Documentation for Emulex® LightPulse® Adapters

User Guide

Version 11.4
September 11, 2017

Corporate Headquarters
San Jose, CA

Website
www.broadcom.com

Table of Contents

Emulex® Software Kit Migration

User Guide

Version 11.2
December 30, 2016

pub-005841

Corporate Headquarters
San Jose, CA

Website
www.broadcom.com

Chapter 1: Introduction

Beginning with software release 11.2, Emulex® LightPulse® adapters and OneConnect® adapters have independent software kits.

This document provides special instructions and considerations for using the driver kits for LightPulse and OneConnect adapters.

Special cases include those in which pre-11.2 (original) drivers and applications are replaced by the new 11.2 drivers and applications, and cases in which inbox drivers are replaced by the new 11.2 out-of-box (OOB) drivers.

NOTES

- Detailed instructions for installing LightPulse kits are provided in the driver and applications user guide for LightPulse adapters.
- Detailed instructions for installing OneConnect kits are provided in the driver and applications user guide for OneConnect adapters.

The installation instructions in the user guides apply to "fresh" installations of the LightPulse or OneConnect drivers, in which the pre-11.2 drivers were never installed. They also apply going forward, after the pre-11.2 drivers have been replaced by either LightPulse or OneConnect drivers.

1.1 Supported Adapters

LightPulse software kits supports the following adapters:

- LPe12000 series (FC protocol)
- LPe16202/OCe15100 adapters (FC or NIC+FCoE protocol, depending on the SFPs installed)
- LPe16000 series (FC protocol)
- LPe31000 series (FC protocol)
- LPe32000 series (FC protocol)

OneConnect software kits supports the following adapters:

- OCe14000 series (NIC, iSCSI, and FCoE protocols)
- OCe11000 series (NIC, iSCSI, and FCoE protocols)

1.2 Considerations

If you have both LightPulse and OneConnect adapters in your system and you are upgrading to 11.2, you must install both kits before rebooting your system.

If you are downgrading from 11.2, you must first uninstall the 11.2 drivers and applications before rebooting your system.

Existing driver configurations are maintained during the upgrade to 11.2.

1.2.1 Inbox Drivers

- Mixing Linux or ESXi inbox drivers with out-of-box drivers is not a supported configuration. If inbox drivers are used, make sure to use inbox drivers for all protocols on all LightPulse and OneConnect adapters in the system.
- Management applications will work with the inbox drivers.

1.2.2 Management Applications

- Install LightPulse management applications from the LightPulse software kit to manage LightPulse adapters.
- Install OneConnect management applications from the OneConnect software kit to manage OneConnect adapters.
- Both sets of management applications can run on the same system at the same time.
- Management applications prior to the 11.2 release cannot be run with the 11.2 LightPulse or OneConnect management tools on the same system. For instructions on reverting applications to a version earlier than 11.2, see Chapter 3 Reverting Drivers and Applications to Previous Versions.

1.2.3 Firmware

A firmware update to version 11.2 is required for all OCe14000-series adapters, OCe11000-series adapters, and LPe16202/OCe15100 adapters (in NIC+FCoE mode) in the system for UEFI boot code to operate properly.

1.3 Driver Kit Designations and Application Kit Names

The following table provides the original and new designations of the driver kits for each operating system.

Table 1 Driver Kit Designations

Kit Names	Windows	Linux	ESXi	Solaris
Original FC kit	elxdrvr-fc	elx-lpfc	lpfc	N/A
New LightPulse FC kit	elxdrvr-fc	elx-lpfc	lpfc	N/A
New OneConnect FC kit	N/A	N/A	N/A	N/A
Original FCoE kit	elxdrvr-fcoe	elx-lpfc	lpfc	elxfc
New LightPulse FCoE kit	elxdrvr-fcoe	elx-lpfc (same kit for FC)	lpfc	lpfcoe
New OneConnect FCoE kit	brcmdrvr-fcoe	brcmfcoe	brcmfcoe	elxfc
Original NIC kit	elxdrvr-nic	be2net	elxnet	elxnic
New LightPulse NIC kit	elxdrvr-nic	lpnic	lpnic	lpnic
New OneConnect NIC kit	brcmdrvr-nic	be2net	elxnet	elxnic
Original iSCSI kit	elxdrvr-iscsi	be2iscsi	be2iscsi	N/A
New LightPulse iSCSI kit	N/A	N/A	N/A	N/A
New OneConnect iSCSI kit	brcmdrvr-iscsi	be2iscsi	be2iscsi	N/A

The following table provides the new names of the management applications:

Table 2 Application Kit Names

LightPulse Name	OneConnect Name	Description
OCManager	BrcmOCManager	Management GUI
HbaCmd	BrcmHbaCmd	Management CLI
OneCapture™	BrcmOneCapture	Dump capture
Elxflash	brcmflash	Firmware download
linlpcfg	brcmlinlpcfg	Linux management tool
winlpcfg	brcmwinlpcfg	Windows management tool
elxcli-fc elxlpfccli lpniccli lpniccli	elxcli-fcoe brcmfcoecli elxcli elxnetcli	esxcli for ESXi 5.5 FC or FCoE esxcli for ESXi 6.0/6.5 FC or FCoE esxcli for ESXi 5.5 NIC esxcli for ESXi 6.0/6.5 NIC
ELXOCM-VMware-vCenter	BRCMOCM-VMware-vCenter	Management GUI and CLI for VMware vCenter
esx-provider-emulex	esx-provider-brcm	CIM Provider

1.4 Additional Information

Appendix A provides information about original and new names of driver files, which are embedded in the driver kits.

Appendix B provides illustrations that show the 11.2 kit names, driver names, operating systems, and applications associated with the LightPulse and OneConnect software kits.

Appendix C contains special instructions on situations that might arise during installation.

Chapter 2: Installing Drivers

This section provides instructions for installing and uninstalling the drivers in special situations.

2.1 Installing Split Out-of-Box Drivers over Inbox Drivers

2.1.1 Windows

There are no special procedures needed to install the new out-of-box drivers over existing inbox drivers on Windows. Follow the installation instructions in the appropriate manual:

- For LightPulse adapters, follow the instructions in the *Emulex Drivers for Windows for LightPulse Adapters User Guide*, located on the Broadcom website at www.broadcom.com.
- For OneConnect adapters, follow the instructions in the *Emulex Drivers for Windows for OneConnect Adapters User Guide*, located on the Broadcom website at www.broadcom.com.

2.1.2 Linux

2.1.2.1 LightPulse FC and FCoE Driver with Only LightPulse Adapters Installed

NOTE The existing inbox lpfc driver cannot be removed; it can only be superseded by an out-of-box driver. When a new out-of-box lpfc driver is installed, it is placed in either the `updates` folder or the `extras` folder. The Linux kernel then checks both of these folders. If an lpfc driver is found, that driver is loaded into memory. Otherwise, the inbox lpfc driver is loaded.

To install the LightPulse FC and FCoE driver with only LightPulse adapters installed in the system, do the following:

1. Issue the `tar xvf elx-lpfc-dd-<Linux distribution version>-<driver version>.tar.gz` command.
2. Do one of the following:
   — Run the `elx_lpfc_install.sh` installation script.
   — Issue the `rpm -ivh elx-lpfc-<driver version>_<Linux distribution version>.rpm` command.

2.1.2.2 OneConnect FCoE Driver with Only OneConnect Adapters Installed

To install the OneConnect FCoE driver with only OneConnect adapters installed in the system, do the following:

1. Issue the `tar xvf brcmfcoe-dd-<Linux distribution version><driver version>.tar.gz` command.
2. Do one of the following:
   — Run the `brcmfcoe_install.sh` installation script.
   — Issue the `rpm -ivh brcmfcoe-<driver version>_<Linux distribution version>.rpm` command.

The OneConnect FCoE driver, brcmfcoe, attaches to OneConnect adapters only. Because the inbox lpfc driver also supports these adapters, the OneConnect out-of-box FCoE driver cannot load if the inbox lpfc driver has already been loaded. Therefore, following the steps above will cause the LightPulse driver to be blacklisted to prevent it from loading.

2.1.2.3 LightPulse FC and FCoE Driver and OneConnect FCoE Driver with a Combination of LightPulse and OneConnect Adapters Installed

To install both the LightPulse FC and FCoE driver and the OneConnect FCoE driver on a system that has both LightPulse and OneConnect adapters installed, perform the following steps:

1. For the LightPulse driver:
   a. Issue the `tar xvf elx-lpfc-dd-<Linux distribution version>-<driver version>.tar.gz` command.
   b. Do one of the following:
      - Run the `elx_lpfc_install.sh` installation script.
      - Issue the `rpm -ivh elx-lpfc-<driver version>_<Linux distribution version>.rpm` command.
2. For the OneConnect driver:
   a. Issue the `tar xvf brcmfcoe-dd-<Linux distribution version><driver version>.tar.gz` command.
   b. Do one of the following:
      - Run the `brcmfcoe_install.sh` installation script.
      - Issue the `rpm -ivh brcmfcoe-<driver version>_<Linux distribution version>.rpm` command.

      NOTE During the brcmfcoe driver installation, if LightPulse adapters are detected, the lpfc driver is not blacklisted. Instead, a message prompts you to install the out-of-box 11.2 or later lpfc driver to ensure the operation of both LightPulse and OneConnect adapters on the system.
3. Reboot the system.

2.1.2.4 LightPulse NIC Driver with Only LightPulse Adapters Installed

NOTE The existing inbox be2net driver cannot be removed; it can only be superseded by an out-of-box driver. When a new out-of-box be2net driver is installed, it is placed in either the `updates` folder or the `extras` folder. The Linux kernel then checks both of these folders. If a be2net driver is found, that driver is loaded into memory. Otherwise, the inbox be2net driver is loaded.

To install the LightPulse NIC driver with only LightPulse adapters installed in the system, do the following:

1. Issue the `tar xvf elx-lpnic-dd-<Linux distribution version>-<driver version>.tar.gz` command.
2. Do one of the following:
   — Run the `elx_lpnic_install.sh` installation script.
   — Issue the `rpm -ivh elx-lpnic-<driver version>_<Linux distribution version>.rpm` command.

The LightPulse NIC driver, lpnic, attaches to LightPulse adapters only. Because the inbox be2net driver also supports these adapters, the OneConnect out-of-box be2net driver cannot load if the inbox be2net driver has already been loaded. Therefore, following the steps above will cause the OneConnect driver to be blacklisted to prevent it from loading.

2.1.2.5 OneConnect NIC Driver with Only OneConnect Adapters Installed

To install the OneConnect NIC driver with only OneConnect adapters installed in the system, do the following:

1. Issue the `tar xvf elx-be2net-dd-<Linux distribution version><driver version>.tar.gz` command.
2. Do one of the following:
   — Run the `elx_be2net_install.sh` installation script.
   — Issue the `rpm -ivh elx-be2net-<driver version>_<Linux distribution version>.rpm` command.

2.1.2.6 LightPulse NIC Driver and OneConnect NIC Driver with a Combination of LightPulse and OneConnect Adapters Installed

To install both the LightPulse NIC driver and the OneConnect NIC driver on a system that has both LightPulse and OneConnect adapters installed, perform the following steps:

1. For the OneConnect driver:
   a. Issue the `tar xvf elx-be2net-dd-<Linux distribution version><driver version>.tar.gz` command.
   b. Do one of the following:
      - Run the `elx_be2net_install.sh` installation script.
      - Issue the `rpm -ivh elx-be2net-<driver version>_<Linux distribution version>.rpm` command.
2. For the LightPulse driver:
   a. Issue the `tar xvf elx-lpnic-dd-<Linux distribution version>-<driver version>.tar.gz` command.
   b. Do one of the following:
      - Run the `elx_lpnic_install.sh` installation script.
      - Issue the `rpm -ivh elx-lpnic-<driver version>_<Linux distribution version>.rpm` command.

      NOTE During the lpnic driver installation, if OneConnect adapters are detected, the be2net driver is not blacklisted. Instead, a message prompts you to install the out-of-box 11.2 or later be2net driver to ensure the operation of both LightPulse and OneConnect adapters on the system.
3. Reboot the system.

2.1.3 VMware Drivers

2.1.3.1 LightPulse FC and FCoE Driver with Only LightPulse Adapters Installed

This driver maintains the lpfc driver binary name so it will update the existing inbox driver.

To install this driver, type:

```
esxcli software vib install -d <path to driver kit>/Emulex-FCoE-FC-lpfc-<driver version>-offline-bundle-<VMware#>.zip

reboot
```

When the server has rebooted, check the vmhba enumeration to driver binary as follows:

```
esxcli storage core adapter list
```

For special instructions, see Appendix C: Section C.2, VMware.

2.1.3.2 OneConnect FCoE Driver with Only OneConnect Adapters Installed

This driver introduces the brcmfcoe driver binary name. This driver only supports OneConnect adapters. Because this driver has a different binary name, it will not affect the existing inbox driver on installation.

To install this driver, first remove the lpfc driver, and then install the brcmfcoe driver. This prevents lpfc from claiming its known PCI IDs in the server.

1. Remove the lpfc driver by typing:

   ```
   esxcli software vib remove –n lpfc
   ```

2. Install the brcmfcoe driver by typing:

   ```
   esxcli software vib install –d <path to driver kit>/brcm-FCoE-brcmfcoe-<driver version>-offline-bundle-<VMware#>.zip


   reboot
   ```

3. When the server has rebooted, check the vmhba enumeration to driver binary as follows:

   ```
   esxcli storage core adapter list
   ```

For special instructions, see Appendix C: Section C.2, VMware.

2.1.3.3 LightPulse FC and FCoE Driver and OneConnect FCoE Driver with a Combination of LightPulse and OneConnect Adapters Installed

If your server configuration has both LightPulse and OneConnect adapters, you need to install both of the LightPulse and OneConnect OOB drivers. Installing both drivers provides complete coverage for all adapter families.

NOTE Because your configuration is using LightPulse and OneConnect adapters, the inbox lpfc driver is updated with an OOB LightPulse driver to acquire the LightPulse adapters, and then the OOB OneConnect driver is added to acquire the OneConnect adapters.

1. To install these drivers, type the following:

   ```
   esxcli software vib install –d <path to driver kit>/Emulex-FCoE-FC-lpfc-<driver version>-offline-bundle-<VMware#>.zip
   esxcli software vib install –d <path to driver kit>/brcm-FCoE-brcmfcoe-<driver version>-offline-bundle-<VMware#>.zip
   ```

2. When the server has rebooted, check the vmhba enumeration to driver binary as follows:

   ```
   esxcli storage core adapter list
   ```

For special instructions, see Appendix C: Section C.2, VMware.

2.1.3.4 LightPulse NIC Driver with Only LightPulse Adapters Installed

This driver introduces the lpnic driver binary name. This driver only supports LightPulse adapters. Because this driver has a different binary name, it will not affect the existing inbox driver on installation.

To install this driver, first remove the elxnet driver, and then install the lpnic driver. This prevents elxnet from claiming its known PCI IDs in the server.

1. Remove the elxnet driver by typing:

   ```
   esxcli software vib remove –n elxnet
   ```

2. Install the lpnic driver by typing:

   ```
   esxcli software vib install –d <path to driver kit>/esx-<OS version>-lpnic-<driver version>.zip
   ```

```
reboot
```

3. When the server has rebooted, check the vmhba enumeration to driver binary as follows:

```
esxcli storage core adapter list
```

For special instructions, see Appendix C: Section C.2, VMware.

2.1.3.5 OneConnect NIC Driver with Only OneConnect Adapters Installed

This driver maintains the elxnet driver binary name so it will update the existing inbox driver.

To install this driver, type:

```
esxcli software vib install –d <path to driver kit>/esx-<OS
version>-elxnet-<driver version>.zip


reboot
```

When the server has rebooted, check the vmhba enumeration to driver binary as follows:

```
esxcli storage core adapter list
```

For special instructions, see Appendix C: Section C.2, VMware.

2.1.3.6 LightPulse NIC Driver and OneConnect NIC Driver with a Combination of LightPulse and OneConnect Adapters

If your server configuration has both LightPulse and OneConnect adapters, you need to install both of the LightPulse and OneConnect OOB drivers. Installing both drivers provides complete coverage for all adapter families.

NOTE Because your configuration is using LightPulse and OneConnect adapters, the inbox elxnet driver is updated with an OOB OneConnect driver to acquire the OneConnect adapters, and then the OOB LightPulse driver is added to acquire the LightPulse adapters.

1. To install these drivers, type the following:

```
esxcli software vib install –d <path to driver kit>/esx-<OS
version>-elxnet-<driver version>.zip
esxcli software vib install –d <path to driver kit>/esx-<OS
version>-lpnic-<driver version>.zip
```

2. When the server has rebooted, check the vmhba enumeration to driver binary as follows:

```
esxcli storage core adapter list
```

For special instructions, see Appendix C: Section C.2, VMware.

2.2 Installing Split Out-of-Box Drivers over Existing 11.1 or Earlier Out-of-Box Drivers

2.2.1 Windows

There are no special procedures needed to install the new out-of-box drivers over existing out-of-box drivers on Windows. Follow the installation instructions in the appropriate manual:

- For LightPulse adapters, follow the instructions in the *Emulex Drivers for Windows for LightPulse Adapters User Guide*, located on the Broadcom website at www.broadcom.com.

- For OneConnect adapters, follow the instructions in the *Emulex Drivers for Windows for OneConnect Adapters User Guide*, located on the Broadcom website at www.broadcom.com.

2.2.2 Linux

2.2.2.1 LightPulse FC and FCoE Driver with Only LightPulse Adapters Installed

NOTE An existing out-of-box lpfc driver cannot be removed; it can only be superseded by another out-of-box driver. When a new out-of-box lpfc driver is installed, it is placed in the `updates` or the `extras` folder. The Linux kernel then checks both of these folders. If a new lpfc driver is found, that driver is loaded into memory. Otherwise, the older lpfc driver is loaded.

To install the LightPulse FC and FCoE driver with only LightPulse adapters installed in the system, do the following:

1. Issue the `tar xvf elx-lpfc-dd-<Linux distribution version>-<driver version>.tar.gz` command.
2. Do one of the following:
   — Run the `elx_lpfc_install.sh` installation script.
   — Issue the `rpm -ivh elx-lpfc-<driver version>_<Linux distribution version>.rpm` command.

2.2.2.2 OneConnect FCoE Driver with Only OneConnect Adapters Installed

To install the OneConnect FCoE driver with only OneConnect adapters installed in the system, do the following:

1. Issue the `tar xvf brcmfcoe-dd-<Linux distribution version><driver version>.tar.gz` command.
2. Do one of the following:
   — Run the `brcmfcoe_install.sh` installation script.
   — Issue the `rpm -ivh brcmfcoe-<driver version>_<Linux distribution version>.rpm` command.

