatter/gather entry array.

DATA: None

ACTION: None required.

**4.1.14.59 elx\_mes2528: Mailbox command <vpi> cannot issue**

DESCRIPTION: The mailbox command could not be issued because the mailbox interrupt is disabled.

DATA: (1) mbxCommand, (2) sli\_flag, (3) flag

ACTION: None required.

**4.1.14.60 elx\_mes2529: Mailbox command <vpi> cannot issue**

DESCRIPTION:

DATA: (1) mbxCommand, (2) sli\_flag, (3) flag

ACTION: None required.

**4.1.14.61 elx\_mes2530: Mailbox command <vpi> cannot issue**

DESCRIPTION: The SLI layer in the driver is inactive.

DATA: (1) mb.mbxCommand, (2) sli\_flag, (3) flag

ACTION: None required.

**4.1.14.62 elx\_mes2531: Mailbox command <cpi> cannot issue**

DESCRIPTION:

DATA: (1) mb.mbxCommand, (2) sli\_flag, (3) flag

ACTION: None required.

**4.1.14.63 elx\_mes2532: Mailbox command <vpi> (<mbxCommand>) cannot issue**

DESCRIPTION: The mailbox bootstrap code detected that the SLI layer is active.

DATA: (1) sli4\_mbox\_opcode, (2) sli\_flag, (3) MBX\_POLL

ACTION: None required.

**4.1.14.64 elx\_mes2533: Mailbox command <vpi> (<mbxCommand>) cannot issue**

DESCRIPTION:

DATA: (1) sli4\_mbox\_opcode, (2) sli\_flag, (3) MBX\_NOWAIT

ACTION: None required.

---

**4.1.14.65 elx\_mes2535: Unsupported RQ count. (<entry\_count>).**

DESCRIPTION: The receive queue ring can only be 512, 1024, 2048, or 4096.

DATA: None

ACTION: None required.

**4.1.14.66 elx\_mes2536: Unsupported RQ count. (<entry\_count>).**

DESCRIPTION: The receive queue ring can only be 512, 1024, 2048, or 4096.

DATA: None

ACTION: None required.

**4.1.14.67 elx\_mes2537: Receive Frame Truncated!**

DESCRIPTION: The receive unsolicited handler detected a truncated frame.

DATA: None

ACTION: None required.

**4.1.14.68 elx\_mes2540: Ring <value> handler: unexpected Rctl <value> Type <value> received**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.14.69 elx\_mes2543: Mailbox command <vpi> (<mbxCommand>) cannot issue**

DESCRIPTION: The mailbox command does not have all of the fields set correctly.

DATA: (1) sli4\_mbx\_opcode, (2) sli\_flag, (3) flag

ACTION: None required.

**4.1.14.70 elx\_mes2544: Mailbox command <vpi> (<mbxCommand>) cannot issue**

DESCRIPTION: The adapter cannot be accessed on the PCI bus.

DATA: (1) sli4\_mbx\_opcode, (2) sli\_flag, (3) flag

ACTION: None required.

**4.1.14.71 elx\_mes2546: New FCF found index <index> tag <event\_tag>**

DESCRIPTION: A new FCF has been found.

DATA: None

ACTION: None required.

**4.1.14.72 elx\_mes2547: Issue FCF scan read FCF mailbox command failed**

DESCRIPTION: Could not read the FCF mailbox command from the firmware.

DATA: None

ACTION: None required.

---

**4.1.14.73 elx\_mes2548: FCF Table full count <count> tag <event\_tag>**

DESCRIPTION: The FCF table is full.

DATA: None

ACTION: None required.

**4.1.14.74 elx\_mes2549: FCF disconnected from network index <index> tag <event\_tag>.**

DESCRIPTION: The FCF has disconnected from the network.

DATA: None

ACTION: None required.

**4.1.14.75 elx\_mes2550: UNREG\_FCFI mbxStatus error <u.mb.mbxStatus> HBA state <port\_state>.**

DESCRIPTION: The unregistered FCFI has failed.

DATA: None

ACTION: None required.

**4.1.14.76 elx\_mes2551: UNREG\_FCFI mbox allocation failed HBA state <port\_state>.**

DESCRIPTION: The allocation for the UNREG\_FCFI mailbox command has failed.

DATA: None

ACTION: None required.

**4.1.14.77 elx\_mes2552: Unregister FCFI command failed rc <rc> HBA state <port\_state>.**

DESCRIPTION: The unregister FCFI mailbox command has failed.

DATA: None

ACTION: None required.

**4.1.14.78 elx\_mes2553: lpfc\_unregister\_unused\_fcf failed to read FCF record HBA state.**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.14.79 elx\_mes2554: Could not allocate memory for fcf record**

DESCRIPTION: The driver was unable to allocate memory for the fcf record.

DATA: None

ACTION: None required.

**4.1.14.80 elx\_mes2555: UNREG\_VFI mbxStatus error <u.mb.mbxStatus> HBA state <port\_state>**

DESCRIPTION: The unregister virtual fabric index mailbox command has failed.

DATA: None

ACTION: None required.

---

**4.1.14.81 elx\_mes2556: UNREG\_VFI mbox allocation failed HBA state <port\_state>**

DESCRIPTION: The driver could not allocate memory for UNREG\_VFI mailbox command.

DATA: None

ACTION: None required.

**4.1.14.82 elx\_mes2557: UNREG\_VFI issue mbox failed rc <rc> HBA state <port\_state>**

DESCRIPTION: The driver could not issue the UNREG\_VFI mailbox command.

DATA: None

ACTION: None required.

**4.1.14.83 elx\_mes2558: ADD\_FCF\_RECORD mailbox failed with status<shdr\_status> add\_status <shdr\_add\_status>**

DESCRIPTION: The ADD\_FCF\_RECORD mailbox command has failed.

DATA: None

ACTION: None required.

**4.1.14.84 elx\_mes2560: Failed to allocate xri for scsi buffer**

DESCRIPTION: The driver failed to allocate an XRI for the SCSI buffer.

DATA: None

ACTION: None required.

**4.1.14.85 elx\_mes2561: Allocated DMA memory size (<allocrlen>) is less than the requested DMA memory size (<reqrlen>)**

DESCRIPTION: The driver could not get the memory required for the number of XRIIs that are trying to be posted.

DATA: None

ACTION: None required.

**4.1.14.86 elx\_mes2562: Failure to allocate an ELS sgl entry: <value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.14.87 elx\_mes2563: Failure to allocate an ELS mbuf: <value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.14.88 elx\_mes2564: POST\_SGL\_BLOCK mailbox command failed status <shdr\_status> add\_status <shdr\_add\_status> mbx status <rc>**

DESCRIPTION: The list of XRI SGEs failed to register with the firmware.

DATA: None

ACTION: None required.

---

**4.1.14.89 elx\_mes2566: Failed to allocate connection table entry**

DESCRIPTION: The driver failed to allocate the connection table entry.

DATA: None

ACTION: None required.

**4.1.14.90 elx\_mes2567: Config region 23 has bad signature**

DESCRIPTION: Configuration region 23 has an invalid signature.

DATA: None

ACTION: None required.

**4.1.14.91 elx\_mes2568: Config region 23 has bad version**

DESCRIPTION: Configuration region 23 has an invalid version.

DATA: None

ACTION: None required.

**4.1.14.92 elx\_mes2572: Failed allocate memory for fast-path per-EQ handle array**

DESCRIPTION: The driver failed to allocate memory for the fast-path per-event queue handle array.

DATA: None

ACTION: None required.

**4.1.14.93 elx\_mes2573: Failed allocate memory for msi-x interrupt vector entries**

DESCRIPTION: The driver failed to allocate memory for MSI-X interrupt vector entries.

DATA: None

ACTION: None required.

**4.1.14.94 elx\_mes2574: Not enough EQs (<sli4\_hba.max\_cfg\_param.max\_eq>) from the pci function for supporting FCP EQs (<cfg\_fcp\_eq\_count>)**

DESCRIPTION: The driver failed to create the minimum number of fast-path event queues.