The OneConnect FCoE driver, brcmfcoe, attaches to OneConnect adapters only. Because the pre-11.2 lpfc driver also supports these adapters, the 11.2 or later OneConnect out-of-box FCoE driver cannot load if the pre-11.2 lpfc driver has already been loaded. Therefore, following the steps above will cause the 11.2 or later LightPulse driver to be blacklisted to prevent it from loading.

2.2.2.3 LightPulse FC and FCoE Driver and OneConnect FCoE Driver with a Combination of LightPulse and OneConnect Adapters Installed

To install both the LightPulse FC and FCoE driver and the OneConnect FCoE driver on a system that has both LightPulse and OneConnect adapters installed, perform the following steps:

1. For the LightPulse driver:
   a. Issue the `tar xvf elx-lpfc-dd-<Linux distribution version>-<driver version>.tar.gz` command.
   b. Do one of the following:
      - Run the `elx_lpfc_install.sh` installation script.
      - Issue the `rpm -ivh elx-lpfc-<driver version>_<Linux distribution version>.rpm` command.

2. For the OneConnect driver:
   a. Issue the `tar xvf brcmfcoe-dd-<Linux distribution version><driver version>.tar.gz` command.
   b. Do one of the following:
      - Run the `brcmfcoe_install.sh` installation script.
      - Issue the `rpm -ivh brcmfcoe-<driver version>_<Linux distribution version>.rpm` command.

      NOTE During the 11.2 or later brcmfcoe driver installation, if LightPulse adapters are detected, the pre-11.2 lpfc driver is not blacklisted. Instead, a message prompts you to install the out-of-box 11.2 or later lpfc driver to ensure the operation of both LightPulse and OneConnect adapters on the system.
3. Reboot the system.

2.2.2.4 LightPulse NIC Driver with Only LightPulse Adapters Installed

NOTE An existing out-of-box be2net driver cannot be removed; it can only be superseded by an out-of-box driver. When a new out-of-box be2net driver is installed, it is placed in either the `updates` folder or the `extras` folder. The Linux kernel then checks both of these folders. If a be2net driver is found, that driver is loaded into memory. Otherwise, the inbox be2net driver is loaded.

To install the LightPulse NIC driver with only LightPulse adapters installed in the system, do the following:

1. Issue the `tar xvf elx-lpnic-dd-<Linux distribution version>-<driver version>.tar.gz` command.
2. Do one of the following:
   — Run the `elx_lpnic_install.sh` installation script.
   — Issue the `rpm -ivh elx-lpnic-<driver version>_<Linux distribution version>.rpm` command.

The LightPulse NIC driver, lpnic, attaches to LightPulse adapters only. Because the pre-11.2 be2net driver also supports these adapters, the OneConnect out-of-box 11.2 or later be2net driver cannot load if the pre-11.2 be2net driver has already been loaded. Therefore, following the steps above will cause the OneConnect out-of-box 11.2 or later driver to be blacklisted to prevent it from loading.

2.2.2.5 OneConnect NIC Driver with Only OneConnect Adapters Installed

To install the OneConnect NIC driver with only OneConnect adapters installed in the system, do the following:

1. Issue the `tar xvf elx-be2net-dd-<Linux distribution version><driver version>.tar.gz` command.
2. Do one of the following:
   — Run the `elx_be2net_install.sh` installation script.
   — Issue the `rpm -ivh elx-be2net-<driver version>_<Linux distribution version>.rpm` command.

2.2.2.6 LightPulse NIC Driver and OneConnect NIC Driver with a Combination of LightPulse and OneConnect Adapters Installed

To install both the LightPulse NIC driver and the OneConnect NIC driver on a system that has both LightPulse and OneConnect adapters installed, perform the following steps:

1. For the OneConnect driver:
   a. Issue the `tar xvf elx-be2net-dd-<Linux distribution version><driver version>.tar.gz` command.
   b. Do one of the following:
      - Run the `elx_be2net_install.sh` installation script.
      - Issue the `rpm -ivh elx-be2net-<driver version>_<Linux distribution version>.rpm` command.
2. For the LightPulse driver:
   a. Issue the `tar xvf elx-lpnic-dd-<Linux distribution version>-<driver version>.tar.gz` command.
   b. Do one of the following:
      - Run the `elx_lpnic_install.sh` installation script.
      - Issue the `rpm -ivh elx-lpnic-<driver version>_<Linux distribution version>.rpm` command.

      NOTE During the 11.2 or later lpnic driver installation, if OneConnect adapters are detected, the pre-11.2 out-of-box be2net driver is not blacklisted. Instead, a message prompts you to install the out-of-box 11.2 or later be2net driver to ensure the operation of both LightPulse and OneConnect adapters on the system.
3. Reboot the system.

2.2.3 VMware

2.2.3.1 LightPulse FC and FCoE Driver with Only LightPulse Adapters Installed

This driver maintains the lpfc driver binary name so it will update the out-of-box driver.

To install this driver, type:

```
esxcli software vib install -d <path to driver kit>/Emulex-FCoE-FC-lpfc-<driver version>-offline-bundle-<VMware#>.zip

reboot
```

When the server has rebooted, check the vmhba enumeration to driver binary as follows:

```
esxcli storage core adapter list
```

For special instructions, see Appendix C: Section C.2, VMware.

2.2.3.2 OneConnect FCoE Driver with Only OneConnect Adapters Installed

This driver introduces the brcmfcoe driver binary name. This driver only supports OneConnect adapters. Because this driver has a different binary name, it will not affect the existing out-of-box driver on installation.

To install this driver, first remove the lpfc driver, and then install the brcmfcoe driver. This prevents lpfc from claiming its known PCI IDs in the server.

1. Remove the lpfc driver by typing:

```
esxcli software vib remove –n lpfc
```

2. Install the brcmfcoe driver by typing:

```
esxcli software vib install –d <path to driver kit>/brcm-FCoE-brcmfcoe-<driver
version>-offline-bundle-<VMware#>.zip

reboot
```

3. When the server has rebooted, check the vmhba enumeration to driver binary as follows:

```
esxcli storage core adapter list
```

For special instructions, see Appendix C: Section C.2, VMware.

2.2.3.3 LightPulse FC and FCoE Driver and OneConnect FCoE Driver with a Combination of LightPulse and OneConnect Adapters Installed

If your server configuration has both LightPulse and OneConnect adapters, you need to install both of the LightPulse and OneConnect OOB drivers. Installing both drivers provides complete coverage for all adapter families.

NOTE Because your configuration is using LightPulse and OneConnect adapters, the existing out-of-box lpfc driver is updated with an out-of-box LightPulse driver to acquire the LightPulse adapters, and then the out-of-box OneConnect driver is added to acquire the OneConnect adapters.

1. To install these drivers, type the following:

```
esxcli software vib install –d <path to driver
kit>/Emulex-FCoE-FC-lpfc-<driver version>-offline-bundle-<VMware#>.zip
esxcli software vib install –d <path to driver kit>/brcm-FCoE-brcmfcoe-<driver
version>-offline-bundle-<VMware#>.zip

reboot
```

2. When the server has rebooted, check the vmhba enumeration to driver binary as follows:

```
esxcli storage core adapter list
```

For special instructions, see Appendix C: Section C.2, VMware.

2.2.3.4 LightPulse NIC Driver with Only LightPulse Adapters Installed

This driver introduces the lpnic driver binary name. This driver only supports LightPulse adapters. Because this driver has a different binary name, it will not affect the existing out-of-box driver on installation.

To install this driver, first remove the elxnet driver, and then install the lpnic driver. This prevents elxnet from claiming its known PCI IDs in the server.

1. Remove the elxnet driver by typing:

```
esxcli software vib remove –n elxnet
```

2. Install the lpnic driver by typing:

```
esxcli software vib install -d <path to driver kit>/esx-<OS
version>-lpnic-<driver version>.zip

reboot
```

3. When the server has rebooted, check the vmhba enumeration to driver binary as follows:

```
esxcli storage core adapter list
```

For special instructions, see Appendix C: Section C.2, VMware.

2.2.3.5 OneConnect NIC Driver with Only OneConnect Adapters Installed

This driver maintains the elxnet driver binary name so it will update the existing out-of-box driver.

To install this driver, type:

```
esxcli software vib install -d <path to driver kit>/esx-<OS
version>-elxnet-<driver version>.zip

reboot
```

When the server has rebooted, check the vmhba enumeration to driver binary as follows:

```
esxcli storage core adapter list
```

For special instructions, see Appendix C: Section C.2, VMware.

2.2.3.6 LightPulse NIC Driver and OneConnect NIC Driver with a Combination of LightPulse and OneConnect Adapters Installed

If your server configuration has both LightPulse and OneConnect adapters, you need to install both of the LightPulse and OneConnect OOB drivers. Installing both drivers provides complete coverage for all adapter families.

NOTE Because your configuration is using LightPulse and OneConnect adapters, the existing out-of-box elxnet driver is updated with an out-of-box OneConnect driver to acquire the OneConnect adapters. Then the out-of-box LightPulse driver is added to acquire the LightPulse adapters.

1. To install these drivers, type the following:

```
esxcli software vib install -d <path to driver kit>/esx-<OS
version>-elxnet-<driver version>.zip
esxcli software vib install -d <path to driver kit>/esx-<OS
version>-lpnic-<driver version>.zip
```

2. When the server has rebooted, check the vmhba enumeration to driver binary as follows:

```
esxcli storage core adapter list
```

For special instructions, see Appendix C: Section C.2, VMware.

2.2.4 Solaris

2.2.4.1 LightPulse FCoE Driver with Only LightPulse Adapters Installed

Install the LightPulse out-of-box driver over the combined out-of-box driver.

1. Uninstall the combined out-of-box driver (elxfc) if it is present:

```
/elxfc_remove
```

2. Install the lpfcoe driver for the LightPulse adapter.

   Untar the driver kit:

   ```
   tar –xf lpfcoe_kit-<driver version>-<OS version>-s11-i386.tar
   ```

   Install the driver:

   ```
   /lpfcoe_install
   ```

2.2.4.2 OneConnect FCoE Driver with Only OneConnect Adapters Installed

Install the elxfc driver for the OneConnect adapter.

Untar the driver kit:

```
tar –xf elxfc_kit-<driver version>-<OS version>-s11-i386.tar
```

Install the driver:

```
/elxfc_install
```

2.2.4.3 LightPulse FC and FCoE Driver and OneConnect FCoE Driver with a Combination of LightPulse and OneConnect Adapters Installed

Install the LightPulse out-of-box driver over the combined out-of-box driver:

1. Uninstall the combined out-of-box driver (elxfc) if it is present:

   ```
   /elxfc_remove
   ```

2. Install the lpfcoe driver for the LightPulse adapter:
   a. Untar the driver kit:

      ```
      tar –xf lpfcoe_kit-<driver version>-<OS version>-i386.tar
      ```

   b. Install the driver:

      ```
      /lpfcoe_install
      ```

3. Install the elxfc driver for the OneConnect adapter.
   a. Untar the driver kit:

      ```
      tar –xf elxfc_kit-<driver version>-<OS version>-i386.tar
      ```

   b. Install the driver:

      ```
      /elxfc_install
      ```

2.2.4.4 LightPulse NIC Driver with Only LightPulse Adapters Installed

Install the LightPulse out-of-box driver over the combined out-of-box driver.

1. Uninstall the combined out-of-box driver (elxnic) if it is present:

   ```
   /elxnic_remove
   ```

2. Install the lpnic driver for the LightPulse adapter.

   Untar the driver kit:

   ```
   tar –xf lpnic_kit-<driver version>-<OS version>-i386.tar
   ```

   Install the driver:

   ```
   /lpnic_install
   ```

2.2.4.5 OneConnect NIC Driver with Only OneConnect Adapters Installed

Install the elxnic driver for the OneConnect adapter.

Untar the driver kit:

```
tar –xf elxnic_kit-<driver version>-<OS version>-i386.tar
```

Install the driver:

```
/elxnic_install
```

2.2.4.6 LightPulse NIC Driver and OneConnect NIC Driver with a Combination of LightPulse and OneConnect Adapters Installed

Install the LightPulse out-of-box driver over the combined out-of-box driver:

1. Uninstall the combined out-of-box driver (elxfc) if it is present:

   ```
   /elxfc_remove
   ```

2. Install the lpfcoe driver for the LightPulse adapter:
   a. Untar the driver kit:

      ```
      tar –xf lpfcoe_kit-<driver version>-<OS version>-i386.tar
      ```

   b. Install the driver:

      ```
      /lpfcoe_install
      ```

3. Install the elxfc driver for the OneConnect adapter.
   a. Untar the driver kit:

      ```
      tar –xf elxfc_kit-<driver version>-<OS version>-i386.tar
      ```

   b. Install the driver:

      ```
      /elxfc_install
      ```

Chapter 3: Reverting Drivers and Applications to Previous Versions

This section describes the downgrade process if you have 11.2 driver and OneCommand® Manager for LightPulse or OneCommand CNA Manager for OneConnect application kits installed and you want to revert to an earlier version. This applies to both LightPulse and OneConnect 11.2 kits.

NOTE You must uninstall *all* 11.2 components before installing pre-11.2 components. Mixing 11.2 and pre-11.2 components will cause unexpected results, such as adapters not being discovered by drivers and management applications.

3.1 Windows

You must uninstall all 11.2 Windows components (drivers and management applications) before installing a pre-11.2 Windows kit.

To uninstall all 11.2 Windows components, do the following:

1. Launch the Control Panel from the start menu and select Programs and Features.
2. Click on all of the kits listed below that appear in the list of installed programs.

   NOTE All of the 11.2 driver and management application kits must be removed. Any combination of pre-11.2 and 11.2 kits is not supported.

   — LightPulse driver kits:
   - Emulex/FC kit – 11.2.<version>
   - Emulex/FCoE kit – 11.2.<version>
   - Emulex/NIC kit – 11.2.<version>

   — OneConnect driver kits:
   - Broadcom/FCoE kit – 11.2.<version>
   - Broadcom/NIC kit – 11.2.<version>
   - Broadcom/iSCSI kit – 11.2.<version>

   — LightPulse OneCommand® Manager kits:
   - Emulex OneCommand Manager Enterprise 11.2.<version>
   - Emulex OneCommand Manager CLI 11.2.<version>

   — OneConnect OneCommand CNA Manager kits:
   - Broadcom OneCommand Manager Enterprise 11.2.<version>
   - Broadcom OneCommand Manager CLI 11.2.<version>
3. Select uninstall. The selected components are uninstalled.

After the uninstallation process is complete, verify that the 11.2 drivers were uninstalled:

From the Computer Management screen, open the Device Manager (see Figure 1).

Figure 1 Windows Device Manager

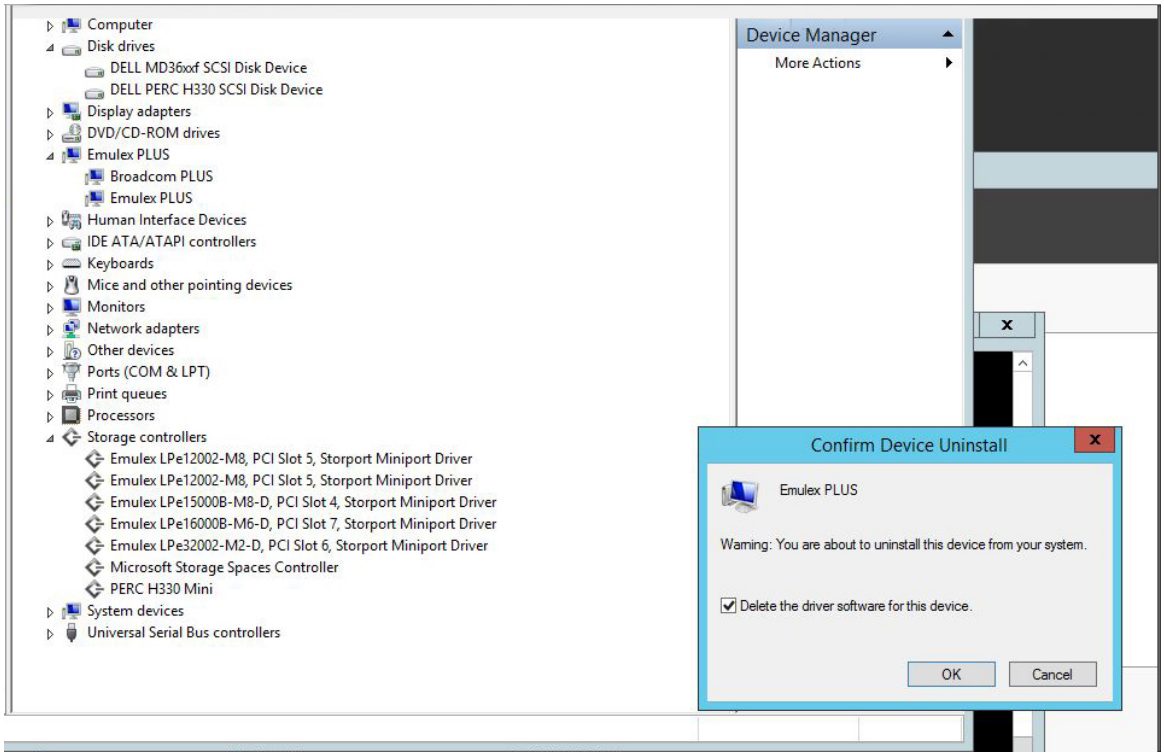

a. Click on **Storage controllers**, and then click on an adapter. A window displays the current driver version for the selected adapter. If the 11.2 kit was uninstalled correctly, the driver version shown should not be 11.2. If the 11.2 kit was not uninstalled, the driver version will still be 11.2 (see the example in Figure 2). If the driver version is 11.2, click on **Uninstall** to remove the driver.

b. Click on **Emulex PLUS**.
   - Click on Broadcom PLUS and then select an adapter. A window displays the current driver version for the selected adapter. If the driver version is 11.2, click on Uninstall to remove the driver.
   - Click on **Emulex PLUS** and then select an adapter. A window displays the current driver version for the selected adapter. If the driver version is 11.2, click on **Uninstall** to remove the driver.

c. Click on **Network adapters** and then select an adapter. A window displays the current driver version for the selected adapter. If the driver version is 11.2, click on **Uninstall** to remove the driver.