DATA: None

ACTION: None required.

**4.1.14.95 elx\_mes2576: Failed allocate memory for fast-path EQ record array**

DESCRIPTION: The driver failed to allocate memory for the fast-path event queue record array.

DATA: None

ACTION: None required.

**4.1.14.96 elx\_mes2577: Failed allocate memory for fast-path CQ record array**

DESCRIPTION: The driver failed to allocate memory for the fast-path completion queue record array.

DATA: None

ACTION: None required.

---

**4.1.14.97 elx\_mes2578: Failed allocate memory for fast-path WQ record array**

DESCRIPTION: The driver failed to allocate memory for the fast-path WQ record array.

DATA: None

ACTION: None required.

**4.1.14.98 elx\_mes2581: Not enough WQs (<sli4\_hba.max\_cfg\_param.max\_wq>) from the pci function for supporting FCP WQs (<cfg\_fcp\_wq\_count>)**

DESCRIPTION: The driver was not configured with the minimum number of fast-path work queues.

DATA: None

ACTION: None required.

**4.1.14.99 elx\_mes2598: Adapter Link is disabled.**

DESCRIPTION: The adapter link is disabled.

DATA: None

ACTION: None required.

**4.1.14.100 elx\_mes2599: Adapter failed to issue DOWN\_LINK mbox command rc <rc>.**

DESCRIPTION: The adapter failed to issue a DOWN\_LINK mailbox command.

DATA: None

ACTION: None required.

**4.1.14.101 elx\_mes2600: failed to allocate mailbox memory**

DESCRIPTION: The adapter failed to allocate mailbox memory.

DATA: None

ACTION: None required.

**4.1.14.102 elx\_mes2605: lpfc\_dump\_static\_vport: memory allocation failed**

DESCRIPTION: Memory allocation failed.

DATA: None

ACTION: None required.

**4.1.14.103 elx\_mes2606: No NPIV Fabric support**

DESCRIPTION: No NPIV Fabric support.

DATA: None

ACTION: None required.

**4.1.14.104 elx\_mes2607: Failed to allocate init\_vpi mailbox**

DESCRIPTION: The adapter failed to allocate the init\_vpi mailbox.

DATA: None

ACTION: None required.

---

**4.1.14.105 elx\_mes2608: Failed to issue init\_vpi mailbox**

DESCRIPTION: The driver failed to issue the init\_vpi mailbox.

DATA: None

ACTION: None required.

**4.1.14.106 elx\_mes2609: Init VPI mailbox failed <u.mb.mbxStatus>**

DESCRIPTION: Initialization of the VPI mailbox has failed.

DATA: None

ACTION: None required.

**4.1.14.107 elx\_mes2610: FCF <value> reached driver's book keeping dimension: <value>.**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.14.108 elx\_mes2619: Config region 23 has bad signature**

DESCRIPTION: Configuration region 23 has an invalid signature.

DATA: None

ACTION: None required.

**4.1.14.109 elx\_mes2620: Config region 23 has bad version**

DESCRIPTION: Configuration region 23 has an invalid version.

DATA: None

ACTION: None required.

**4.1.14.110 elx\_mes2707: Ring <Ring#> handler: Failed to allocate iocb Rctl <fh\_rctl> Type <fh\_type> received**

DESCRIPTION: The driver could not allocate an IOCB with which to associate this received frame.

DATA: None

ACTION: None required.

**4.1.14.111 elx\_mes2710: PCI channel disable preparing for reset**

DESCRIPTION: The driver is resetting the PCI slot for this port - starting preparations.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.14.112 elx\_mes2711: PCI channel permanent disable for failure**

DESCRIPTION: The driver has detected a fatal port error - disabling the PCI channel.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

---

**4.1.14.113 elx\_mes2712: lpfc\_aer support attribute value <value> out of range, allowed values are 0|1, setting it to default value of 1**

DESCRIPTION: The lpfc\_aer support attribute value was out of range. The adapter is resetting it to the default value of 1.

DATA: None

ACTION: None required.

**4.1.14.114 elx\_mes2718: Clear Virtual Link Received for VPI <index> tag <event\_tag>**

DESCRIPTION: A Clear virtual link command was received from the Fabric for this VPI.

DATA: None

ACTION: None required.

**4.1.14.115 elx\_mes2719: Invalid response length: tgt <TGT\_ID> lun <LUN> cmnd <CMD> rsplen <RSPLLEN>**

DESCRIPTION: The response length for this FCP command is not supported.

DATA: None

ACTION: None required.

**4.1.14.116 elx\_mes2723 PCI channel I/O abort preparing for recovery**

DESCRIPTION: The driver is preparing the port PCI channel for reset/recovery after an I/O error.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.14.117 elx\_mes2726: READ\_FCF\_RECORD Indicates empty FCF table**

DESCRIPTION: The driver requested the firmware provide a list of FCF entries to connect to and the firmware responded that the FCF table is empty.

DATA: None

ACTION: None required.

**4.1.14.118 elx\_mes2729: Unable to dma\_map\_single request\_buffer: <value>**

DESCRIPTION: The driver was unable to map the SCSI command scatter-gather buffer.

DATA: (1) dma mapping error.

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.14.119 elx\_mes2731: Cannot find fabric controller node.**

DESCRIPTION: The driver was unable to find the Fabric controller node in its data base.

DATA: None

ACTION: None required.

**4.1.14.120 elx\_mes2732: Failed to issue INIT\_VPI mailbox command.**

DESCRIPTION: The driver wanted to send a INIT\_VPI mailbox command to initialize a vPort, but failed to send the mailbox command due to the state of the adapter.

DATA: None

ACTION: None required.

---

**4.1.14.121 elx\_mes2745: Failed to allocate mbox for requesting FCF rediscover**

DESCRIPTION: The driver is trying to rediscover the FCF table, but failed to allocate the memory needed.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.14.122 elx\_mes2746: Requesting for FCF rediscovery failed status <value> add\_status <value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.14.123 elx\_mes2747: Issue FCF scan read FCF mailbox command failed**

DESCRIPTION: The driver wanted to send a read FCF record mailbox command to start fast FCF failover on FCF scan, but failed to send the mailbox command due to the state of the adapter.

DATA: None

ACTION: None required.

**4.1.14.124 elx\_mes2748 Failed to prepare for unregistering HBA's FCF record: rc=<value>**

DESCRIPTION: The driver encountered an initialization error when preparing to rescan the FCF tables and needed to unregister an old FCF record.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.14.125 elx\_mes2749 Failed to prepare for unregistering HBA's FCF record: rc=<value>**

DESCRIPTION: The driver encountered an initialization error when preparing to unregister an FCF and needed to prepare the command.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.14.126 elx\_mes2751: Adapter failed to restart, status reg <status>, FW Data: A8 <0xA8> AC <0xAC>**

DESCRIPTION: The adapter has failed to restart.

DATA: None

ACTION: If the issue persists, report the error to Broadcom technical support.

**4.1.14.127 elx\_mes2752: KILL\_BOARD command failed retval <retval>**

DESCRIPTION: The KILL BOARD mailbox command failed to complete.

DATA: None

ACTION: If the issue persists, report the error to Broadcom technical support.

**4.1.14.128 elx\_mes2753: PLOGI failure DID:<DID> Status:<Status>/<Extended Status>.**

DESCRIPTION: A port login to <DID> was failed either by the driver, the firmware, or the target. The <status> and <extended status> indicate why the port login failed.

DATA: None

ACTION: If the issue persists, report the error to Broadcom technical support.

---

**4.1.14.129 elx\_mes2754: PRLI failure DID:<DID> Status:<Status>/<Extended Status>.**

DESCRIPTION: A process login to <DID> failed either by the driver, the firmware, or the target. The <status> and <extended status> indicate why the process login failed.

DATA: None

ACTION: If the issue persists, report the error to Broadcom technical support.

**4.1.14.130 elx\_mes2755: ADISC failure DID:<DID> Status:<Status>/<Extended Status>.**

DESCRIPTION: An address discovery to <DID> failed either by the driver, the firmware, or the target. The <status> and <extended status> indicate why the address discovery failed.

DATA: None

ACTION: If the issue persists, report the error to Broadcom technical support.

**4.1.14.131 elx\_mes2756: LOGO failure DID:<DID> Status:<Status>/<Extended Status>.**

DESCRIPTION: An N\_Port logout to <DID> failed either by the driver, the firmware, or the target. The <status> and <extended status> indicate why the N\_Port logout failed.

DATA: None

ACTION: If the issue persists, report the error to Broadcom technical support.

**4.1.14.132 elx\_mes2757: Protocol failure detected during processing of FCP I/O op: tgt <tgt ID> lun <LUN> cmnd <CMD> rsplInfo3 <rsplInfo3>**

DESCRIPTION: The FCP response from a target indicated that the response length is valid, but rsplInfo3 indicates that there is no Failure. This is an FCP specification violation by the target.

DATA: None

ACTION: If the issue persists, report the error to Broadcom technical support.

**4.1.14.133 elx\_mes2758: Failed to allocate mbox for READ\_FCF command.**

DESCRIPTION: The driver failed to allocate memory from the mempool for issuing an FCF read mailbox command during the round robin FCF bmask update.

DATA: None

ACTION: None required.

**4.1.14.134 elx\_mes2759: Failed to allocate memory for round robin FCF failover bmask.**

DESCRIPTION: The driver failed to allocate memory for the round robin FCF failover bmask.

DATA: None

ACTION: Make sure the system has enough kernel memory, you may need to reload the driver after the memory issue is resolved.

**4.1.14.135 elx\_mes2762: FCF <value> reached driver's book keeping dimension: <value>**

DESCRIPTION:

DATA: None

ACTION: Inform Broadcom about this.

---

**4.1.14.136 elx\_mes2763: Failed to allocate mbox for READ\_FCF cmd.**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.14.137 elx\_mes2765 Mailbox command READ\_FCF\_RECORD failed to retrieve a FCF record**

DESCRIPTION: The driver failed to find an FCF record when the FCF table scan completed.

DATA:

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.14.138 elx\_mes2772: Issue FCF rediscover mailbox command failed, fault through to FCF dead event**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.14.139 elx\_mes2774: Issue FCF rediscover mailbox command failed, through to CVL event**

DESCRIPTION:

DATA: None

ACTION: None required.

## 4.1.15 New Events

**4.1.15.1 elx\_mes2796: Mailbox memory allocation failed**

DESCRIPTION: The driver failed to get memory resources to release an RPI.

DATA: None.

ACTION: None. The driver's heap is exhausted. A server reboot is required to fix the exhaustion. Contact Broadcom technical support if the issue persists.

**4.1.15.2 elx\_mes2798: Unreg\_vpi failed vpi <value>, mb status = <value>**

DESCRIPTION: The driver attempted to unregister a virtual port index and failed. The failure status is shown.

DATA: None.

ACTION: This condition is not catastrophic, but is unexpected. If the issue persists, contact Broadcom technical support.

**4.1.15.3 elx\_mes2813: Mgmt IO is Blocked <value> - mbox cmd <value> still active.**

DESCRIPTION: The HBA's management interface is marked as blocked in preparation for an online or offline state transition. All user space access to the HBA via libdfc interface will be blocked.

DATA: None.

ACTION: None. Notification of a run-state change only.

---

**4.1.15.4 elx\_mes2822: IOCB failed <value> iotag <value> xri <value>**

DESCRIPTION: The driver attempted to drain an internal queue and failed. The failure reason and some state variables are written to the console.

DATA: None

ACTION: None required. This should be a transient condition. If not, contact Broadcom technical support.

**4.1.15.5 elx\_mes2823: txq empty and txq\_cnt is <value>**

DESCRIPTION: The driver has detected a discrepancy between the elements queued to the txq and the counter tracking the number of items.

DATA: None

ACTION: None required. The driver has nothing to do except correct the counter - the txq is empty.

**4.1.15.6 elx\_mes2824: Cannot re-enable interrupt after slot reset.**

DESCRIPTION: The driver failed to re-enable interrupts following a PCI slot reset command.

DATA: None

ACTION: A system reboot may be required to fully recover. Contact Broadcom technical support if the issue persists.

**4.1.15.7 elx\_mes2825: Unknown PCI error state: <value>**

DESCRIPTION: The driver writes this message to the console if the PCI subsystem has detected an error on a Broadcom port and called the driver. The driver reacts by resetting the port.

DATA: None

ACTION: None required. The driver resets the device in an attempt to recover. Contact Broadcom technical support if the issue persists.

**4.1.15.8 elx\_mes2826: PCI channel disable preparing for reset**

DESCRIPTION: The driver writes this message to the console if it is preparing the port for a reset operation.

DATA: None

ACTION: None required. This message is notification of a corrective measure. Contact Broadcom technical support if the issue persists.

**4.1.15.9 elx\_mes2827: PCI channel permanent disable for failure**

DESCRIPTION: The driver writes this message to the console if a recovery mechanism has failed and the driver wants to mark the port with a permanent failure.

DATA: None

ACTION: A system reboot may correct the failure. If not, contact Broadcom technical support.

**4.1.15.10 elx\_mes2828: PCI channel I/O abort preparing for recovery**

DESCRIPTION: The driver writes this message to the console when it is preparing the port for a recovery operation.

DATA: None

ACTION: None required. This is a notification message for the recovery action.

---

**4.1.15.11 elx\_mes2831: FLOGI response with cleared Fabric bit fcf\_index <value> Switch Name <value> Fabric Name <value>**

DESCRIPTION: When the driver completed a FLOGI, the common service parameters did not indicate an FPort or NPort remote node. The driver treats this as an error.

DATA: None

ACTION: Validate the external cable connection and FPort/Nport configuration. Contact Broadcom technical support if the issue persists.

**4.1.15.12 elx\_mes2856: Config Port Security Crypto Error: <value>, pmb->u.mb.un.varCfgPort.sec\_err**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.13 elx\_mes2858: FLOGI failure Status:<value>/<value> TMO:<value>**

DESCRIPTION: The driver issued a FLOGI, but never received any completion within the timeout period. The driver is marking the FLOGI as failed and stops discovery.

DATA: None

ACTION: Check your Fabric to verify that it is operating correctly. Contact Broadcom technical support if the issue persists.

**4.1.15.14 elx\_mes2860: SLI authentication is required for INIT\_LINK but has not been done yet**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.15 elx\_mes2862: FCF (<value>) matches property of in-use FCF (<value>)**

DESCRIPTION: The driver has found an FCF record that matches the properties of the current FCF record, except for the VLAN ID and Index. The driver will attempt to use this FCF.

DATA: None

ACTION: None required. The driver is in its FCF discovery phase and is trying to recover a match to its in-use FCF.

**4.1.15.16 elx\_mes2863: New FCF (<value>) matches property of in-use FCF (<value>)**

DESCRIPTION: The driver has found a new FCF record that matches the properties of the current FCF record, but the record instance numbers do not match.

DATA: None

ACTION: None required. The driver is in its FCF discovery phase and is trying to recover a match to its in-use FCF.

**4.1.15.17 elx\_mes2877: FCP XRI exchange busy wait time: <value> seconds**

DESCRIPTION: An FCP exchange cannot be released - no port completion. The driver is waiting.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

---

**4.1.15.18 elx\_mes2878: ELS XRI exchange busy wait time: <value> seconds**

DESCRIPTION: An extended link service exchange cannot be released - no port completion. The driver is waiting.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.15.19 elx\_mes2881: RRQ failure DID:<value> Status:<value>/<value>**

DESCRIPTION: A Driver Reinstate Recovery Qualifier request failed - driver write target DID and status values.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.15.20 elx\_mes2882: RRQ completes to NPort <value> with no ndlp. Data: <value> <value> <value>**

DESCRIPTION: The driver completes a Reinstate Recovery Qualifier, but there is no node association.