Figure 2 Windows Driver Version

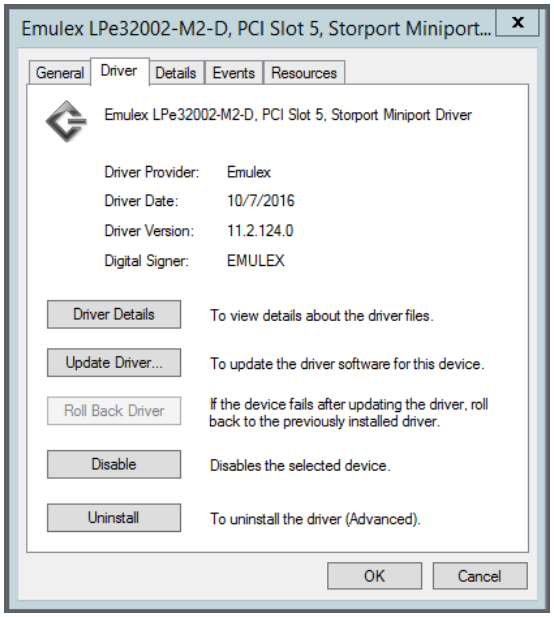

NOTE When the 11.2 drivers have been uninstalled, if you recheck an adapter, you might see an inbox driver displayed.

Verify that the OneCommand Manager kits were uninstalled:

a. From the Start menu, launch the Control Panel
b. Select Programs and Features.
c. Verify that the OneCommand Manager applications have been removed.

If the 11.2 kits were not uninstalled, perform the uninstall procedure again (see steps 1 through 3).

After all 11.2 Windows components have been uninstalled, install the pre-11.2 drivers and applications as instructed in the *Emulex Drivers for Windows User Guide,* the *OneCommand Manager Application User Guide,* and the *OneCommand Manager Command Line Interface User Guide* that accompanied the kit you are installing.

3.2 Linux

You must uninstall all 11.2 Linux components (drivers and management applications) before installing a pre-11.2 Linux kit.

To uninstall the 11.2 Linux FC and FCoE drivers for LightPulse adapters:

NOTE You must run the uninstall script that shipped with the 11.2 driver kit.

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

   ```
   ./elx_lpfc_install.sh --uninstall
   ```
4. Verify that the 11.2 drivers were uninstalled.

a. To determine whether a driver is installed, type the following at the command prompt:

```
lsmod | grep lpfc
```

If a driver is installed, output similar to the following appears:

```
lpfc  693873 0
scsi_transport_fc 64159 1 lpfc
crc tl0dif 12563 2 sd_mod,lpfc
scsi mod 240575 17
iscsi_tcp,libiscsi,scsi_transport_iscsi,st,sr_mod,usb_storage,sg,sd_m
od,scsi_dh_
alua,scsi_dh_emc,scsi_dh_rdac,scsi_dh_hp_sw,scsi_dh,lpfc,scsi_transpo
rt_fc,scsi
tgt,hpsa
```

If a driver is not installed, typing `lsmod | grep lpfc` only displays the command prompt, indicating that the uninstall was successful.

b. If a driver is installed, determine its version by typing

```
/sys/module/lpfc/version
```

Output similar to the following appears:

```
0:11.1.123.456
```

If the version is 11.2.x.x, the driver uninstall was not successful. Perform the uninstall procedure again.

To uninstall the 11.2 Linux NIC driver for LightPulse adapters:

NOTE You must run the uninstall script that shipped with the 11.2 driver kit.

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

```
./elx_lpnic_install.sh --uninstall
```

4. Verify that the 11.2 drivers were uninstalled.

a. To determine whether a driver is installed, type the following at the command prompt:

```
lsmod | grep lpnic
```

If a driver is installed, output similar to the following appears:

```
lpnic 693873 0
scsi_transport_fc 64159 1 lpnic
crc tl0dif 12563 2 sd_mod,lpfc
scsi mod 240575 17
iscsi_tcp,libiscsi,scsi_transport_iscsi,st,sr_mod,usb_storage,sg,sd_m
od,scsi_dh_
alua,scsi_dh_emc,scsi_dh_rdac,scsi_dh_hp_sw,scsi_dh,lpfc,scsi_transpo
rt_fc,scsi
tgt,hpsa
```

If a driver is not installed, typing `lsmod | grep lpnic` only displays the command prompt, indicating that the uninstall was successful.

b. If a driver is installed, determine its version by typing

```
/sys/module/lpnic/version
```

Output similar to the following appears:

```
0:11.1.123.456
```

If the version is 11.2.x.x, the driver uninstall was not successful. Perform the uninstall procedure again.

To uninstall the 11.2 Linux FCoE driver for OneConnect adapters:

NOTE You must run the uninstall script that shipped with the 11.2 driver kit.

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

   ```
   ./brcmfcoe_install.sh --uninstall
   ```
4. Verify that the 11.2 drivers were uninstalled.
   a. To determine whether a driver is installed, type the following at the command prompt:

      ```
      lsmod | grep brcmfcoe
      ```

      If a driver is installed, output similar to the following appears:

      ```
      brcmfcoe 693873 0
      scsi_transport_fc 64159 1 brcmfcoe
      crc tl0dif 12563 2 sd_mod,brcmfcoe
      scsi mod 240575 17
      iscsi_tcp,libiscsi,scsi_transport_iscsi,st,sr_mod,usb_storage,sg,sd_m
      od,scsi_dh_
      alua,scsi_dh_emc,scsi_dh_rdac,scsi_dh_hp_sw,scsi_dh,lpfc,scsi_transpo
      rt_fc,scsi
      tgt,hpsa
      ```

      If a driver is not installed, typing `lsmod | grep brcmfcoe` only displays the command prompt, indicating that the uninstall was successful.
   b. If a driver is installed, determine its version by typing

      ```
      /sys/module/brcmfcoe/version
      ```

      Output similar to the following appears:

      ```
      0:11.1.123.456
      ```

      If the version is 11.2.x.x, the driver uninstall was not successful. Perform the uninstall procedure again.

To uninstall the 11.2 Linux NIC drivers for OneConnect adapters:

NOTE You must run the uninstall script that shipped with the 11.2 driver kit.

1. Log in as root.
2. If possible, exit all applications that use Ethernet-attached drives, then unmount the drives. If you cannot exit all applications that use Ethernet-attached drives, the uninstall script works properly, but you must reboot after the uninstall is complete.
3. Run the `elx_be2net_install.sh` script with the --uninstall option:

   ```
   ./elx_be2net_install.sh --uninstall
   ```
4. Verify that the 11.2 drivers were uninstalled.
   a. To determine whether a driver is installed, type the following at the command prompt:

      ```
      lsmod | grep elx_be2net
      ```

      If a driver is installed, output similar to the following appears:

      ```
      elx_be2net 693873 0
      scsi_transport_fc 64159 1 elx_be2net
      crc tl0dif 12563 2 sd_mod,elx_be2net
      scsi mod 240575 17
      ```

```
iscsi_tcp,libiscsi,scsi_transport_iscsi,st,sr_mod,usb_storage,sg,sd_m
od,scsi_dh_
alua,scsi_dh_emc,scsi_dh_rdac,scsi_dh_hp_sw,scsi_dh,lpfc,scsi_transpo
rt_fc,scsi
tgt,hpsa
```

If a driver is not installed, typing `lsmod | grep be2net` only displays the command prompt, indicating that the uninstall was successful.

b. If a driver is installed, determine its version by typing

```
/sys/module/be2net/version
```

Output similar to the following appears:

```
0:11.1.123.456
```

If the version is 11.2.x.x, the driver uninstall was not successful. Perform the uninstall procedure again.

To uninstall the LightPulse 11.2 OneCommand Manager enterprise kit:

1. Log on as root.
2. Change to the `elxocm-<platform>-<version>` installation directory.
3. Type the following:

```
./uninstall
```

4. Verify that the 11.2 OneCommand Manager enterprise kit was uninstalled:
   a. Log on as root.
   b. Run the following command from the command prompt:

```
/usr/sbin/ocmanager/ocmanager
```

The following message is displayed if the OneCommand Manager was successfully uninstalled:

```
Bash: /usr/sbin/ocmanager/ocmanager no such a file or directory
```

To uninstall the LightPulse 11.2 OneCommand Manager core kit:

1. Log on as root.
2. Do one of the following:
   — Run the `uninstall_ocmanager.sh` script located in `/usr/sbin/hbanyware/scripts`.
   — Run the `uninstall.sh` script located in the installation tar file.
3. Verify that the 11.2 OneCommand Manager core kit was uninstalled:
   a. Log on as root.
   b. Run the following command from the command prompt:

```
/usr/sbin/ocmanager/ocmanager
```

The following message is displayed if the OneCommand Manager core kit was successfully uninstalled:

```
Bash: /usr/sbin/ocmanager/ocmanager no such a file or directory
```

To uninstall the OneConnect 11.2 OneCommand CNA Manager enterprise kit:

1. Log on as root.
2. Change to the `brcmocm-<platform>-<version>` installation directory.
3. Type the following:

```
./uninstall
```

4. Verify that the 11.2 OneCommand CNA Manager enterprise kit was uninstalled:
   a. Log on as root.
   b. Run the following command:

```
/usr/sbin/brcmocmanager/brcmocmanager
```

The following message is displayed if the OneCommand CNA Manager was successfully uninstalled:

```
Bash: /usr/sbin/brcmocmanager/brcmocmanager no such a file or
directory
```

To uninstall the OneConnect 11.2 OneCommand CNA Manager core kit:

1. Log on as root.
2. Do one of the following:
   — Run the `uninstall_brcmocmanager.sh` script located in `/usr/sbin/brcmocmanager/scripts`.
   — Run the `uninstall.sh` script located in the installation tar file.
3. Verify that the 11.2 OneCommand CNA Manager core kit was uninstalled:
   a. Log on as root.
   b. Run the following command from the command prompt:

      ```
      /usr/sbin/brcmocmanager/brcmocmanager
      ```

      The following message is displayed if the OneCommand CNA Manager core kit was successfully uninstalled:

      ```
      Bash: /usr/sbin/brcmocmanager/brcmocmanager no such a file or
      directory
      ```

After all 11.2 Linux driver and OneCommand Manager kits have been uninstalled, install the pre-11.2 drivers and applications as instructed in the *Emulex Drivers for Linux User Guide,* the *OneCommand Manager Application User Guide,* and the *OneCommand Manager Command Line Interface User Guide* that accompanied the kit you are installing.

3.3 VMware

You must uninstall all 11.2 VMware components (drivers, CIM Provider, and OneCommand Manager for VMware vCenter) before installing a pre-11.2 VMware kit.

To uninstall the 11.2 drivers, see the VMware Patch Download page for instructions.

To uninstall the CIM Provider for LightPulse adapters:

1. Type the following:

   ```
   esxcli software vib remove
   --vibname=vmware-esx-provider-emulex-cim-provider-<OS
   version>-<11.2.version>-<VMware#>.vib
   ```

To verify that the 11.2 CIM provider was uninstalled, type the following using the CLI:

```
esxcli software vib remove –n <vibname>
```

For example:

```
esxcli software vib remove –n Emulex-cim-provider
```

To uninstall the CIM provider using Windows:

From the Control Panel, click Programs and Features, right-click the CIM provider, and click **Uninstall**.

To uninstall the CIM Provider for OneConnect adapters:

Type the following:

```
esxcli software vib remove
--vibname=vmware-esx-provider-brcm-cim-provider-<OS
version>-<11.2.version>-<VMware#>.vib
```

To verify that the 11.2 CIM provider was uninstalled, type the following using the CLI:

```
esxcli software vib remove –n <vibname>
```

For example:

```
esxcli software vib remove –n brcm-cim-provider
```

To uninstall the CIM provider using Windows:

From the Control Panel, click Programs and Features, right-click the CIM provider, and click **Uninstall**.

To uninstall the OneCommand Manager for VMware vCenter plug-in for LightPulse Adapters:

1. Unregister the plug-in with the registered VMware vCenter Server:
   a. Launch the shortcut **OCM for VMware vCenter Registration** located on the desktop.
   b. Provide the details for the registered vCenter server.
   c. Click on **Unregister**.

   Repeat steps 1 to 3 if the plug-in is registered with multiple vCenter servers.
2. Uninstall the plug-in server by opening the Control Panel and uninstalling the following kit:
   — Emulex OneCommand Manager for VMware vCenter 11.2.<version> (see Figure 3)

Figure 3 Uninstalling the 11.2 OneCommand Manager for VMware vCenter

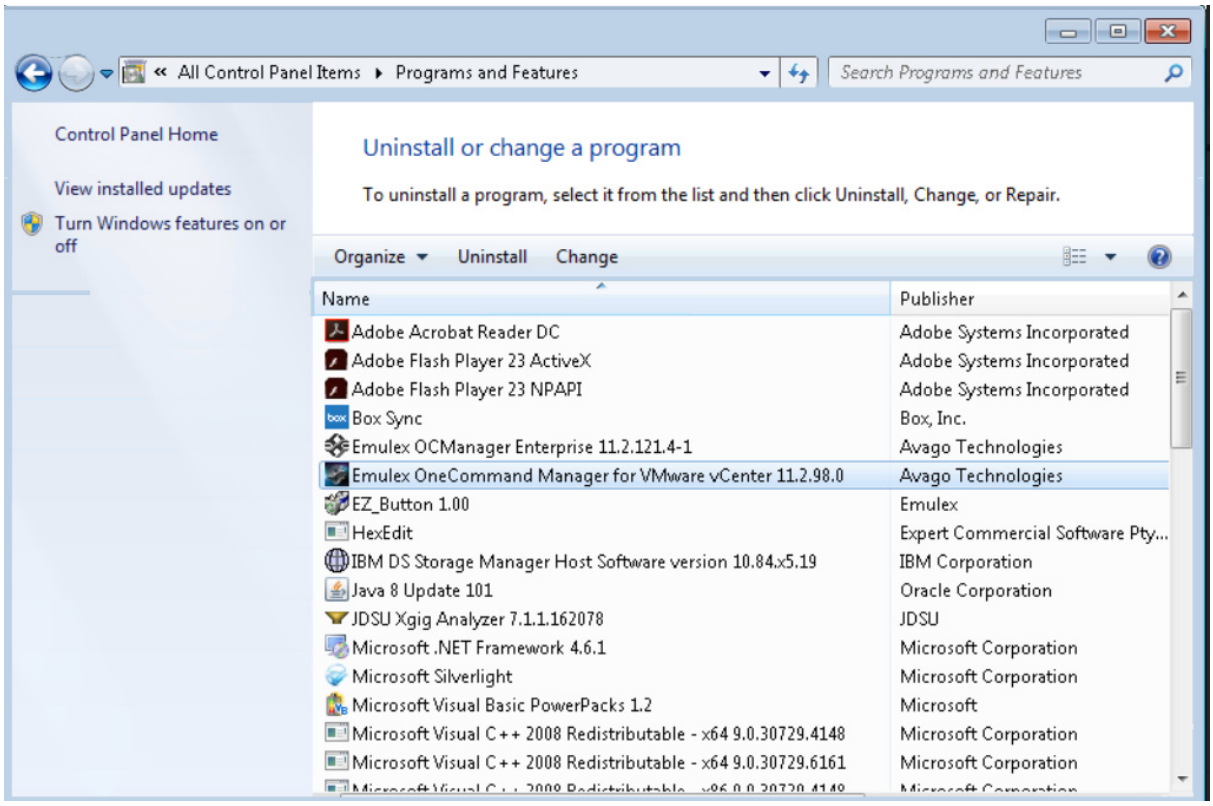

To verify that the 11.2 OneCommand Manager for VMware vCenter was uninstalled, type the following using the CLI:

```
esxcli software vib remove –n <vibname>
```

For example:

```
esxcli software vib remove –n ELXOCM-VMware-vCenter
```

To uninstall the OneCommand CNA Manager for VMware vCenter plug-in for OneConnect Adapters:

1. Unregister the plug-in with the registered VMware vCenter Server:
   a. Launch the shortcut **Broadcom OCM CNA for VMware vCenter Registration** located on the desktop.
   b. Provide the details for the registered vCenter server.
   c. Click on **Unregister**.

   Repeat steps 1 to 3 if the plug-in is registered with multiple vCenter servers.
2. Uninstall the plug-in server by opening the Control Panel and uninstalling the following kit:
   — Broadcom OneCommand CNA Manager for VMware vCenter 11.2.<version>

   To verify that the 11.2 OneCommand CNA Manager for VMware vCenter was uninstalled, type the following using the CLI:

   ```
   esxcli software vib remove –n <vibname>
   ```

   For example:

   ```
   esxcli software vib remove –n BRCMOCM-VMware-vCenter
   ```

After the 11.2 drivers, CIM Provider, and OneCommand Manager for VMware vCenter kits have been uninstalled, install the pre-11.2 drivers, CIM Provider, and OneCommand Manager for VMware vCenter kits as instructed in the *Emulex Drivers for VMware User Guide, CIM Provider Package Installation Guide,* and the *OneCommand Manager for VMware vCenter User Guide* that accompanied the kit you are installing.

3.4 Solaris

You must uninstall all 11.2 Solaris components (drivers and either OneCommand Manager for Solaris or FCA utilities) before installing a pre-11.2 Solaris kit.

To uninstall the 11.2 Solaris FCoE driver for LightPulse adapters:

1. Log in as root.
2. Remove the driver:
   — For Solaris 10: type `pkgrm lpfcoe`
   — For Solaris 11: type one of the following:
   - `./lpfcoe_remove`
   - `pkg uninstall lpfcoe`

   If the uninstall was successful, the following messages are displayed:

   ```
   Removal of <lpfcoe> was successful.
   <Cleaning directory>
   <lpfcoe_remove complete>
   <uninstall complete>
   ```
3. Reboot the system by typing

   ```
   reboot
   ```

To uninstall the 11.2 Solaris NIC driver for LightPulse adapters:

1. Log in as root.
2. Remove the driver:
   - For Solaris 10: type `pkgrm lpnic`
   - For Solaris 11: type one of the following:
     - `./lpnic_remove`
     - `pkg uninstall lpnic`

   If the uninstall was successful, the following messages are displayed:

   ```
   Removal of <lpnic> was successful.
   <Cleaning directory>
   <lpnic_remove complete>
   <uninstall complete>
   ```
3. Reboot the system by typing

   ```
   reboot
   ```

To uninstall the 11.2 Solaris FCoE driver for OneConnect adapters:

1. Log in as root.
2. Remove the driver:
   - For Solaris 10: type `pkgrm elxfc`
   - For Solaris 11: type one of the following:
     - `./elxfc_remove`
     - `pkg uninstall elxfc`

   If the uninstall was successful, the following messages are displayed:

   ```
   Removal of <elxfc> was successful.
   <Cleaning directory>
   <elxfc_remove complete>
   <uninstall complete>
   ```
3. Reboot the system by typing

   ```
   reboot
   ```

To uninstall the 11.2 Solaris NIC driver for OneConnect adapters:

1. Log in as root.
2. Remove the driver:
   - For Solaris 10: type `pkgrm elxnic`
   - For Solaris 11: type one of the following:
     - `elxnic_remove`
     - `pkg uninstall elxnic`

   If the uninstall was successful, the following messages are displayed:

   ```
   Removal of <elxnic> was successful.
   <Cleaning directory>
   <elxnic_remove complete>
   <uninstall complete>
   ```
3. Reboot the system by typing

   ```
   reboot
   ```

To uninstall the LightPulse 11.2 OneCommand Manager enterprise kit:

1. Log in as root.
2. Run the following script:

   ```
   /opt/ELXocm/scripts/uninstall
   ```

   If the uninstallation is successful, the following messages are displayed:

   ```
   Removal of <ELXocm> was successful.
   <Cleaning directory>
   <ELXocm_remove complete>
   <uninstall complete>
   ```

To uninstall the LightPulse 11.2 OneCommand Manager core kit:

1. Log on as root.
2. Do one of the following:
   - Run `/opt/ELXocm/scripts/uninstall`
   - Run the uninstall script located in the installation tar file.
   - Enter the command `pkgrm ELXocmcore`.