DATA: (1) Status (2) Reason (3) IoTag

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.15.21 elx\_mes2884: Vport array allocation failed**

DESCRIPTION: The driver could not create a buffer list of VPorts.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

**4.1.15.22 elx\_mes2885: Port Status Event: port status reg <value>, port smphr reg <value>, error 1=<value>, error 2=<value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.23 elx\_mes2886: HBA Error Attention on unsupported if type <value>.**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.24 elx\_mes2887: Reset Needed: Attempting Port Recovery**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.25 elx\_mes2888: Unrecoverable port error following POST: port status reg <value>, port smphr reg <value>, error 1=<value>, error 2=<value>**

DESCRIPTION:

DATA: None

ACTION: None required.

---

**4.1.15.26 elx\_mes2889: Port overtemperature event, taking port**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.27 elx\_mes2890: Port error detected during port reset: wait\_tmo: <value>, port status reg <value>, error 1=<value>, error 2=<value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.28 elx\_mes2891: Init VFI mailbox failed <value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.29 elx\_mes2892: Failed to allocate init\_vfi mailbox**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.30 elx\_mes2893: Failed to issue init\_vfi mailbox**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.31 elx\_mes2894: SLI\_INTF reg contents invalid sli\_intf reg <value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.32 elx\_mes2895: Non FC link Event detected.(<value>)**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.33 elx\_mes2897: The mboxq allocation failed**

DESCRIPTION:

DATA: None

ACTION: None required.

---

**4.1.15.34 elx\_mes2898: The lpfc\_dmabuf allocation failed**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.35 elx\_mes2899: The mbuf allocation failed**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.36 elx\_mes2919: Failed to release resource extents for type <value> - Status <value> Add'l Status <value>. Resource memory not released.**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.37 elx\_mes2920: Failed to alloc Resource IDs rc = <value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.38 elx\_mes2925: Failed to issue CT ABTS RSP <value> on xri <value>, Data <value>**

DESCRIPTION: The driver tried and failed to issue a response to an unsolicited abort sequence from the SAN.

DATA: None

ACTION: None required.

**4.1.15.39 elx\_mes2929: Resource Extent Opcode <value> is unsupported**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.40 elx\_mes2930: Failed to get resource extents Status <value> Add'l Status <value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.41 elx\_mes2936: Could not find Vport mapped to vpi <value>**

DESCRIPTION:

DATA: None

ACTION: None required.

---

**4.1.15.42 elx\_mes2945: SLI\_CONFIG(mse) rd, ext\_buf\_cnt(<value>) out of range(<value>)**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.43 elx\_mes2946: SLI\_CONFIG(hbd) rd, ext\_buf\_cnt(<value>) out of range(<value>)**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.44 elx\_mes2950: Failed SLI\_CONFIG(hbd) rd (<value>)**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.45 elx\_mes2953: SLI\_CONFIG(mse) wr, ext\_buf\_cnt(<value>) out of range(<value>)**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.46 elx\_mes2954: SLI\_CONFIG(hbd) wr to interface type:<value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.47 elx\_mes2999: Unsupported SLI4 Parameters Extents and RPI headers enabled.**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.48 elx\_mes3008: No available Resource Extents for resource type <value>: Count: <value>, Size <value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.49 elx\_mes3010: Run link diag test mailbox failed with mbx\_status <value> status <value>, add\_status <value>**

DESCRIPTION:

DATA: None

ACTION: None required.

---

**4.1.15.50 elx\_mes3029: SLI\_CONFIG(hbd) rd to interface type:<value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.51 elx\_mes3042: Failed SLI\_CONFIG(hbd) wr (<value>)**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.52 elx\_mes3045: SLI\_CONFIG(hbd) wr, ext\_buf\_cnt(<value>) out of range(<value>)**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.53 elx\_mes3061: Last IDX <value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.54 elx\_mes3069: Clearing FCP rules**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.55 elx\_mes3070: ipc\_clr\_ndlps\_pri: fcp\_priority sz = <value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.56 elx\_mes3071: no memory for priority rules**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.57 elx\_mes3084: Allocated DMA memory size (<value>) is less than the requested DMA memory size (<value>)**

DESCRIPTION:

DATA: None

ACTION: None required.

---

**4.1.15.58 elx\_mes3089: Failed to allocate queues**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.59 elx\_mes3103: Adapter Link is disabled.**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.60 elx\_mes3104: Adapter failed to issue**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.61 elx\_mes3105: failed to allocate mailbox memory**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.62 elx\_mes3112: lpfc\_link\_speed attribute cannot be set to <value>. Speed is not supported in loop mode.**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.63 elx\_mes3113: Loop mode not supported at speed <value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.64 elx\_mes3125: Not receiving unsolicited event**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.65 elx\_mes3141: Loopback mode: <value> not supported**

DESCRIPTION:

DATA: None

ACTION: None required.

---

**4.1.15.66 elx\_mes3142: Failed loopback test issue iocb: <value> iocb\_stat: <value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.67 elx\_mes3143: Port Down: Firmware Restarted**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.68 elx\_mes3144: Port Down: Debug Dump**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.69 elx\_mes3145: Port Down: Provisioning**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.70 elx\_mes3147: Fast-path EQs not allocated**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.71 elx\_mes3148: Fast-path FCP CQ array not allocated**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.72 elx\_mes3149: Fast-path FCP WQ array not allocated**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.73 elx\_mes3150: No privilege to perform the requested access: <value>**

DESCRIPTION: The driver lacks the necessary privileges for the requested action.

DATA: Requested action

ACTION: Contact Broadcom technical support.

---

**4.1.15.74 elx\_mes3151: PCI bus read access failure: <value>**

DESCRIPTION: A read to the adapter's PCI registers failed.

DATA: Contents read from requested register.

ACTION: Contact Broadcom technical support.

**4.1.15.75 elx\_mes3152: Unrecoverable error, bring the port offline**

DESCRIPTION: Unable to recover the adapter port post hardware error. Taking the port offline.

DATA: None

ACTION: Contact Broadcom technical support.

**4.1.15.76 elx\_mes3153: Fail to perform the requested access: <value>**

DESCRIPTION: The driver failed to execute the requested management action.

DATA: Register action.

ACTION: Contact Broadcom technical support.

**4.1.15.77 elx\_mes3154: BLS ABORT RSP failed, data <value/value>**

DESCRIPTION: The driver issued BLS ABORT Response failed to complete.

DATA: I/O status and I/O reason

ACTION: Contact Broadcom technical support.

**4.1.15.78 elx\_mes3161: Failure to post els sgl to port.**

DESCRIPTION: The driver failed to port the scatter gather list to the adapter.

DATA: None

ACTION: Contact Broadcom technical support.

**4.1.15.79 elx\_mes3172: SCSI layer issued Host Reset Data: <value>**

DESCRIPTION: The SCSI layer issued a host reset request to the driver.

DATA: Reset result

ACTION: Contact Broadcom technical support.

**4.1.15.80 elx\_mes3175: Failed to enable interrupt**

DESCRIPTION: The driver failed to get interrupts re-enabled after an adapter reset.

DATA: None

ACTION: Contact Broadcom technical support.

**4.1.15.81 elx\_mes3176: Misconfigured Physical Port - Port Name <value>**

DESCRIPTION: The driver has detected an unknown firmware name.

DATA: Detected firmware name

ACTION: Contact Broadcom technical support.

---

**4.1.15.82 elx\_mes3177: Nport <value>, DID <value> in wrong state for discovery to start. Failing request Data <value>**

DESCRIPTION: The driver is trying to restart SAN discovery and detected an NPort is in the wrong initial state.

DATA: NPort ID, FCID, NPort state

ACTION: Reset the link. If this fails, contact Broadcom technical support.

**4.1.15.83 elx\_mes3249: Unable to allocate memory for QUERY\_FW\_CFG mailbox command**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.84 elx\_mes3250: QUERY\_FW\_CFG mailbox failed with status <value> add\_status <value>, mbx status <value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.85 elx\_mes3252: WQ doorbell offset not supported**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.86 elx\_mes3262: RQ doorbell format not supported**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.87 elx\_mes3263: WQ failed to memmap pci barset: <value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.88 elx\_mes3265: WQ doorbell format not supported: <value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.89 elx\_mes3269: RQ failed to memmap pci barset: <value>**

DESCRIPTION:

DATA: None

ACTION: None required.

---

**4.1.15.90 elx\_mes3270: RQ doorbell format not supported: <value>**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.91 elx\_mes3279: Invalid provisioning of rpi**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.92 elx\_mes3280: Invalid provisioning of vpi**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.93 elx\_mes3281: Invalid provisioning of xri**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.94 elx\_mes3282: Invalid provisioning of vfi**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.95 elx\_mes3300: In-use FCF modified, perform FCF rediscovery**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.96 elx\_mes3303: Failed to obtain vport vpi**

DESCRIPTION:

DATA: None

ACTION: None required.

**4.1.15.97 elx\_mes3317: HBA not functional: IP Reset Failed after <value> retries, try: echo fw\_reset > board\_mode**

DESCRIPTION:

DATA: None

ACTION: None required.

#### 4.1.15.98 elx\_mes3321: Recovered mailbox <value>

DESCRIPTION:

DATA: None

ACTION: None required.

#### 4.1.15.99 elx\_mes9000: Allocated DMA memory size (<value>) is less than the requested DMA memory size (<value>)

DESCRIPTION:

DATA: None

ACTION: None required.

## 4.2 Troubleshooting the NIC Driver

This section provides troubleshooting information for the NIC driver. Be sure to check the readme.txt file located on CD1 for other troubleshooting issues.

[Table 16](#) provides troubleshooting information for the NIC driver.

**Table 16 Troubleshooting the NIC Driver**

Issue	Resolution
1. When there is a great deal of network traffic in some VMs, a few VMs appear to have lost network connectivity. 2. Several "alloc_skb() failed" messages appear in the log file: /proc/vmware/log	This could be due to low configured value for netPktHeapMaxSize. Try increasing it to a higher value. To read the current value, run: <code># esxcfg-advcfg -j netPktHeapMaxSize</code> (A value of 0 indicates the default - 64 MB) To increase the size to (for example, 128 MB), run: <code># esxcfg-advcfg -k 128 netPktHeapMaxSize</code> (netPktHeapMaxSize can also be configured through the VI Client using Configuration > Advanced Settings > VMKernel.) After configuring the size, reboot the system.
Unable to ping from one VM to another VM.	The NIC driver creates two vmnic interfaces, one for each port. If these interfaces are configured as uplinks in two separate vSwitches, the VMs in each of these switches are in separate networks with no network path between them. Thus, pinging between the VMs in the two groups fails. If you want all these VMs in the same network, configure them as teaming uplinks to one vSwitch option. Each of the vmnics, vmnic1 to vmnic16, must be configured in a separate vSwitch. In this configuration, there is no network path between the vSwitches and pinging between these VMs does not work.
When inserting or removing a 1Gb SFP-RJ45 module on an adapter without RJ45 copper cables attached, the operating system indicates link up or down status. When inserting or removing copper cables attached to a switch to the module, link up or down events are not reported to the operating system.	No solution to this issue exists as the link status is not reported due to a PHY limitation on the card.
Flow control setting is not stored per port after rebooting the system.	With flow control, there is no persistence across reboot. TX and RX are always on when the system starts. For persistence, run a config command from an RC file at reboot.

## 4.2.1 NIC Event/Error Logging for ESXi 5.1

### 4.2.1.1 Retrieving ESXi Server NIC Error Log Codes

For ESXi Server systems, the NIC driver generates error codes to the /var/log/vmkernel log file. The vmkernel log file is an ASCII text file that can be viewed and searched with a text editor such as vim. The vmkernel log file is automatically rotated as it gets larger, and the rotated log files are named vmkernel.x, where x is an integer.

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

```
#cd /var/log  
#less vmkernel
```

For example, you might see the following message:

```
Sep 9 19:48:04 esx-server vmkernel: WARNING: Found a OneConnect card in Gen 1 x8  
PCI-e slot. Should be in Gen 2, x8 slot for best performance.
```

### 4.2.1.2 ESXi Server NIC Event Log Entries

[Table 17](#) below provides a list of ESXi Server network event log error messages. It includes the severity of the error, the message displayed, and the message description. When reporting a problem with the adapter to Broadcom, check the message log (/proc/vmware/log) and report any of these entries that may be present.

**NOTE**

In [Table 17](#), <D>, <DD>, or <DDD> in the 'Message Displayed' column refers to decimal values that appear in the error messages.

**Table 17 ESXi Server NIC Event Log Entries**

Severity	Displayed Message	Description
Error	OneConnect POST failed	The power-on self-test of the adapter failed. This indicates either a hardware or a firmware problem. Try rebooting the system after a reset.
Error	OneConnect initialization failed	Either the initialization of the adapter or the allocation of some resource for initializing the driver failed. In most cases, this message is accompanied by another more specific error message. Try rebooting the system after a power cycling. If the problem persists, this could indicate a hardware problem or corrupted firmware.
Error	RSS cannot be supported on this interface as SRIOV is enabled in the BIOS	RSS cannot be enabled when SR-IOV is enabled in the system BIOS.
Error	RSS cannot be supported on this interface as VFs are created	On LPe16202/OCe15100 and LPe16000-series adapters, RSS is not supported when VFs are already created on the interface.
Error	RSS cannot be supported when msix is disabled	MSI-X must be enabled to enable RSS.
Error	RSS is not supported on this interface	RSS is not supported on this interface.
Error	RSS cannot be supported on 1G interface	RSS cannot be enabled on 1GbE interfaces.
Error	RSS cannot be supported as the host does not have minimum required CPUs.	The host currently does not have enough CPUs to enable RSS. The host must have at least four cores and 4 to 8GB of RAM.
Warning	Using INTx interrupts. NetQueues feature will be disabled	The driver could not allocate the MSI-X vector for interrupts. The driver might continue to work, but the performance may be impacted.

**Table 17 ESXi Server NIC Event Log Entries (Continued)**

Severity	Displayed Message	Description
Warning	WARNING: Found a OneConnect card in Gen <D> x<D> PCIe slot. Should be in Gen 2, x8 slot for best performance	The adapter is an x8, Gen2 PCIe® device. For best performance, the adapter should be installed in a Gen2 PCIe slot 8 or 16 channels wide. The driver displays this warning if it finds the device in a slower or narrower PCIe slot. The device continues to work with reduced performance.
Warning	Command to get pause frame settings failed	The firmware command to get the PAUSE settings failed.
Warning	Command to set pause frame settings failed	The firmware command to change the PAUSE settings failed.
Warning	Command to apply MAC address filter failed	The driver could not set the MAC address filter on the hardware. The device continues to work. A performance impact may occur.
Warning	Command to delete MAC address filter failed	The firmware command to delete a MAC address filter failed. The device should continue to work.
Warning	Unable to get Firmware Version	The command to get the firmware revision number failed. The version number is not shown. The device will continue to work.
Warning	Did not receive completions for all TX requests	While the driver was unloading, some outstanding transmit requests were found. This is an indication that the hardware is not functioning properly.
Warning	Failed to register char device	Could not create the char device used for certain management functions. The driver will continue to work. You may not be able to use HBACMD to interact with the device.
Warning	alloc_skb failed. Try increasing netPktHeapMaxSize	Could not allocate the skb structure to send a frame received from the network to the operating system. This is a transient failure that can be ignored. If this message appears continually, you may need to allocate more memory to the network heap. For example, to increase the heap size to 128MB, run:  <code># esxcfg-advcfg -k 128 netPktHeapMaxSize</code>
Warning	Invalid MTU requested. MTU must be between 64 and 9000 bytes.	There was an invalid MTU size in the MTU configuration IOCTL. The MTU will not be changed.
Warning	Invalid vlan priority labeled. Must be 0 - 7	A request to set a VLAN priority tag was made with an invalid value.
Warning	Failed to allocate memory for pass through command	The memory allocation for a pass through command failed. The driver will continue to function. The configuration utility that issued the pass through IOCTL will fail.
Warning	Pass through command failed. opcode <DDD>, status 0x<XXX>	The pass through firmware command with the indicated opcode failed. The driver should continue to function. The configuration utility that issued the pass through IOCTL will fail.
Warning	Command to modify EQ delay failed	The firmware command to change the event queue delay failed. The driver will continue to function. Adaptive interrupt coalescing does not function correctly.
Warning	Unqualified SFP+ module detected on Port <adapter port number> from <vendor name> part number: <vendor part number>	The specified port has an unqualified SFP+ module inserted into it.