   If the uninstallation is successful, the following messages are displayed:

   ```
   Removal of <ELXocmcore> was successful.
   <Cleaning directory>
   <ELXocmcore_remove complete>
   <uninstall complete>
   ```

To uninstall the OneConnect 11.2 OneCommand CNA Manager enterprise kit:

1. Log on as root.
2. Run the following script:

   ```
   /opt/BRCMocm/scripts/uninstall
   ```

   If the uninstallation is successful, the following messages are displayed:

   ```
   Removal of <BRCMocm> was successful.
   <Cleaning directory>
   <BRCMocm_remove complete>
   <uninstall complete>
   ```

To uninstall the OneConnect 11.2 OneCommand CNA Manager core kit:

1. Log on as root.
2. Do one of the following:
   - Run `/opt/brcmocm/scripts/uninstall`
   - Run the uninstall script located in the installation tar file.
   - Enter the command `pkgrm brcmocmcore`.

   If the uninstallation is successful, the following messages are displayed:

   ```
   Removal of <brcmocmcore> was successful.
   <Cleaning directory>
   <brcmocmcore_remove complete>
   <uninstall complete>
   ```

To uninstall the LightPulse 11.2 emlxu kit:

1. Log on as root.
2. Do one of the following:
   — For Solaris 10: type one of the following:
      - `emlxu_remove`
      - `pkgrm` EMLXemlxu
   — For Solaris 11: type one of the following:
      - `emlxu_remove`
      - `pkg uninstall emlxu`

   If the uninstallation was successful, the following messages are displayed:

```
Removal of <EMLXemlxu> was successful.
<Cleaning directory>
<emlxu_remove complete>
<uninstall complete>
```

To uninstall the OneConnect 11.2 brcmxu kit:

1. Log on as root.
2. Do one of the following:
   — For Solaris 10: type one of the following:
      - `brcmxu_remove`
      - `pkgrm brcmxu`
   — For Solaris 11: type one of the following:
      - `brcmxu_remove`
      - `pkg uninstall brcmxu`

   If the uninstallation was successful, the following messages are displayed:

```
Removal of <brcmxu> was successful.
<Cleaning directory>
<brcmxu_remove complete>
<uninstall complete>
```

After all 11.2 Solaris drivers, OneCommand Manager kits, and FCA utilities kits have been uninstalled, install the pre-11.2 drivers and applications as instructed in the *Emulex Drivers for Solaris User Guide,* the *OneCommand Manager Application User Guide,* the *OneCommand Manager Command Line Interface User Guide,* and the *FCA Utilities for Solaris User Guide* that accompanied the kit you are installing.

Chapter 4: Technical Tips

This section provides information on situations that may be encountered during installation.

4.1 General

4.1.1 LightPulse-Specific or OneConnect-Specific Applications

Applications that are installed for LightPulse adapters will not recognize or see any OneConnect adapters if they are also installed in the system. The same is true for applications that are installed for OneConnect adapters; the application will not recognize or see any LightPulse adapters that are installed in the system.

This applies to the following applications:

- OneCommand Manager and OneCommand CNA Manager
- Elxflash Standalone Utility
- Elxflash and Brcmflash standalone utilities
- ESX CIM Provider
- OneCommand Manager for VMware vCenter and OneCommand CNA Manager for VMware vCenter

NOTE For the Elxflash Standalone Utility, when the elxflash /fc command is used in environments that include both OneConnect and LightPulse adapters, Elxflash only displays LightPulse FC adapters. OneConnect adapters will not be displayed.

When the elxflash /ucna command is used, Elxflash only displays LightPulse FCoE/NIC adapters. OneConnect adapters will not be displayed.

4.1.2 Multiple ASICs

If an adapter has multiple ASICs, and one ASIC is LightPulse and the other is OneConnect, the OneCommand Manager for LightPulse Adapters ignores the OneConnect ASIC, and the OneCommand CNA Manager for OneConnect Adapters ignores the LightPulse ASIC.

4.1.3 Reverting to a Prior Driver Version

4.1.3.1 Linux and Windows, NIC and iSCSI Protocols

If you attempt to install an 11.1 or earlier driver without first uninstalling the 11.2 or later driver, a message will direct you to uninstall the 11.2 or later driver first.

4.1.3.2 Downgrading from Version 11.0 and Higher to 10.6 is Unsuccessful when NPar and SR-IOV are Both Enabled (Dell)

If SR-IOV and nPAR are both enabled on the adapter, a firmware downgrade from 11.0 and higher to version 10.6 will be unsuccessful. Version 10.6 does not support the concurrent use of SR-IOV and nPAR.

Disable either SR-IOV or nPAR and attempt the firmware downgrade again.

4.1.3.3 Downgrading from an 11.2 Kit to a pre-11.2 Kit with the Secure Firmware Jumper Installed in the Adapter

If you attempt to downgrade 11.2 firmware to a pre-11.2 on an adapter with the secure firmware jumper installed using UEFI, the following error message is displayed:

```
Incorrect firmware file or no supported device.
```

Workaround

Remove the secure firmware jumper on the adapter to allow firmware downgrading. (To determine whether the adapter has a secure firmware jumper installed, refer to the installation guide for the adapter.)

4.1.4 Boot Code Functionality

If 11.2 boot code functionality is missing, make sure the firmware for all adapters in the system is at an 11.2 version.

4.2 Windows

4.2.1 Windows OneInstall

The OneInstall kit upgrade fails if the One Command Manager application is open and running.You must close the One Command Manager application before starting the upgrade.

4.2.2 Downgrading from a LightPulse 11.2 Kit with LPe16000-series Adapters

If you are downgrading from the LightPulse 11.2 kit and the LPe16000-series adapters BIOS is downgraded to a version prior to 10.6.193.22, the system will not boot in a boot from SAN (BFS) environment. This can be avoided if you do not downgrade to versions earlier than 10.6.193.22 in Windows, either locally or BFS.

- Windows BFS environment:

  Before booting to the operating system after a firmware downgrade, set the boot topology to point-to-point in the BIOS. The system will be able to boot, but the non-boot port will still show that the link is down.
- Windows local boot environment:

  When the system is not in a BFS environment, the adapter link will be down. However, if you are using an 8G Brocade switch, the topology can be changed to loop-only mode and the system can boot.

4.2.3 Downgrading from a LightPulse 11.2 Kit to a 11.0 Kit with LPe12000-Series Adapters

If you attempt to downgrade from an 11.2 kit to a pre-11.2 kit without first uninstalling the 11.2 kit, the following popup message is displayed:

Figure 4 Windows Security Message

Before performing the 11.2 downgrade, uninstall the 11.2 kit.

4.3 Linux

4.3.1 Unloading the lpfc and the brcmfcoe Drivers

Unloading the lpfc and the brcmfcoe drivers fail if the OneCommand Manager application is not stopped first.

- If the BrcmOCManager brcmhbamgrd is running, the LightPulse lpfc driver cannot be unloaded. Stop the BrcmOCManager brcmhbamgrd by typing the following command:

  `/usr/sbin/brcmocmanager/stop_brcmocmanager`

- If the OCManager elxhbamgrd is running, the OneConnect brcmfcoe driver cannot be unloaded. Stop the OCManager elxhbamgrd by typing the following command:

  `/usr/sbin/ocmanager/stop_ocmanager`

The lpfc and brcmfcoe drivers can now be unloaded.

4.3.2 Installing a pre-11.2 Version of OneCommand Manager

Running a pre-11.2 `ocm_install.sh` script to install a pre-11.2 version of OneCommand Manager for Linux fails.

To install the pre-11.2 OneCommand Manager, first uninstall all of the 11.2 Linux drivers. See Section 3.2 for instructions on uninstalling the drivers.

4.4 VMware

4.4.1 VMware Inbox Driver Installation over 11.2 Drivers

The pre-11.2 driver supports all FC and FCoE adapters, and the OOB LightPulse driver named **lpfc** supports only LightPulse FC and FCoE adapters. The OOB OneConnect driver named **brcmfcoe** supports only OneConnect FCoE adapters.

- Installing the pre-11.2 driver when only the OOB LightPulse driver is installed:
  - The vib upgrade command to replace the OOB LightPulse driver with the pre-11.2 driver will not work because the 11.2 drivers are always higher versions or more recent than the pre-11.2 driver.
  - Remove the existing OOB LightPulse lpfc driver before installing the pre-11.2 driver:

    ```
    esxcli software vib remove –n lpfc
    esxcli software vib install –d <path to
    offline-bundle/lpfc-offline-bundle.zip>
    ```

  - vmhba re-numbering is seen. Manually remove old entries that changed from esx.conf to vmhba numbers (see Section C.2, VMware for instructions on how to remove old entries).
- Installing the pre-11.2 driver when only an OOB OneConnect driver is installed:
  - OneConnect brcmfcoe does not share a common name with the pre-11.2 driver, so the upgrade vib command is not needed.
  - Remove the existing OOB OneConnect brcmfcoe driver before installing the pre-11.2 driver:

    ```
    esxcli software vib remove –n  brcmfcoe
    esxcli software vib install –d <path to
    offline-bundle/lpfc-offline-bundle.zip>
    ```

  - vmhba re-numbering is not usually seen if only brcmfcoe was previously installed.
- Installing the pre-11.2 driver when both the OOB LightPulse and OneConnect drivers are present:
  - Remove the existing 11.2 OneConnect brcmfcoe driver and 11.2 LightPulse driver (if any, in any order) before installing the pre-11.2 driver:

    ```
    esxcli software vib remove –n brcmfcoe
    esxcli software vib remove –n lpfc
    esxcli software vib install –d <path to
    offline-bundle/lpfc-offline-bundle.zip>
    ```

  - vmhba re-numbering is seen. Manually remove old entries that changed from esx.conf to vmhba numbers (see Section C.2, VMware for instructions on how to remove old entries).
- If the pre-11.2 driver is installed before the 11.2 drivers are removed:
  - The pre-11.2 driver will not install if LightPulse lpfc is present using either vib install or vib upgrade (the esxcli vib install command states that lpfc is already present, or there is nothing to upgrade).
  - If only brcmfcoe is present, then the pre-11.2 driver will install using the vib install command.

4.4.1.1 11.2 Drivers Installation

The server has no inbox or OOB pre-11.2 driver installed, or the pre-11.2 driver has been removed and the system has been rebooted before the 11.2 drivers are installed.

- Both of the 11.2 drivers are installed at the same time:
  - Use the vib install command to install the 11.2 drivers, and then reboot:

    ```
    esxcli software vib install –d <path to
    offline-bundle/<lpfc/brcmfcoe>-offline-bundle.zip>
    ```
  - Brcmfcoe claims the lower vmhba numbering for its adapters.
- LightPulse lpfc is installed first and the system was rebooted before the installation of brcmfcoe.
  - This can force ESXi to load lpfc first, and then brcmfcoe, so the vmhba numbering is different than when both lpfc and brcmfcoe are installed at the same time; or, when brcmfcoe is installed first and then lpfc is installed after a reboot.
  - LightPulse lpfc claims its HBAs first, so it gets lower vmhba numbering for its adapters.
- Brcmfcoe is installed, the system is rebooted, and then LightPulse lpfc is installed.
  - This can be the same situation as when both of the 11.2 drivers are installed together before a reboot.

NOTE In Boot from SAN configurations, make sure that on reboot there will be at least one driver installed that claims the adapter which is configured to boot from SAN. For OneConnect adapters, there must be at least an FCoE driver and a NIC driver on reboot for Boot from SAN setups.

4.5 Solaris

4.5.1 Running the Install Script of a pre-11.2 Kit Does Not Uninstall the Broadcom OneCommand CNA Manager

If the 11.2 OneCommand CNA Manager kits have not been removed and you run the `install.sh` of a pre-11.2 kit, the installation script detects the Emulex OneCommand CNA Manager and prompts you to uninstall the kit. However, the pre-11.2 script does not detect Broadcom OneCommand CNA Manager, and so it does not uninstall it. After the pre-11.2 kit is installed, it will only manage LightPulse adapters.

4.6 Dell Applications Downgrade

LightPulse and OneConnect DUP/MUP application packages do not support automatic downgrading from 11.2 to a prior version. If this is attempted, the following message is displayed:

```
A later version of [OCM Version] is already installed on this machine. The
setup cannot continue.
```

In some situations, when you attempt to downgrade from an 11.2.xx kit to a prior version, the pop-up message will not be displayed, and the installation of the prior version will succeed. This does not remove the prior versions, but will install the prior version kit alongside the others. Functionality of the OneCommand Manager application and the drivers being used by adapters in the system is inconsistent.

To downgrade to version 11.1 and earlier from version 11.2 and later, all kits (LightPulse and OneConnect) need to be uninstalled prior to installing the older kit version.

To uninstall 11.2.xx.xx kits for both LightPulse and OneConnect:

1. Run the `setup.exe` file that was used to install the current kit on the host.
2. The installation wizard will display an options dialog. Select the **Remove** option.
3. Repeat steps one and two for the remaining LightPulse or OneConnect kits.

NOTE The kits can also be uninstalled by using the Control Panel.

1. From the Start menu, launch the Control Panel.
2. Select Programs and Features.
3. Right-click on the kit you want to remove and select Uninstall.

After removing the LightPulse or OneConnect 11.2.xx.xx kits from the system, an older kit can be installed.

Appendix A: Driver Names

To ensure independent delivery of the LightPulse and OneConnect software kits, the following changes were made:

- Drivers that support both divisions' products have been separated into two drivers with unique names.
- The original names of LightPulse Fibre Channel drivers and utilities have been retained.
- The original names of OneConnect NIC and iSCSI drivers have been retained.

The following have been given new names:

- OneConnect utilities
- OneConnect Linux and VMware FCoE drivers
- LightPulse NIC drivers
- LightPulse Windows and Solaris FCoE drivers

The following table provides the original and new driver names for each operating system.

Table 3 Original and New Driver Names

Operating System Family	Original Name	LightPulse Name	OneConnect Name
Windows	elxplus	elxplus	brcmplus
Windows	elxfc	elxfc	n/a
Windows	elxcna	lpfcoe	elxcna
Windows	ocnd63	lpnic	ocnd63
Windows	ocnd64	lpnic	ocnd64
Windows	ocnd65	lpnic	ocnd65
Windows	be2nd62	lpnic	be2nd62
Windows	be2nd6x	lpnic	be2nd6x
Windows	be2iscsi	n/a	be2iscsi
Windows	lancutil	lancutil	n/a
Windows	lanc6ut	lanc6ut	n/a
Linux	lpfc	lpfc	brcmfcoe
Linux	be2net	lpnic	be2net
Linux	be2iscsi	n/a	be2iscsi
Linux	ocrdma	n/a	ocrdma
VMware	lpfc	lpfc	brcmfcoe
VMware	elxlpfccli	elxlpfccli	brcmfcoecli
VMware	elxnet	lpnic	elxnet
VMware	elxnetcli	lpniccli	elxnetcli
VMware	be2iscsi	n/a	be2iscsi
VMware	elxiscsi	n/a	elxiscsi
Solaris	elxfc	lpfcoe	elxfc
Solaris (inbox)	emlxs	emlxs	n/a
Solaris	elxnic	lpnic	elxnic

Appendix B: Quick Lookup Guide

B.1 Quick Lookup Guide, LightPulse and OneConnect

The following illustrations provide a guide to the split kits names, driver names, operating systems, and applications associated with the LightPulse and OneConnect kits. Figure 5 shows an overview.

Figure 5 Overview

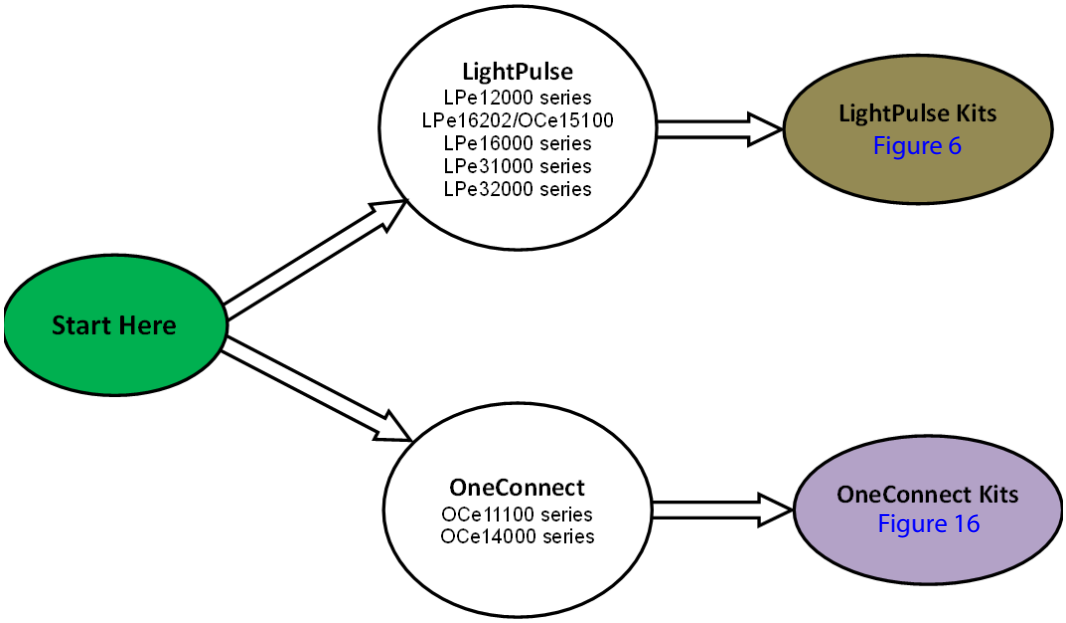

For LightPulse kits, start at flow A (Figure 6).