## 4.2.2 NIC Adapter Firmware Error

The following POST message appears if you have loaded firmware on the adapter that the controller does not support:

POST Error : Firmware halted. This firmware does not support this controller.

## 4.2.3 NIC Informational Log Groups for ESXi 5.5 and 6.0

The elxnet driver informational logs for ESXi 5.5 and 6.0 are categorized into groups. You can enable or suppress the logs of a specific group by setting or clearing the corresponding bit on DebugMask. The DebugMask is a bit-vector (uint32) and each bit represents a group.

[Table 18](#) lists the elxnet driver informational log categories.

**Table 18 NIC Informational Log Groups for ESXi 5.5 and 6.0**

Hexadecimal Value	Description
0x1	DRIVER
0x2	UPLINK
0x4	QUEUE
0x8	INTR
0x10	MCC
0x20	TX
0x40	RX
0x80	MGMT
0x100	WORKER
0x200	SRIOV
0x400	EVENT
0x800	VLAN
0x1000	VXLAN

The default enabled groups include DRIVER, UPLINK, MCC, QUEUE, SRIOV, and VXLAN.

Log groups reduce informational log clutter at the default log level and aid in debugging by enabling or disabling group logs. Warning and error logs are always logged, regardless of the debug mask. Refer to [Table 17, ESXi Server NIC Event Log Entries](#), for more information on the warning and error logs.

The debug masks control which groups are enabled or disabled, and they are not meant to be a replacement of the log level adjustments supported by vsish. The vmkernel log levels control the verbosity of logs at each group. For example, level 0 results in less logging and level 4 results in more verbose logs.

There are two ways to modify the debug mask, with the ModuleParam debugMask command or the esxcli-Plugin dbgmask command.

### 4.2.3.1 ModuleParam debugMask

**NOTE**

You must reload (or reboot) the driver after modifying the debugMask module parameter.

This is a global debugMask command; therefore, the same value applies to all Broadcom Emulex NIC PFs.

Use either of these commands:

```
esxcfg-module -s "debugMask=0x0120" elxnet
```

- Or -

```
esxcli system module parameters set -p "debugMask=0x0120" -m elxnet
```

#### 4.2.3.2 esxcli-Plugin

**NOTE** A driver reload (or reboot) is not required after modifying the dbgmask parameter.

This is a per-PF debugMask command; therefore, the value applies to a specific Broadcom Emulex NIC PF (PF with SBDF is specified with "-p").

Command to Get:

```
esxcli elxnet dbgmask get -p 0000:03:00.1
```

Command to Set:

```
esxcli elxnet dbgmask set -p 0000:03:00.1 0x000000ff
```

## 4.3 Troubleshooting the iSCSI Driver

This section provides troubleshooting information for the iSCSI driver.

[Table 19](#) provides information on troubleshooting the iSCSI driver.

**Table 19 Troubleshooting the iSCSI Driver**

Issue	Resolution
Mutual CHAP authentication cannot be configured with the vSphere client.	One-way CHAP is the only authentication method that can be configured with the vSphere Client. To set Mutual CHAP authentication, you must use iSCSISelect.
IPv6 support is not available in vSphere client for hardware iSCSI adapters.	To set the IPv6 address, you must use iSCSISelect.

### 4.3.1 Retrieving iSCSI Error Log Codes

The iSCSI driver generates error codes to the /var/log/vmkernel log file. The vmkernel log file is an ASCII text file that can be viewed and searched with a text editor such as vim. The vmkernel log file is automatically rotated as it gets larger, and the rotated log files are named vmkernel.x, where x is an integer.

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

```
#cd /var/log  
#vim vmkernel.log
```

For example, you may see the following message:

```
Sep 9 19:48:04 esx-server vmkernel: OneConnect iSCSI Driver: The be2iscsi driver received a Task Management Function that is not supported and rejected this request. The error log entry immediately following this entry will indicate the TMF function code that was rejected.
```

### 4.3.2 iSCSI Error Log Code Entries

The following is a brief description of the error log codes generated by the iSCSI driver. It includes the message displayed, the meaning of the message, and the recommended resolution.

**Table 20 iSCSI Error Log Code Entries**

Severity	Message	Recommended Resolution
Error	The be2iscsi driver failed to load because initialization failed during a power management boot-up.	This failure may be due to the firmware not being present or not running currently. This failure may also indicate a hardware problem.
Error	The be2iscsi driver was unable to map one or more PCI Base Address Register and hence failed to load.	This failure may indicate a low memory condition or a hardware error.
Error	The be2iscsi driver ignored a configuration entry since the entry was invalid.	Check the registry configuration for any recent entries added for driver parameters. The invalid entry must be removed or corrected.
Error	The be2iscsi driver failed to load due to memory allocation failure.	This failure occurred due to a failed memory allocation in the driver. Check low memory conditions.
Error	The be2iscsi driver failed to load because initialization failed during normal boot-up.	This failure may be due to the firmware not being present or not running currently. This failure may also indicate a hardware problem.
Error	An internal API failed in be2iscsi driver during initialization.	This failure may indicate a low memory condition.
Error	An Unrecoverable Error was detected by the be2iscsi driver. Following this error log entry, the next 3 entries will indicate the error codes.	This may be due to hardware errors or due to unhandled exceptions in the hardware or firmware.
Error	The be2iscsi driver failed an IOCTL request since the number of scatter gather elements required for the IOCTL buffer exceeded the adapter's firmware limit. Following this error log entry, the next entry will indicate the IOCTL opcode and the payload length requested.	This error may indicate an incorrect configuration option for the iSCSI driver. It may also indicate a low memory condition.
Error	The be2iscsi driver detected an error during offloading the iSCSI connection. The operation will be retried again. Following this error log entry, the next entry will indicate the session handle and the firmware error code.	This may indicate a target is in error or may point to transient network connectivity issues. It may also indicate a firmware error.
Error	The be2iscsi driver did not receive an iSCSI command window update up to 25 seconds during I/O operations. Following this error log entry, the next entry will indicate the session handle where this error occurred. The be2iscsi driver will trigger a session recovery on the session and continue.	Check for any errors reported at the target. The iSCSI initiator is only supported with certified targets. Check for the latest software at the target vendor's website. Check for the latest software at <a href="http://www.broadcom.com">www.broadcom.com</a> . If the above fails, contact Broadcom technical support.
Warning	The be2iscsi driver received an invalid iSCSI Command Sequence Number update from the target. Following this error log entry, the next three entries will indicate the session handle and the iSCSI parameters - MaxCmdSN and ExpCmdSN respectively.	Check for any errors reported at the target. The iSCSI initiator is only supported with certified targets. Check for the latest software at the target vendor's website. Check for the latest software at <a href="http://www.broadcom.com">www.broadcom.com</a> . If the above fails, contact Broadcom technical support.

**Table 20 iSCSI Error Log Code Entries (Continued)**

Severity	Message	Recommended Resolution
Warning	A connection to the target was lost for a period exceeding the Extended Timeout (ETO). The error log entry immediately following this entry will indicate the session ID of the target that lost the connection. Event log entries from the disk subsystem indicating that the drives were lost will exist. If any I/Os were in progress, the system may see I/O errors or failures.	Check the connection to the target or the state of the target device. If the target is made available, any sessions that existed previously will be reestablished and the devices will be available for I/O.
Warning	The be2iscsi driver received a Task Management Function that is not supported and rejected this request. The error log entry immediately following this entry will indicate the task management function code that was rejected.	The operating system version is not supported.
Warning	The be2iscsi driver received a Task Management Function Abort request for an I/O request that is not present with the driver.	This message may indicate a slow connection to the target. Check network connectivity to the target for any errors.
Error	The be2iscsi driver encountered a mismatched version of the firmware running on the board. This error may be followed by more error codes 0x31840001 or 0x31880001 indicating that the be2iscsi driver failed to load.	This failure indicates that the driver version that is running on the system does not match the version of the firmware flashed on the board. Fix this by installing the desired version.
Error	The be2iscsi driver detected a failure in the hardware during initialization. This error may be followed by more error codes 0x31840001 or 0x31880001 indicating that the be2iscsi driver failed to load.	This failure indicates that the hardware has not been initialized or is malfunctioning. This may also indicate that the firmware is not running correctly.
Warning	Both Port 0 and Port 1 links were down for a period exceeding the Link Down Timeout (LDTO). If the initiator has connection to the target, there will be event log entries from the disk subsystem indicating that the drives were lost. If any I/Os were in progress, the system may see I/O errors or failures.	Check the links to the adapter. If the link is reestablished, any sessions that existed previously will be reestablished and the devices will be available for I/O.
Warning	Both Port 0 and Port 1 links are down.	Check the links to the adapter.

### 4.3.3 Additional iSCSI Driver Messages

The following iSCSI error messages are returned if you specify illegal options when loading the driver:

**Table 21 Additional iSCSI Driver Messages**

Message
be2iscsi <PCI_DEVICE_LOCATION>: WARNING: dic value = %d out of range. Valid range is 0 - 1. Using default value = 1
be2iscsi <PCI_DEVICE_LOCATION>: WARNING: eto value = %d out of range. Valid range is 0 - 3600. Using default value = 30
be2iscsi <PCI_DEVICE_LOCATION>: WARNING: ldt0 value = %d out of range. Valid range is 0 - 3600. Using default value = 20
be2iscsi <PCI_DEVICE_LOCATION>: WARNING: ios_per_ctrl value = %d out of range. Valid range is 1 - 512. Using default value = 512
be2iscsi <PCI_DEVICE_LOCATION>: WARNING: max_io_size value = %d out of range. Valid range is 4 - 512 (KByte). Using default value = 128
be2iscsi <PCI_DEVICE_LOCATION>: WARNING: tmf_reset value= %d out of range. Valid range is 1 - 3. Using default value= 1

**NOTE**

For the <PCI\_DEVICE\_LOCATION>, the driver returns a value in the following format – 0000:00:00.0 (for example, 0000:05:00.4).

%d represents a signed integer outside of the valid range. The driver returns a number in the Warning message.

Any other messages will be returned in the following form:

be2iscsi <PCI\_DEVICE\_LOCATION>: FUNCTION\_NAME:LINE:MESSAGE

where the <PCI\_DEVICE\_LOCATION> is a value returned by the driver in the following format – 0000:00:00.0.

For example:

be2iscsi 0000:05:00.4: 1088: kmalloc failed.

You are asked to find the function <FUNCTION\_NAME> at line <LINE> in the source.

#### 4.3.3.1 Description of Mandatory and Optional Parameters

The following table describes the parameters used in the data string for option 43.

**Table 22 Data String Parameters for Option 43**

Parameter	Description	Field Type
<TargetIP>	Replace it with a valid IPv4 address in dotted decimal notation.	Mandatory
<TargetTCPPort>	Replace it with a decimal number ranging from 1 to 65535 (inclusive). The default TCP port 3260 is assumed, if not specified.	Optional
<LUN>	It is the hexadecimal representation of the Logical Unit Number of the boot device. If not provided, LUN 0 is assumed to be the boot LUN. It is an eight-byte number which should be specified as a hexadecimal number consisting of 16 digits, with an appropriate number of 0's padded to the left, if required.	Optional
<TargetName>	Replace it with a valid iSCSI target IQN name of up to 223 characters.	Mandatory
<InitiatorName>	Replace it with a valid iSCSI IQN name of up to 223 characters. If not provided the default Initiator name (generated by the adapter based on the board's MAC address) will be used.	Optional
<HeaderDigest>	Replace it with either E or D. <ul style="list-style-type: none"> <li>■ E denotes that the header digest is enabled.</li> <li>■ D denotes that it is disabled.</li> </ul>	Optional
<DataDigest>	Replace it with either E or D. <ul style="list-style-type: none"> <li>■ E denotes that the data digest is enabled.</li> <li>■ D denotes that it is disabled.</li> </ul> If not provided it is assumed that Data Digest is disabled by default.	Optional
<AuthenticationType>	If applicable replace it with D, E, or M. <ul style="list-style-type: none"> <li>■ D denotes that authentication is disabled.</li> <li>■ E denotes that one-way CHAP is enabled - the username and secret to be used for oneway CHAP must be specified by non-DHCP means</li> <li>■ M denotes that MutualCHAP is enabled - user name and passwords required for mutual CHAP authentication must be specified by non-DHCP means.</li> </ul> If not specified, this field defaults to authentication disabled.	Optional

##### 4.3.3.1.1 Examples

The following is an example of the Default Initiator name and Data Digest Settings:

iscsi:"192.168.0.2":3261:"000000000000000E":"iqn.2009-4.com:1234567890"::"E"::"E"

where the values are:

- Target IP address: 192.168.0.2
- Target TCP port: 3261

- 
- Target boot LUN: 0x0E
  - Target IQN name: iqn.2009-04.com:1234567890
  - Initiator name: Not specified. Use the Initiator name already configured. Use the default name if none was configured.
  - Header Digest: Enabled
  - Data digest: Not specified. Assume disabled.
  - Authentication Type: 1-way CHAP.

The following is an example of Default TCP Port and Mutual CHAP Settings:

```
iscsi :"192.168.0.2" :: "000000000000000E" :"iqn.2009-4.com:1234567890" :: "E" :"D" :"M"
```

where the values are:

- Target IP address: 192.168.0.2
- Target TCP port: Use default from RFC 3720 (3260)
- Target boot LUN: 0x0E
- Target IQN name: iqn.2009-04.com:1234567890
- Initiator name: Not specified. Use the Initiator name already configured. Use the default name if none was configured.
- Header Digest: Enabled
- Data digest: Data Digest disabled
- Authentication Type: Mutual CHAP

## 4.4 ESXi 5.5 and ESXi 6.0 Native Mode NIC Driver Troubleshooting Support

**Table 23** lists the troubleshooting support provided by the earlier vmklinux driver in terms of the proc interface and the equivalent support provided by the native mode NIC driver through either the esxcli plug-in or the VmkMgmtKeyVal interface.