For OneConnect kits, start at flow K (Figure 16).

B.1.1 LightPulse Adapters and Kits

Figure 6 LightPulse Adapters and Kits

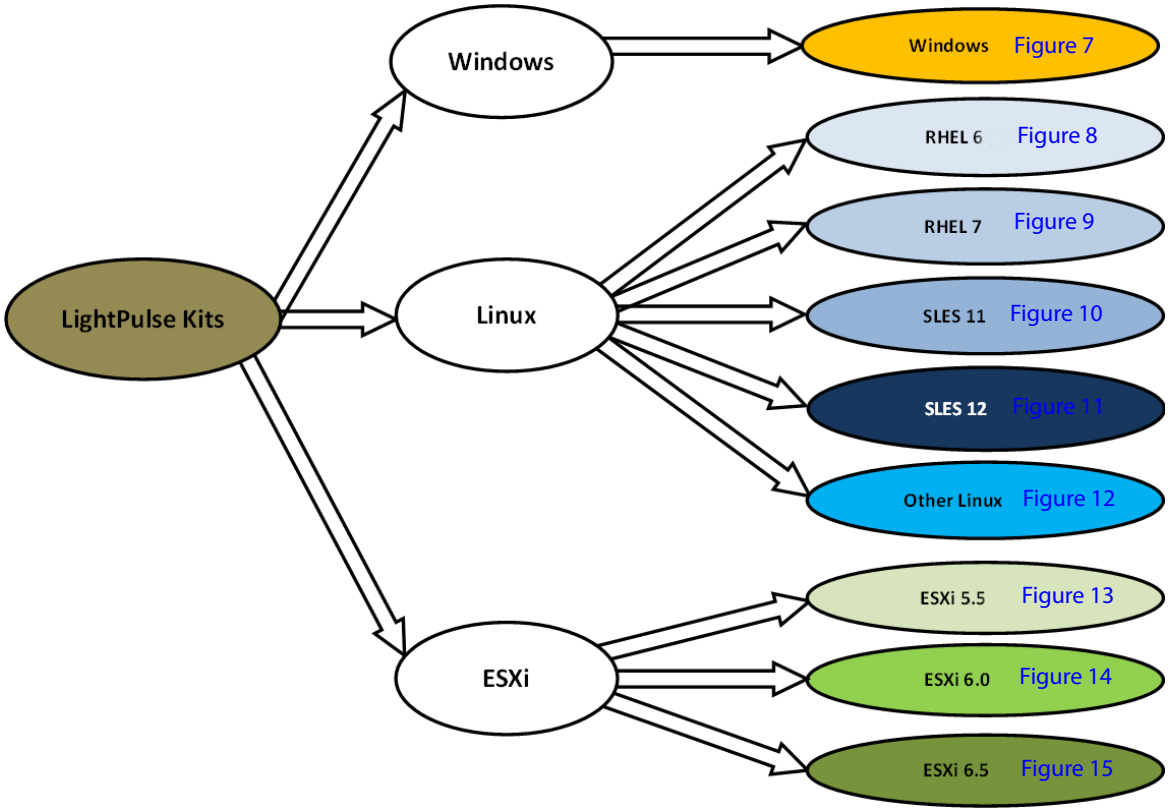

LightPulse	
Windows	Figure 7
RHEL 6	Figure 8
RHEL 7	Figure 9
SLES 11	Figure 10
SLES 12	Figure 11
Other Linux	Figure 12
ESXi 5.5	Figure 13
ESXi 6.0	Figure 14
ESXi 6.5	Figure 15

Figure 7 LightPulse Windows

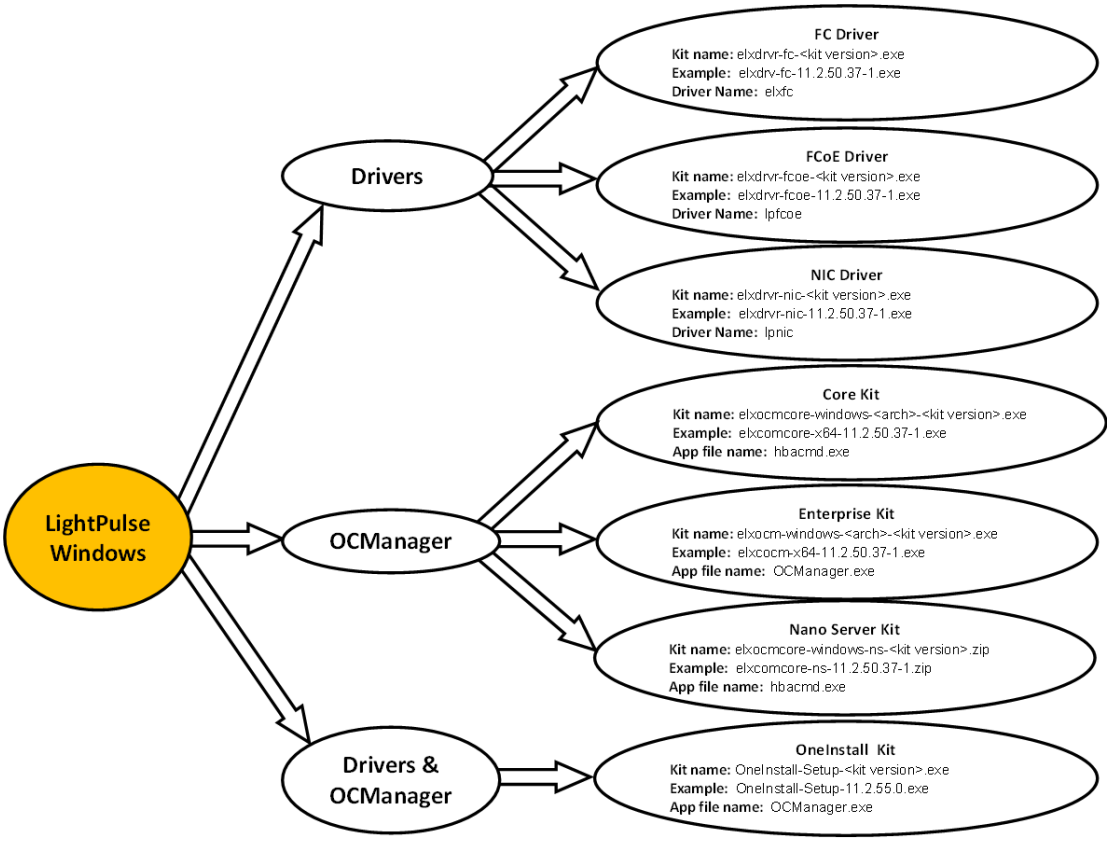

Figure 8 LightPulse Linux RHEL 6

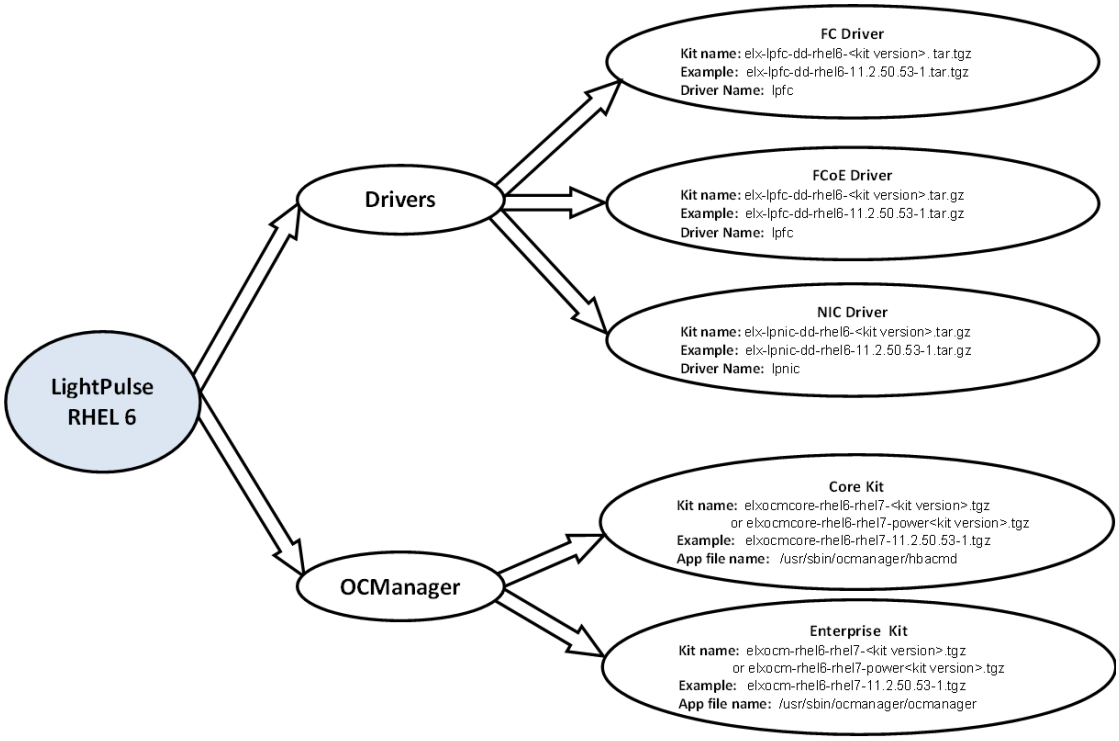

Figure 9 LightPulse RHEL 7

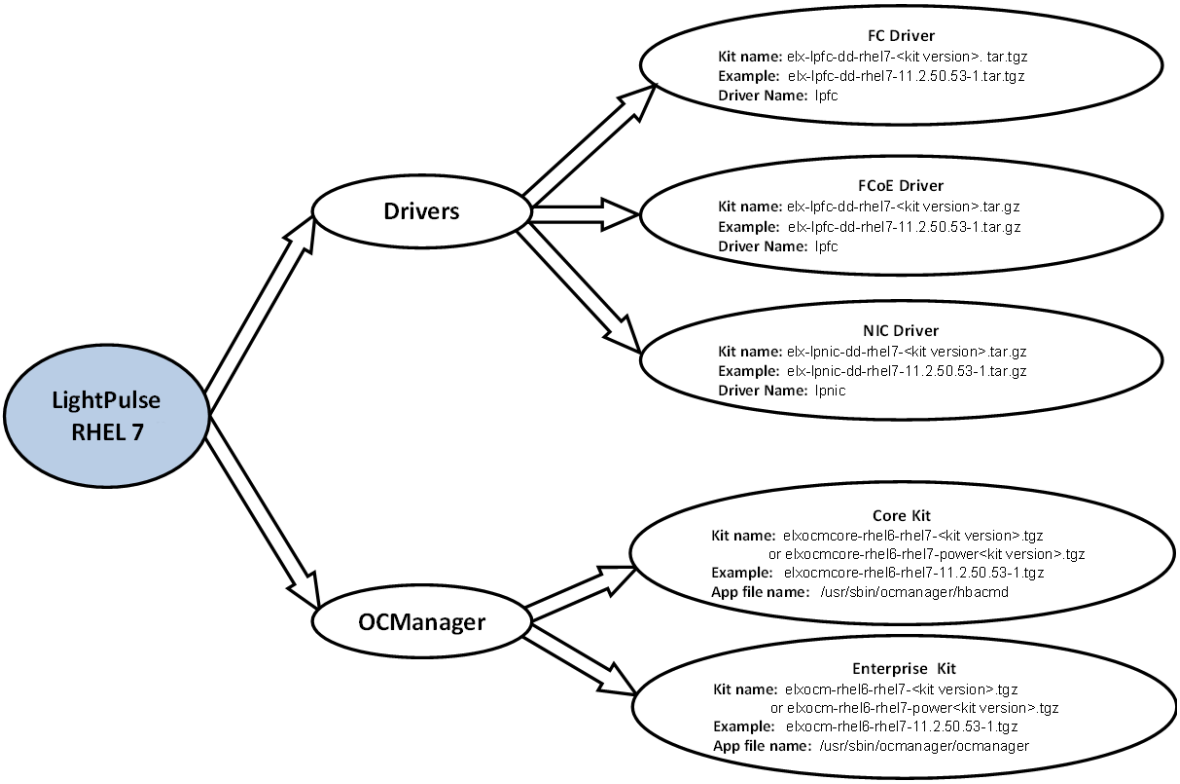

Figure 10 LightPulse SLES 11

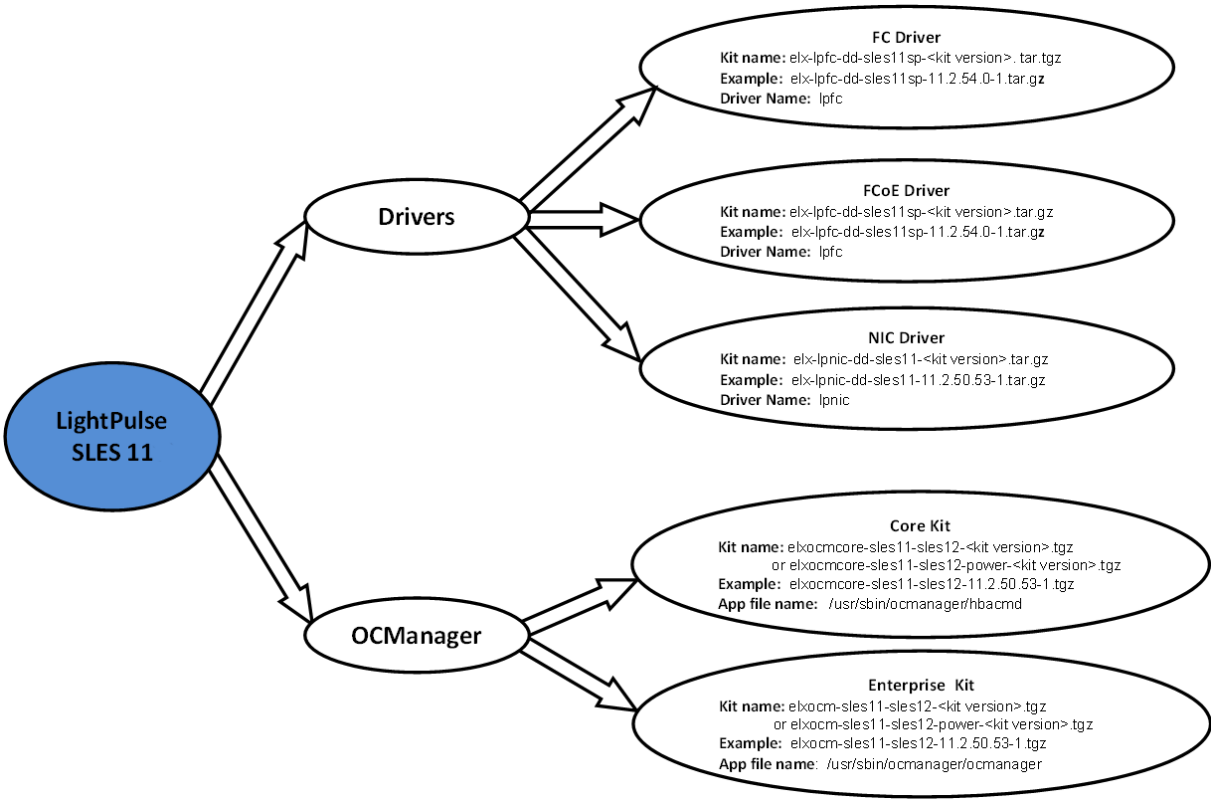

Figure 11 LightPulse SLES 12

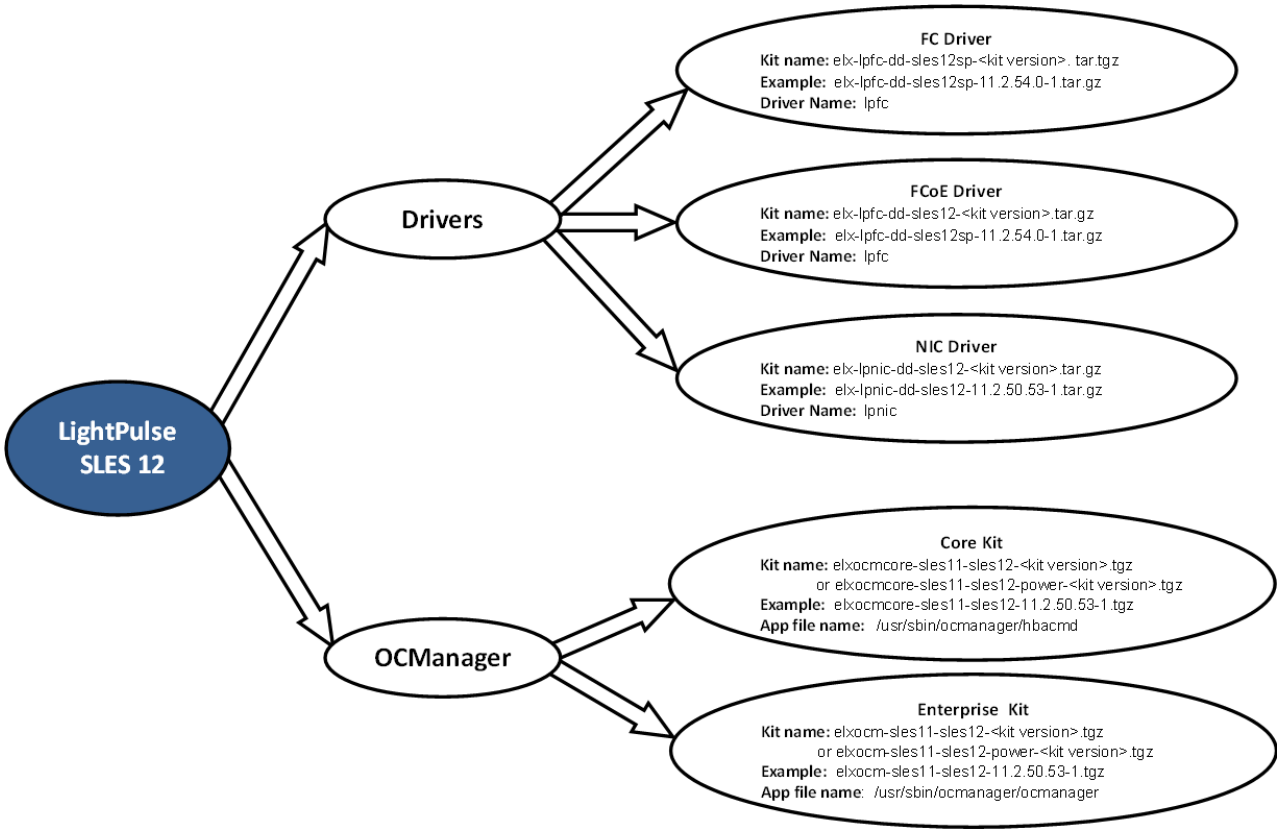

Figure 12 LightPulse Other Linux

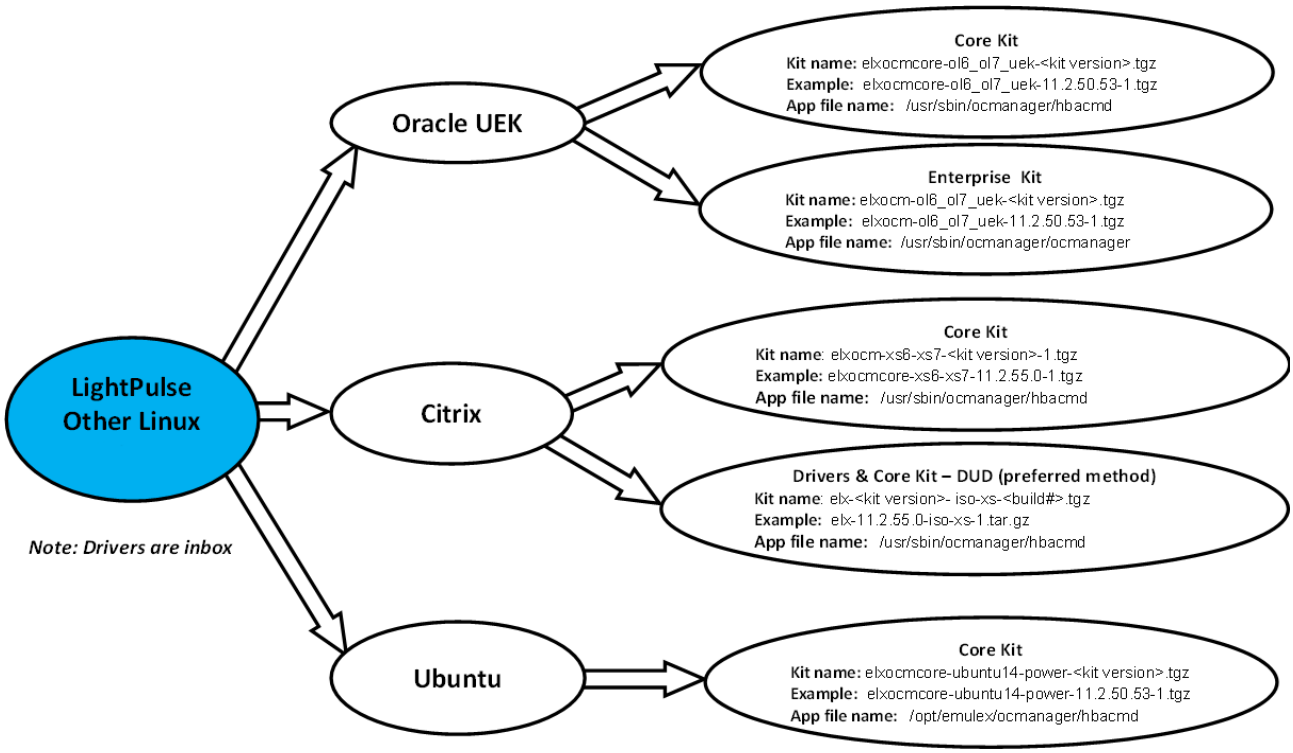

Figure 13 LightPulse ESX 5.5

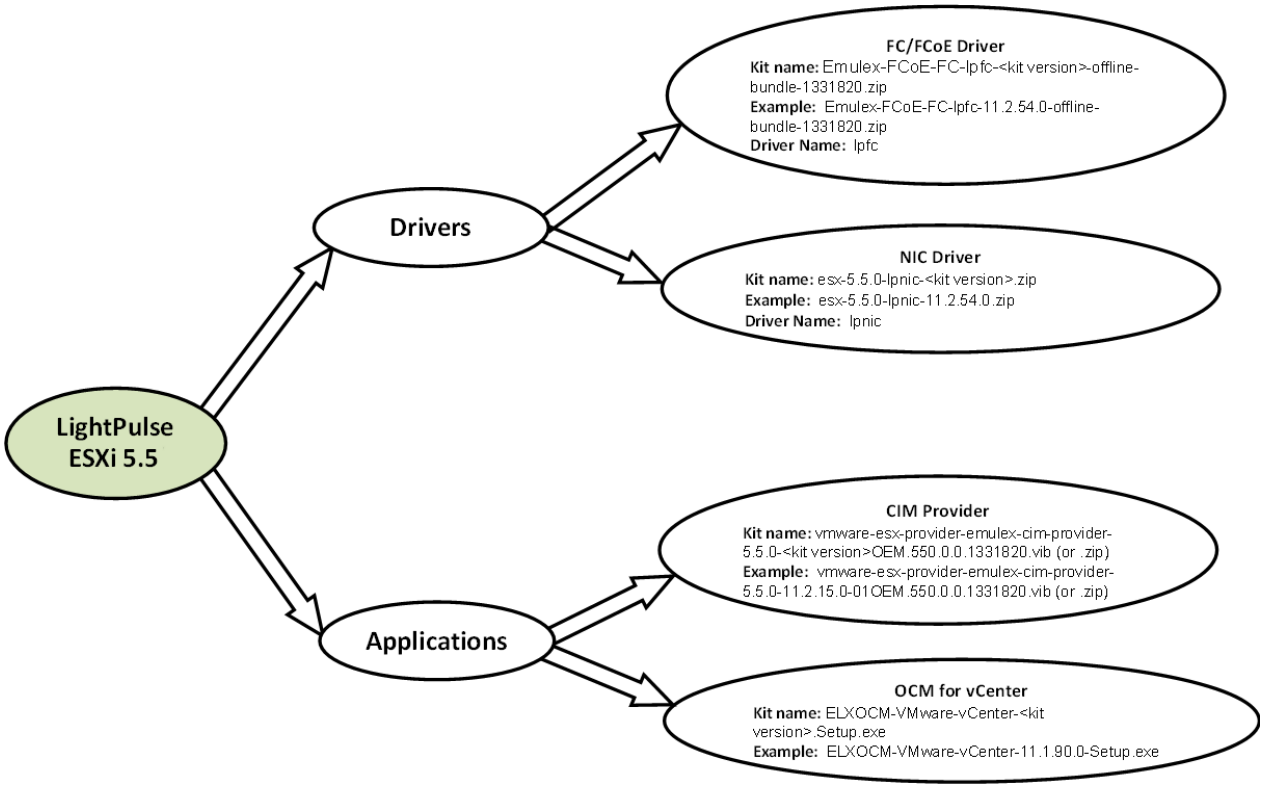

Figure 14 LightPulse ESXi 6.0

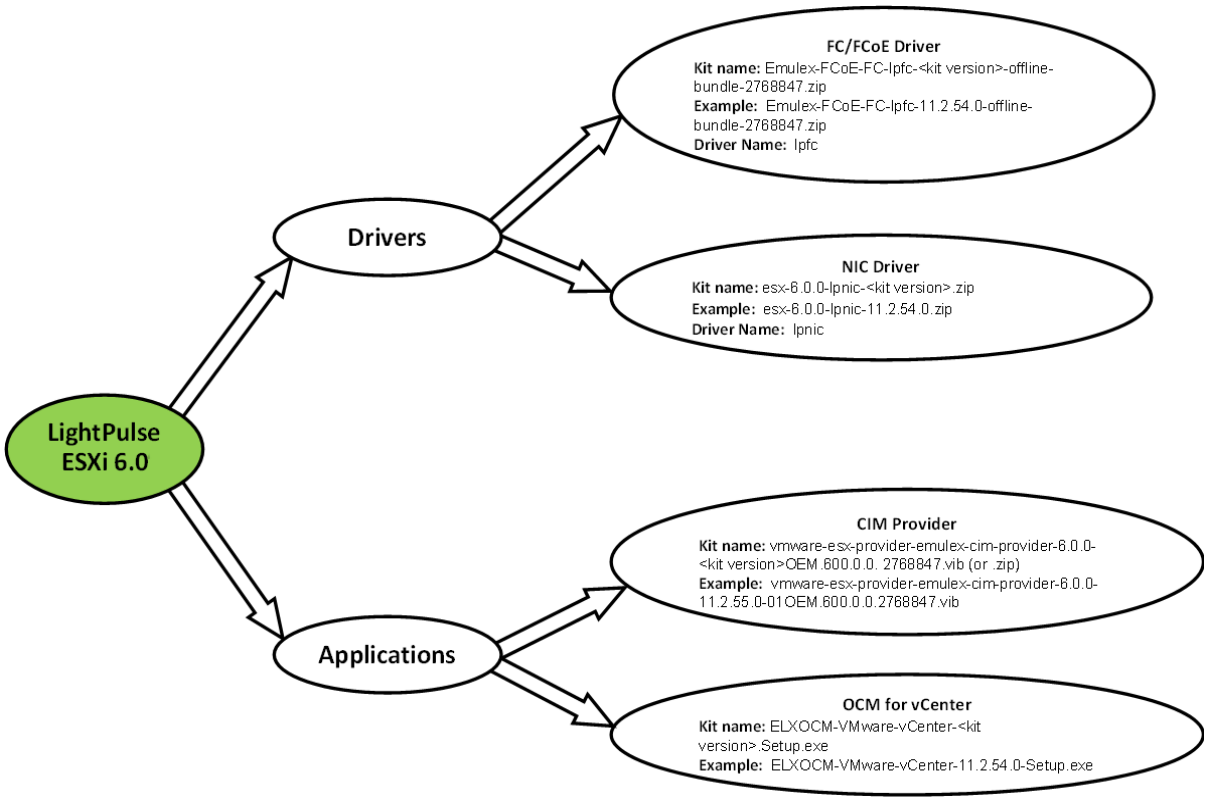

Figure 15 LightPulse ESXi 6.5

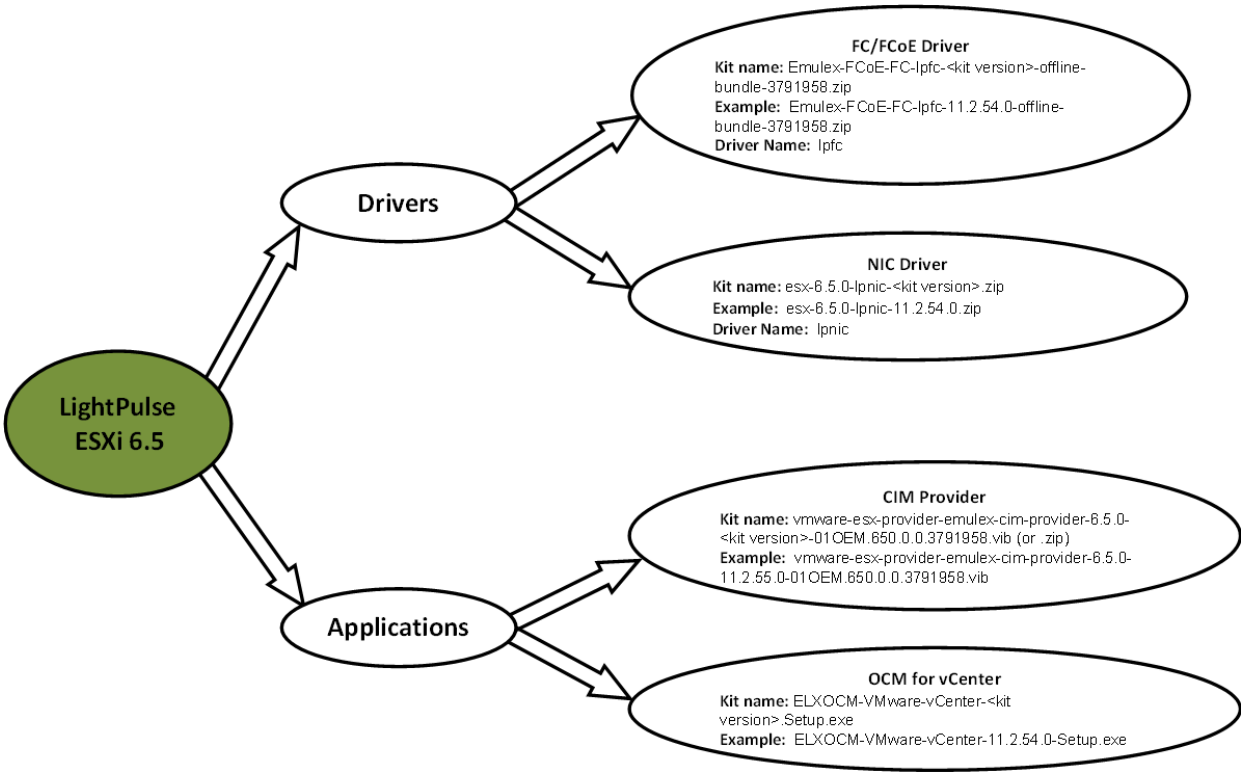

B.1.2 OneConnect Adapters and Kits

Figure 16 OneConnect Kits

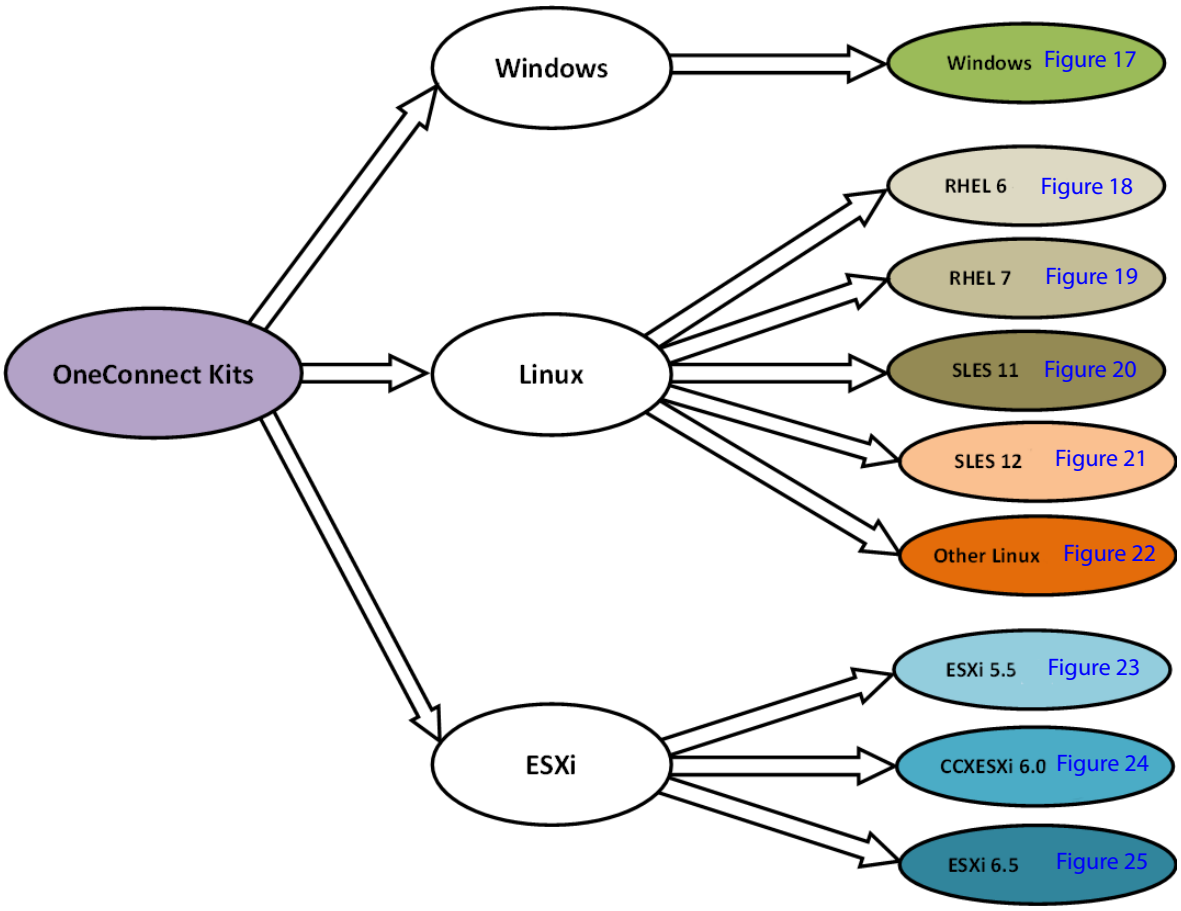

OneConnect	
Windows	Figure 17
RHEL 6	Figure 18
RHEL 7	Figure 19
SLES 11	Figure 20
SLES 12	Figure 21
Other Linux	Figure 22
ESXi 5.5	Figure 23
ESXi 6.0	Figure 24
ESXi 6.5	Figure 25