**Table 23 ESXi 5.5 and ESXi 6.0 Native Mode NIC Driver Troubleshooting Support**

Troubleshooting Items from Earlier vmklinux Driver using the proc Interface	Description	ESXi 5.5 and ESXi 6.0 Native Mode Driver Support		Usage in ESXi 5.5 and ESXi 6.0
		(VmkMgmtKeyVal interface)	(esxcli plug-in)	
csr_read	Read the 32-bit register value from the CSR space at the offset set through the csr_read file.	X	—	<pre>#/usr/lib/vmware/vmkmgmt_keyval/vmk mgmt_keyval -i "vmnic0-KeyValue/emulex" -k "CsrRead" -s "&lt;offset = 0x2d90&gt;" #/usr/lib/vmware/vmkmgmt_keyval/vmk mgmt_keyval -i "vmnic0-KeyValue/emulex" -k "CsrRead" -g Key 'CsrRead': CSR Offset:0x2d90 ==&gt; 0x126</pre>
csr_write	Set the CSR space offset for csr_write.	X	—	<pre>#/usr/lib/vmware/vmkmgmt_keyval/vmk mgmt_keyval -i "vmnic0-KeyValue/emulex" -k "CsrWrite" -s "&lt;offset = 0x2d90&gt; &lt;value = 3&gt;"</pre>
drv_stat	Read the driver statistics.	—	X	Available in custom esxcli plug-in in the asynchronous driver:  <pre>esxcli elxnet stats get -p &lt;pci dev name&gt;</pre>
eth_ring	Get the various Ethernet ring properties.	—	—	Not implemented.
misc_stat	Read the miscellaneous counters, format them with description. These counters are clear on read.	—	Partially implemented	Error stats available in a custom esxcli plug-in in the asynchronous driver:  <pre>esxcli elxnet stats get -p &lt;pci dev name&gt;</pre>
pci_read	Read the 32-bit register value from the PCI space at the offset set through the pci_read file.	X	—	<pre>#/usr/lib/vmware/vmkmgmt_keyval/vmk mgmt_keyval -i "vmnic0-KeyValue/emulex" -k "PciRead" -s "&lt;offset = 0x1f0&gt;" #/usr/lib/vmware/vmkmgmt_keyval/vmk mgmt_keyval -i "vmnic0-KeyValue/emulex" -k "PciRead" -g</pre>

**Table 23 ESXi 5.5 and ESXi 6.0 Native Mode NIC Driver Troubleshooting Support (Continued)**

Troubleshooting Items from Earlier vmklinux Driver using the proc Interface	Description	ESXi 5.5 and ESXi 6.0 Native Mode Driver Support		Usage in ESXi 5.5 and ESXi 6.0
		(VmkMgmtKeyVal interface)	(esxcli plug-in)	
pci_write	Set the PCI space offset for pci_write.	X	—	#/usr/lib/vmware/vmkmgmt_keyval/vmkmgmt_keyval -i "vmnic0-KeyValue/emulex" -k "PciWrite" -s "<offset>"
port_stat	Read the counters for port 0 and port 1 and formats them with descriptions. These counters are clear-on-read.	—	Partially implemented	From vsish node /net/pNics/vmnicXX/stats. Also available in a custom esxcli plug-in in the asynchronous driver: esxcli elxnet stats get -p <pci dev name>
vlan_stat	Get the driver-specific VLAN status.	—	Not implemented	Not relevant since the driver does not use VLAN filtering.

## 4.5 ESXi 5.5 and ESXi 6.0 Native Mode NIC Driver Support for ethtool Commands

Table 24 lists the ethtool support provided by the earlier vmklinux driver and the equivalent support provided through the esxcli and vsish commands.

**Table 24 ESXi 5.5 and 6.0 Native Mode NIC Driver Support for ethtool Commands**

ethtool Command	Description	vmklinux Driver Support (VmkMgmtKeyVal interface)	ESXi 5.5 and 6.0 Native Mode Driver Support (esxcli/ vsish command)
ethtool -a --show-pause DEVNAME	Show pause options.	X	✓ (esxcli network nic get -n vmnic0)
ethtool -A --pause DEVNAME	Set pause options.	X	✓ To enable RX/TX pause: vsish -e set /net/pNics/vmnic1/firmware/pauseParams < autonegsupport> <txpauseenabled> <rxpauseenabled> Enable: vsish -e set /net/pNics/vmnic1/firmware/pauseParams 0 1 1 Disable: vsish -e set /net/pNics/vmnic1/firmware/pauseParams 0 0 0
ethtool -c --show-coalesce DEVNAME	Show coalesce options.	X	✓ (esxcli network nic coalesce get -n vmnicX)
ethtool -C --coalesce DEVNAME	Set coalesce options.	X	✓ (esxcli network nic coalesce set -n vmnicX)

**Table 24 ESXi 5.5 and 6.0 Native Mode NIC Driver Support for ethtool Commands (Continued)**

ethtool Command	Description	vmklinux Driver Support (VmklMgmtKeyVal interface)	ESXi 5.5 and 6.0 Native Mode Driver Support (esxcli/ vsish command)
ethtool -g --show-ring DEVNAME	Query RX/TX ring parameters.	X	esxcli does not support it
ethtool -G --set-ring DEVNAME	Set RX/TX ring parameters.	Not implemented	esxcli does not support it
ethtool -k --show-offload DEVNAME	Get protocol offload information.	X	X (esxcli network nic sg/tso/cso get)
ethtool -K --offload DEVNAME	Set protocol offload.	Not implemented	Not implemented (esxcli network nic sg/tso/cso set -n vmnicX) Supported using: vsish:vsish -e get /net/pNics/vmic<>/hwCapabilities vsish -e set /net/pNics/vmnic0/hwCapabilities/<CAP> <1/0>
ethtool -i --driver DEVNAME	Show driver information.	✓	✓ (esxcli network nic info get)
ethtool -d --register-dump DEVNAME	Dump device registers.	✓	✓ (esxcli elxnet regdump get -p <pcidevname> -f <filepath>)
ethtool -e --eprom-dump DEVNAME	Dump device EEPROM.	✓	Not implemented (esxcli network nic eeprom dump -n vmnicX)
ethtool -E --change-eeprom DEVNAME	Change bytes in device EEPROM.	Not implemented	Not implemented (esxcli network nic eeprom change -n vmnicX)
ethtool -r --negotiate DEVNAME	Restart N-WAY negotiation.	Not implemented	Not implemented (esxcli network nic negotiate restart -n vmnicX)
ethtool -p --identify DEVNAME	Show visible port identification (for example, blinking).	✓	esxcli does not support it
ethtool -t --test DEVNAME	Execute adapter self test.	Not implemented	Not implemented (esxcli network nic selftest run -n vmnicX)
ethtool -S --statistics DEVNAME	Show adapter statistics.	✓	✓ (esxcli network nic stats get) The standard esxcli statistics only expose a limited set of statistic counters. More extensive statistics are available through the esxcli plug-in available in the asynchronous driver: esxcli elxnet stats get -p <pci dev name>

# Appendix A: Configuring iSCSI Through DHCP

This section provides information for configuring the iSCSI driver using DHCP.

## A.1 Recommendations

If you are using a DHCP server to obtain an IP address for your iSCSI initiator, set up a reservation. A reservation assigns a specific IP address based on the MAC address of your iSCSI function. If you do not reserve an IP address through DHCP, then you must set the lease length for the iSCSI initiator IP address to unlimited. This prevents the IP address lease from expiring.

### A.1.1 Vendor-Specific Option 43

The following section describes the format for the data returned in DHCP vendor-specific option 43. The method and format for specifying the Vendor ID is outside the scope of this document and is not included here. The initiator offers this Vendor ID to the DHCP server to retrieve data in the format described in the following section.

#### A.1.1.1 Format of Vendor-Specific Option 43

The data returned in DHCP vendor-specific option 43 uses the following format:

```
iscsi : "<TargetIP>" : "<TargetTCPPort>" : "<LUN>" : "<TargetName>" : "<InitiatorName>" :  
"HeaderDigest" : "DataDigest" : "AuthenticationType"
```

The guidelines for creating the data string include:

- Strings shown in quotation marks are part of the syntax and are therefore mandatory.
- Fields enclosed in angular brackets (including the angular brackets) must be replaced with their corresponding values. Some of these fields are optional and may be skipped.
- If an optional field is skipped, a colon must be used as a placeholder to indicate the default value for that field.
- If specified, the value of each parameter must be enclosed in double quotes.

All options are case sensitive.

For more information on the Option 43 values, see [Section 4.3.3.1, Description of Mandatory and Optional Parameters](#).