Figure 17 OneConnect Windows

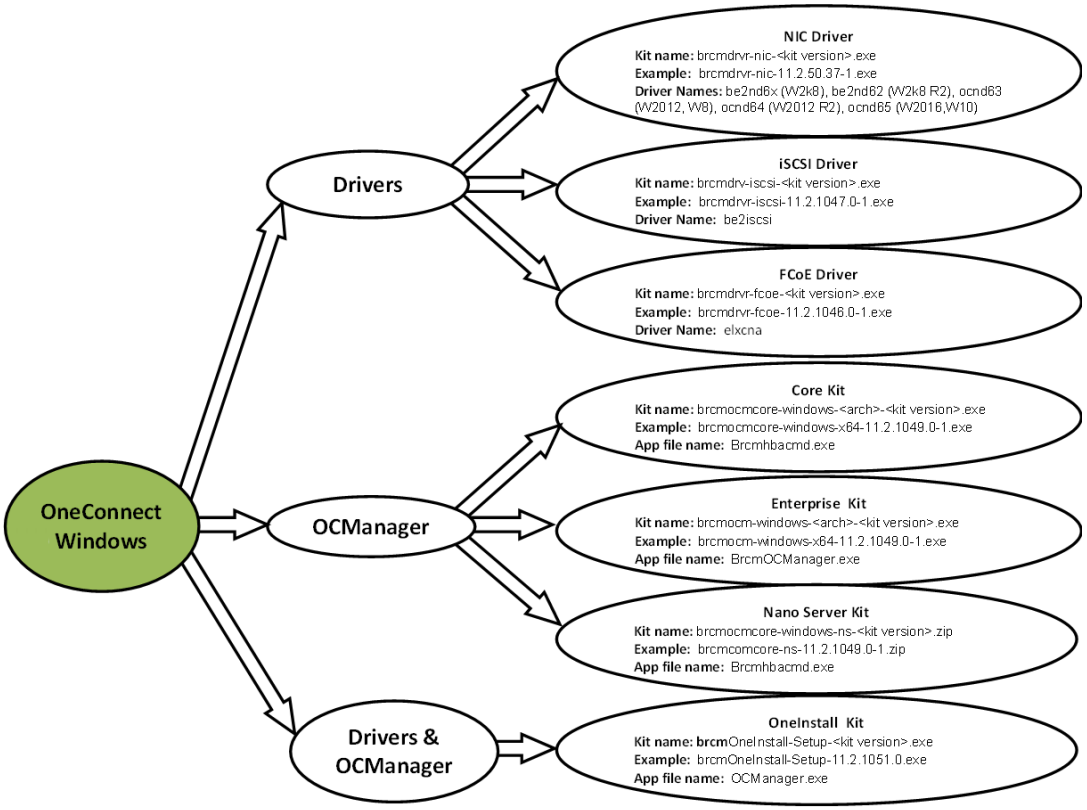

Figure 18 OneConnect RHEL 6

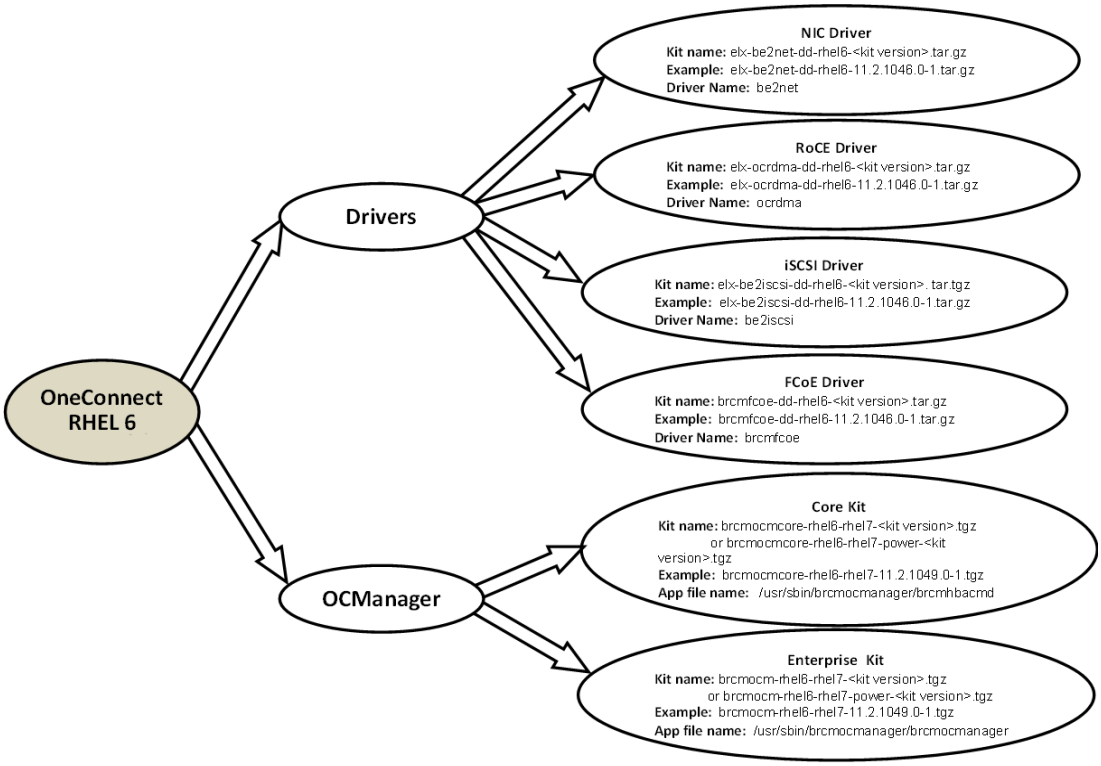

Figure 19 OneConnect RHEL 7

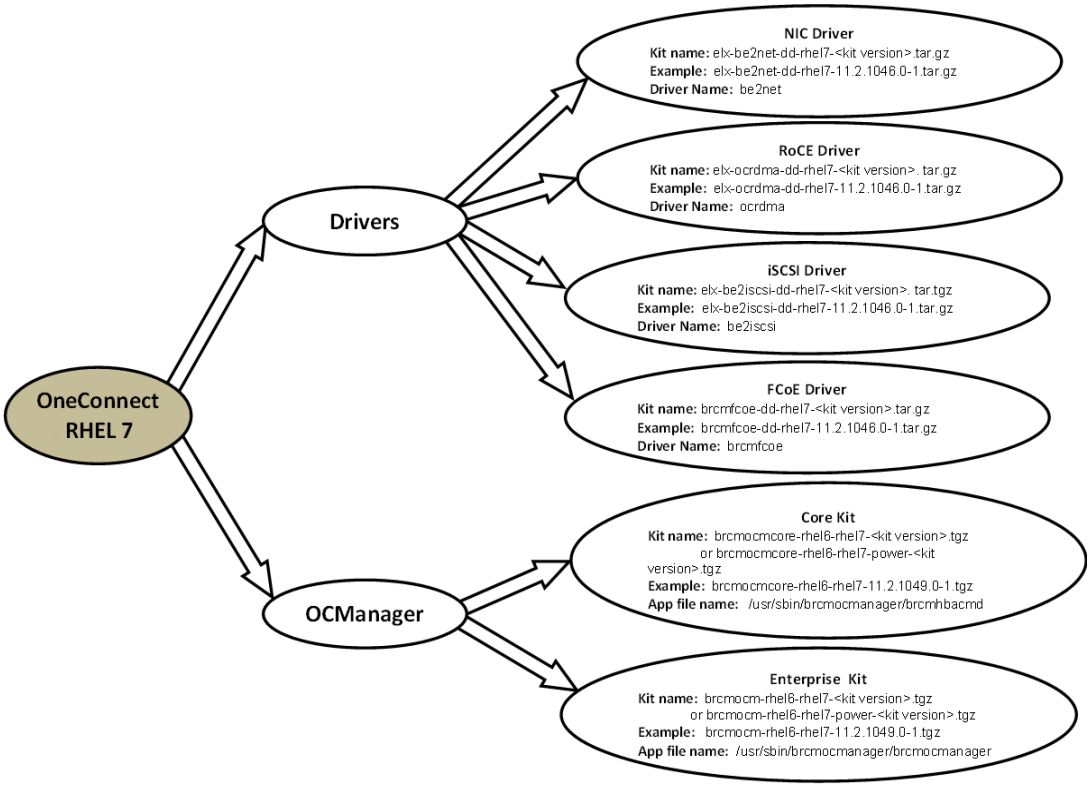

Figure 20 OneConnect SLES 11

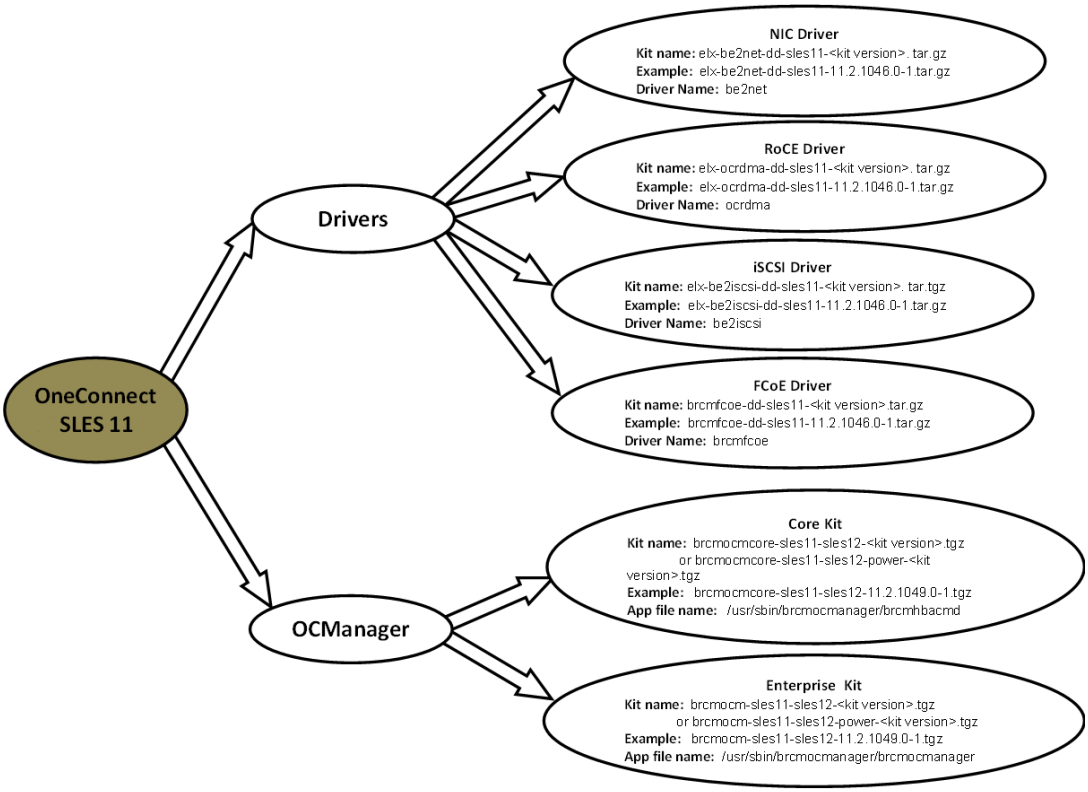

Figure 21 OneConnect SLES 12

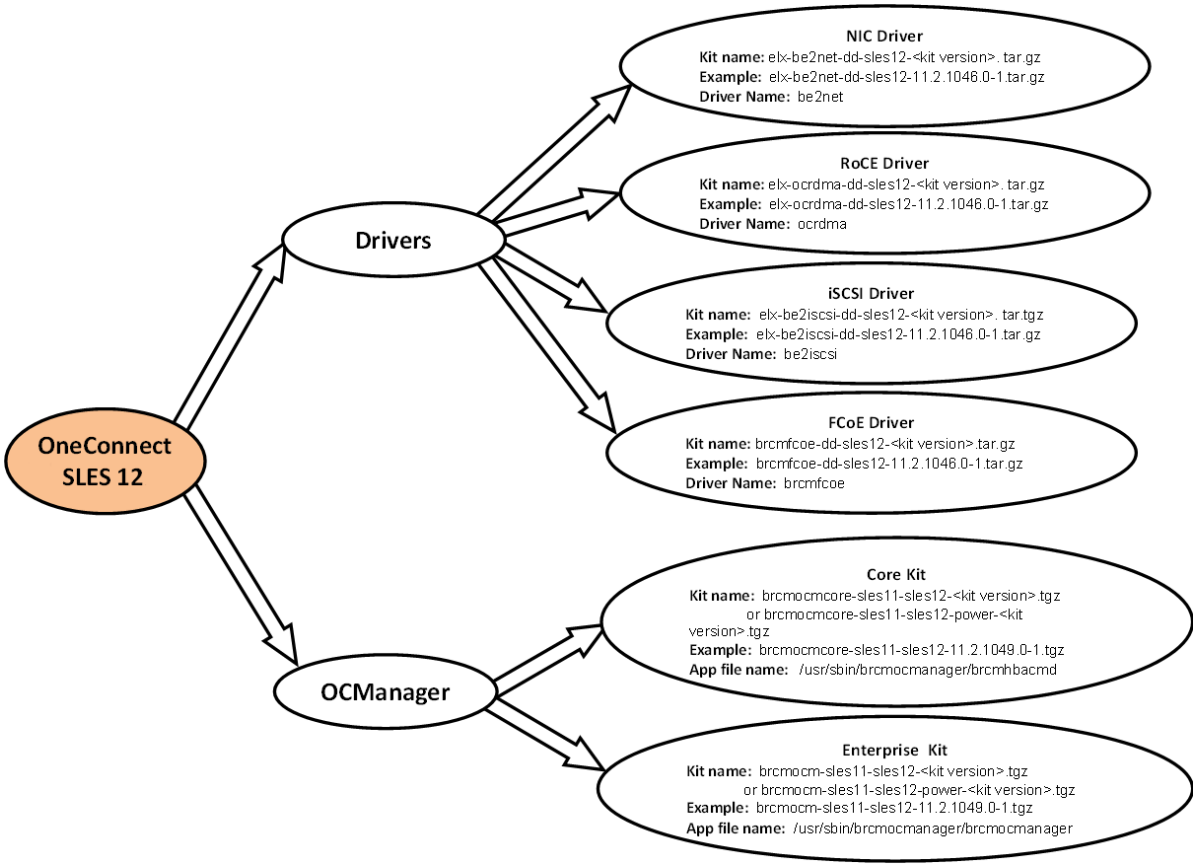

Figure 22 OneConnect Other Linux

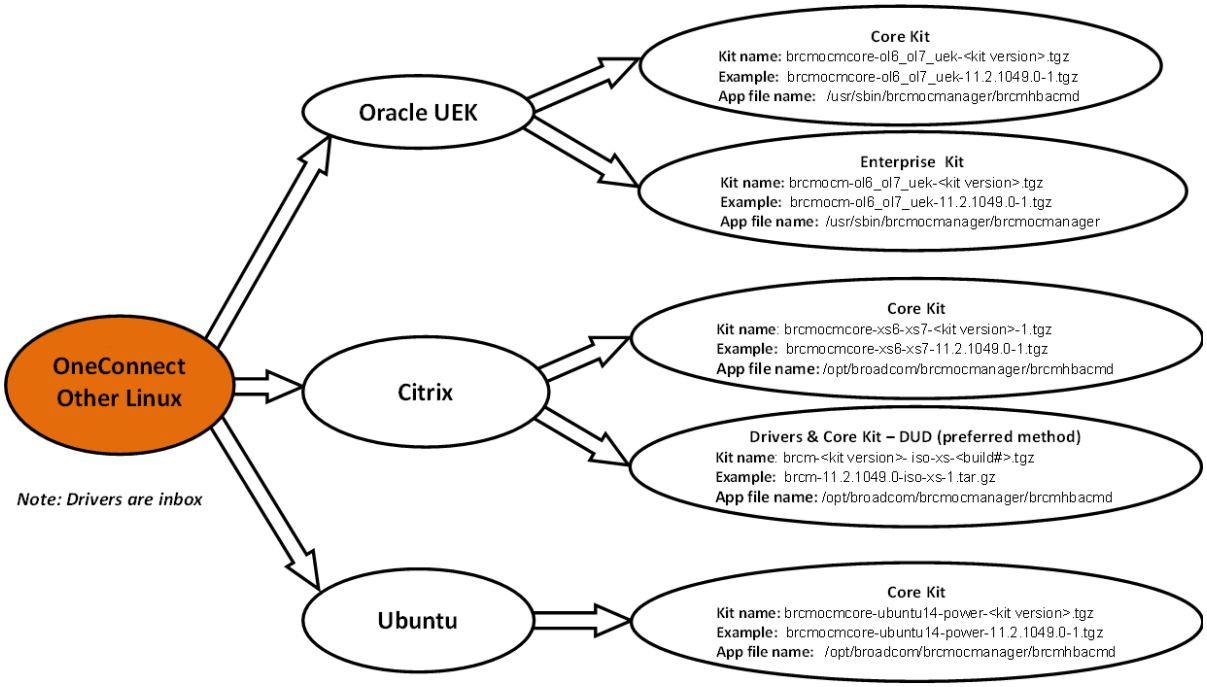

Figure 23 OneConnect ESXi 5.5

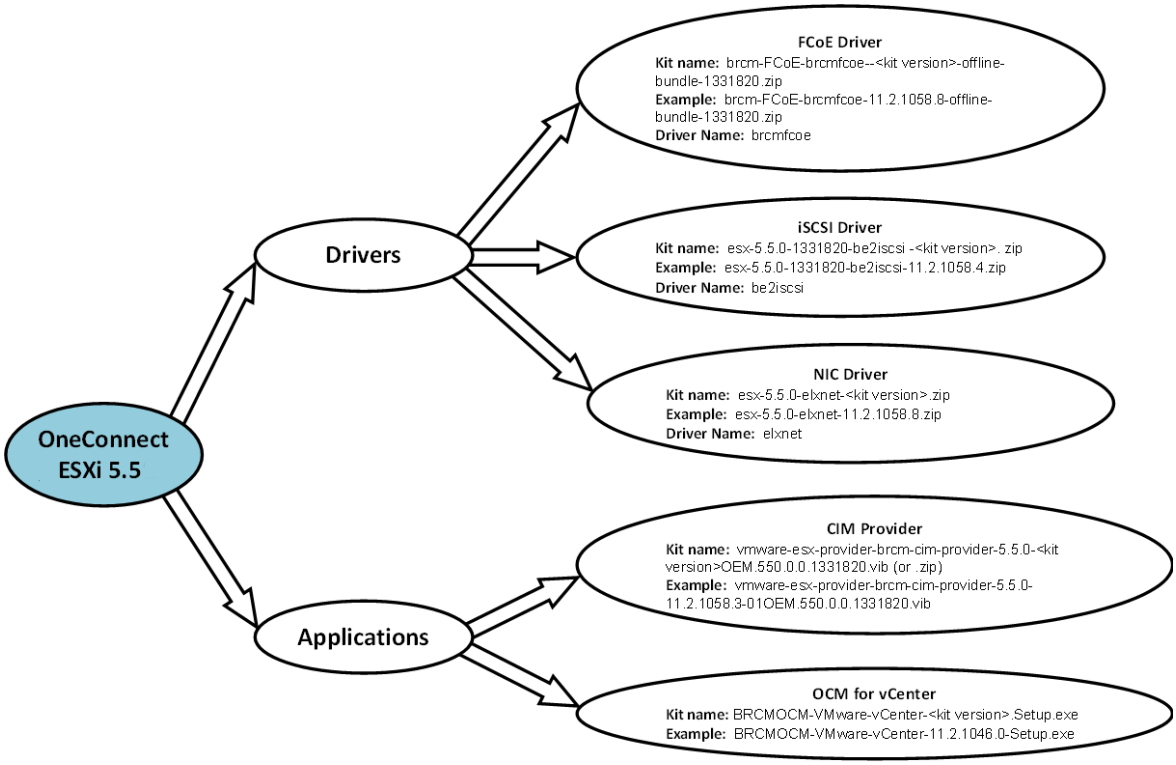

Figure 24 OneConnect ESXi 6.0

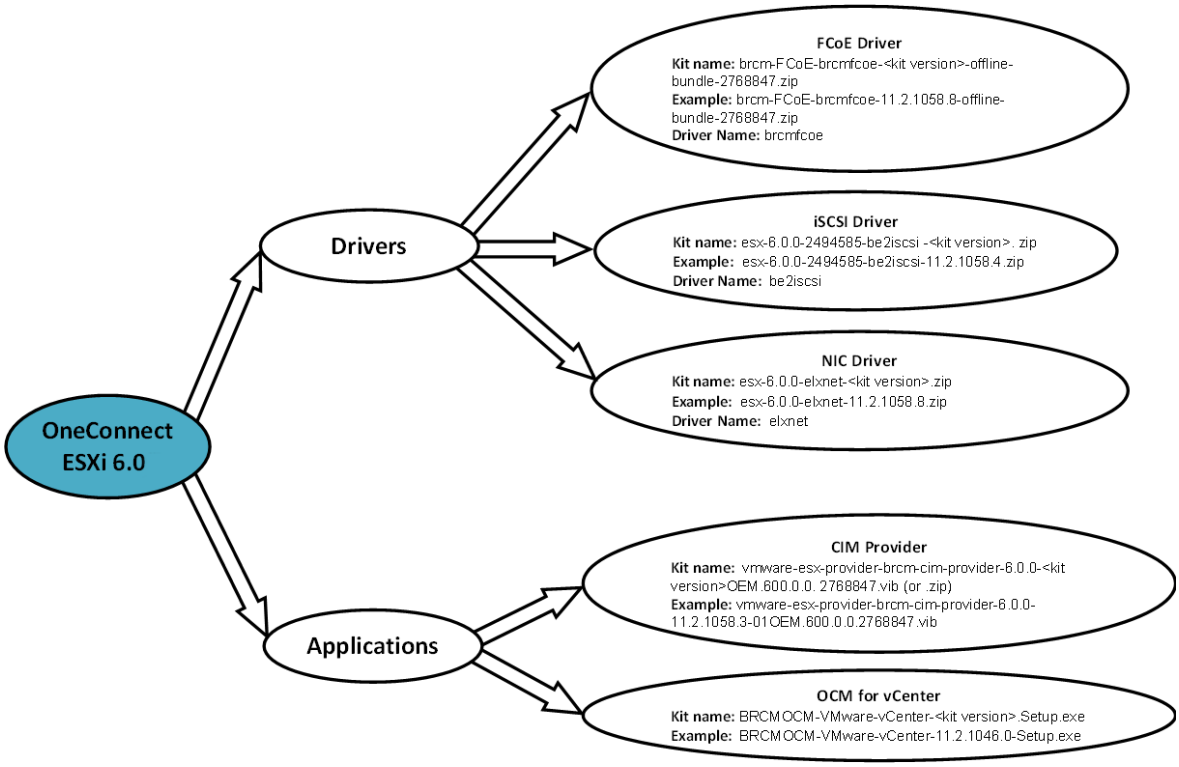

Figure 25 OneConnect ESXi 6.5

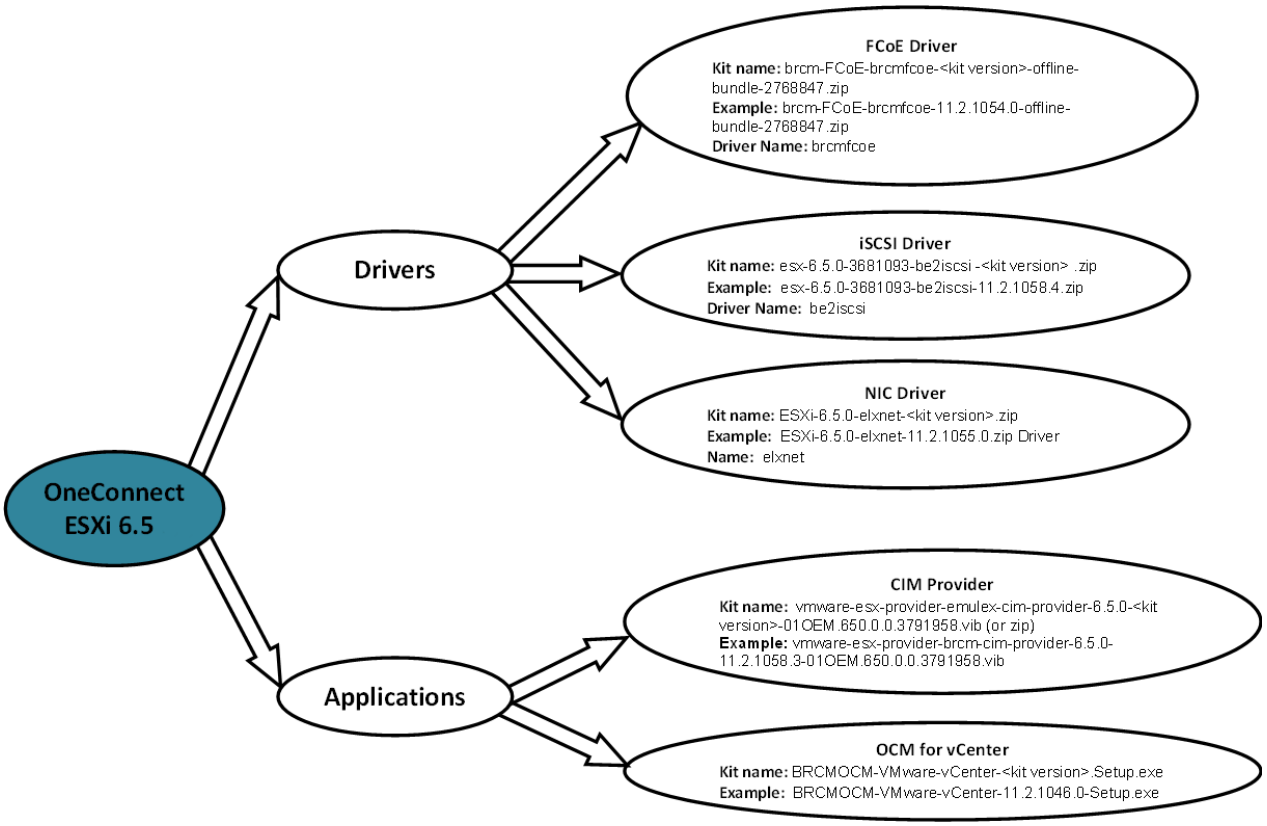

Appendix C: Special Instructions

C.1 Linux

Uninstalling the brcmfcoe driver using the brcmfcoe_install.sh script (`brcmfcoe_install.sh -u`) will remove the brcmfcoe driver and unblacklist the lpfc driver.

C.2 VMware

If you need to restore the inbox driver, you can do this in two ways:

- Reboot your server with a local console connected to the server. As POST is completing, watch for instructions to press `shift R`. This allows the administrator to restore the original pristine kernel.

  When the server has rebooted, check the vmhba enumeration to driver binary as follows:

  ```
  esxcli storage core adapter list
  ```

- VMware provides a **depot** zip file with the original inbox drivers. The original lpfc driver is included in this depot zip file. Remove the OOB driver(s) and install the original lpfc driver bundle by typing:

  ```
  esxcli software vib remove -n lpfc (if necessary)
  esxcli software vib remove -n brcmfcoe (if necessary)
  esxcli software vib install -d <path to inbox bundle zip file>
  reboot
  ```

  When the server has rebooted, check the vmhba enumeration to driver binary as follows:

  ```
  esxcli storage core adapter list
  ```

All installation instructions state:

`"When the server has rebooted, check the vmhba enumeration…".`

The ESXi device manager modifies vmhba numbering because the split driver installation modifies the PCI ID ownership to potentially different driver names. For more information about renumbering, see the following example:

Example

Prior to split kit installation:

NOTE The vmhba instances are vmhba2 to 9. In particular, vmhba2 to 5 are FC and vmhba6 to 9 are FCoE.

```
root@lucic:~] esxcli storage core adapter list
vmhba2  lpfc  link-up   fc.20000090fa021916: (0000:08:00.0) Emulex Corporation Emulex
                        10000090fa021916     LightPulse LPe16000 PCIe Fibre Channel
                                             Adapter
vmhba3  lpfc  link-up   fc.20000090fa021917: (0000:08:00.1) Emulex Corporation Emulex
                        10000090fa021917     LightPulse LPe16000 PCIe Fibre Channel
                                             Adapter
vmhba4  lpfc  link-up   fc.20000000c9f3f9b6: (0000:06:00.0) Emulex Corporation Emulex
                        10000000c9f3f9b6     LPe12000 8Gb PCIe Fibre Channel Adapter
vmhba5  lpfc  link-up   fc.20000000c9f3f9b7: (0000:06:00.1) Emulex Corporation Emulex
                        10000000c9f3f9b7     LPe12000 8Gb PCIe Fibre Channel Adapter
```

```
vmhba6  lpfc  link-up   fc.20000000c9ef3dab: (0000:09:00.2) ServerEngines Corporation
                         10000000c9ef3dab     Emulex OneConnect OCe11100 FCoE Initiator
vmhba7  lpfc  link-up   fc.20000000c9ef3daf: (0000:09:00.3) ServerEngines Corporation
                         10000000c9ef3daf     Emulex OneConnect OCe11100 FCoE Initiator
vmhba8  lpfc  link-n/a  fc.20000090fa5d2f29: (0000:23:00.2) Emulex Corporation Emulex
                         10000090fa5d2f29     OneConnect OCe14000, FCoE Initiator
vmhba9  lpfc  link-n/a  fc.20000090fa5d2f31: (0000:23:00.3) Emulex Corporation Emulex
                         10000090fa5d2f31     OneConnect OCe14000, FCoE Initiator
```

Install the new OOB LightPulse/OneConnect driver bundles:

```
[root@lucic:~] esxcli software vib install -d
Emulex-FCoE-FC-lpfc-11.2.50.60-offline-bundle-2768847.zip
```

Installation Result:

Message: `The update completed successfully, but the system needs to be rebooted for the changes to be effective.`

```
Reboot Required: true
VIBs Installed: EMU_bootbank_lpfc_11.2.50.60-1OEM.600.0.0.2768847
VIBs Removed: Emulex_bootbank_lpfc_11.1.166.0-1OEM.550.0.0.1331820
VIBs Skipped:
[root@lucic:~] esxcli software vib install -d
brcm-FCoE-brcmfcoe-11.2.41.63-offline-bundle-2768847.zip
```

Installation Result:

Message: `The update completed successfully, but the system needs to be rebooted for the changes to be effective.`

```
Reboot Required: true
VIBs Installed: EMU_bootbank_brcmfcoe_11.2.41.63-1OEM.600.0.0.2768847
VIBs Removed:
VIBs Skipped:
[root@lucic:~] reboot
```

After installing OOB LightPulse/OneConnect drivers:

NOTE vmhba2 to 5 are still FC. However, vmhba6 to9 have been renumbered to vmhba64 to 67. The re-enumeration of the vmhbas is current because the paths are still visible.

```
root@lucic:~] esxcli storage core adapter list
HBA Name Driver    Link      UID                  Capabilities        Description
                   State
vmhba2   lpfc      link-up   fc.20000090fa021916: Second Level Lun ID Emulex Corporation
                             10000090fa021916     (0000:08:00.0)      Emulex LightPulse
                                                                      LPe16000 PCIe Fibre
                                                                      Channel Adapter
vmhba3   lpfc      link-up   fc.20000090fa021917: Second Level Lun ID Emulex Corporation
                             10000090fa021917     (0000:08:00.1)      Emulex LightPulse
                                                                      LPe16000 PCIe Fibre
                                                                      Channel Adapter
vmhba64  brcmfcoe  link-up   fc.20000000c9ef3dab: Second Level Lun ID ServerEngines
                             10000000c9ef3dab     (0000:09:00.2)      Corporation Emulex
                                                                      OneConnect OCe11100
                                                                      FCoE Initiator
```

```
vmhba4   lpfc      link-up   fc.20000000c9f3f9b6: Second Level Lun ID  Emulex Corporation
                             10000000c9f3f9b6     (0000:06:00.0)      Emulex LPe12000 8Gb
                                                                      PCIe Fibre Channel
                                                                      Adapter
vmhba65  brcmfcoe  link-up   fc.20000000c9ef3daf: Second Level Lun ID  ServerEngines
                             10000000c9ef3daf     (0000:09:00.3)      Corporation Emulex
                                                                      OneConnect OCe11100
                                                                      FCoE Initiator
vmhba5   lpfc      link-up   fc.20000000c9f3f9b7: Second Level Lun ID  Emulex Corporation
                             10000000c9f3f9b7     (0000:06:00.1)      Emulex LPe12000 8Gb
                                                                      PCIe Fibre Channel
                                                                      Adapter
vmhba66  brcmfcoe  link-n/a  fc.20000090fa5d2f29: Second Level Lun ID  Emulex Corporation
                             10000090fa5d2f29     (0000:23:00.2)      Emulex OneConnect
                                                                      OCe14000, FCoE
                                                                      Initiator
vmhba67  brcmfcoe  link-n/a  fc.20000090fa5d2f31: Second Level Lun ID  Emulex Corporation
                             10000090fa5d2f31     (0000:23:00.3)      Emulex OneConnect
                                                                      OCe14000, FCoE
                                                                      Initiator
```

```
[root@lucic:~] esxcli storage core path list | grep -e vmhba2 -e vmhba64
Runtime Name:vmhba64:C0:T0:L0
Adapter:vmhba64
Runtime Name:vmhba2:C0:T0:L0
Adapter:vmhba2
Runtime Name:vmhba2:C0:T0:L1
Adapter:vmhba2
Runtime Name:vmhba2:C0:T0:L2
Adapter:vmhba2
Runtime Name:vmhba2:C0:T0:L3
Adapter:vmhba2
Runtime Name:vmhba64:C0:T1:L0
Adapter:vmhba64
```

This renumbering can be repaired by editing the `/etc/vmware/esx.conf` file:

1. Change the directory to `/etc/vmware` and save the existing `esx.conf` to `esx.conf.old`.
2. Open the `esx.conf` file for editing. Search for the vmhba6 -> vmhba9 and vmhba64 -> vmhba67 instances (following this example), delete those lines, save the file, and reboot.

After the `esx.conf` file has been fixed, the output looks like the following:

```
[root@lucic:~] esxcli storage core adapter list
HBA Name Driver    Link      UID                  Capabilities        Description
                   State
vmhba2   lpfc      link-up   fc.20000090fa021916: Second Level Lun ID  Emulex Corporation
                             10000090fa021916     (0000:08:00.0)      Emulex LightPulse
                                                                      LPe16000 PCIe Fibre
                                                                      Channel Adapter
vmhba3   lpfc      link-up   fc.20000090fa021917: Second Level Lun ID  Emulex Corporation
                             10000090fa021917     (0000:08:00.1)      Emulex LightPulse
                                                                      LPe16000 PCIe Fibre
                                                                      Channel Adapter
```

```
vmhba4   lpfc      link-up   fc.20000000c9f3f9b6: Second Level Lun ID  Emulex Corporation
                             10000000c9f3f9b6     (0000:06:00.0)       Emulex LPe12000 8Gb
                                                                       PCIe Fibre Channel
                                                                       Adapter
vmhba5   lpfc      link-up   fc.20000000c9f3f9b7: Second Level Lun ID  Emulex Corporation
                             10000000c9f3f9b7     (0000:06:00.1)       Emulex LPe12000 8Gb
                                                                       PCIe Fibre Channel
                                                                       Adapter
vmhba6   brcmfcoe  link-up   fc.20000000c9ef3dab: Second Level Lun ID  ServerEngines
                             10000000c9ef3dab     (0000:09:00.2)       Corporation Emulex
                                                                       OneConnect OCe11100
                                                                       FCoE Initiator
vmhba7   brcmfcoe  link-up   fc.20000000c9ef3daf: Second Level Lun ID  ServerEngines
                             10000000c9ef3daf     (0000:09:00.3)       Corporation Emulex
                                                                       OneConnect OCe11100
                                                                       FCoE Initiator
vmhba8   brcmfcoe  link-n/a  fc.20000090fa5d2f29: Second Level Lun ID  Emulex Corporation
                             10000090fa5d2f29     (0000:23:00.2)       Emulex OneConnect
                                                                       OCe14000, FCoE
                                                                       Initiator
vmhba9   brcmfcoe  link-n/a  fc.20000090fa5d2f31: Second Level Lun ID  Emulex Corporation
                             10000090fa5d2f31     (0000:23:00.3)       Emulex OneConnect
                                                                       OCe14000, FCoE
                                                                       Initiator
```

Helpful VMware KB articles can be found by using the following link:

https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2091560


BROADCOM®

Emulex® OneCommand® Manager Application for LightPulse® Adapters

User Guide

Version 11.4
September 6, 2017

OCM-APP-LPE-UG114-100

Corporate Headquarters
San Jose, CA

Website
www.broadcom.com

Chapter 1: Introduction

The Emulex® OneCommand ® Manager application is a comprehensive management utility for Emulex adapters that provides a powerful, centralized adapter management suite. Adapter management includes discovery, reporting, and management of local and remote adapters from a single console anywhere in the network and across operating system platforms. Remote configuration capability is provided by Transmission Control Protocol/Internet Protocol (TCP/IP) access from IP addresses of remote machines. The OneCommand Manager application contains a graphical user interface (GUI) and a command line interface (CLI). Refer to the *Emulex OneCommand Manager Command Line Interface for LightPulse Adapters User Guide* for information about installing and using the CLI.

NOTE Screenshots in this user guide are for illustrative purposes only. Your system information can vary.

The OneCommand Manager application can be installed on multiple operating systems, including Windows, Linux, and Solaris. For supported versions of operating systems, platforms, and adapters, go to http://www.broadcom.com.

For VMware hosts, use the OneCommand Manager application for VMware vCenter. For more details, refer to the *Emulex OneCommand Manager for VMware vCenter for LightPulse Adapters User Guide*. You can manage adapters using the OneCommand Manager application on Windows, but you must install and use the appropriate Emulex CIM Provider.

1.1 Abbreviations

AL_PA	Arbitrated Loop Physical Address
AP	access point
API	application programming interface
ASIC	application-specific integrated circuit
BIOS	basic input/output system
BOFM	Blade Open Firmware Management Protocol
CND	congestion notification domain
CIM	Common Interface Model
CLI	command line interface
CLP	Command Line Protocol
CRC	cyclic redundancy check
CSV	comma-separated values
D_ID	destination identifier
DCB	Data Center Bridging
DCBX	Data Center Bridging Capabilities Exchange
DDR	double data rate
DH	Diffie-Hellman
DHCHAP	Diffie-Hellman Challenge Handshake Authentication Protocol
DHCP	Dynamic Host Control Protocol
DID	device ID
DMA	direct memory access

EDD	Enhanced Disk Drive
EFI	Extensible Firmware Interface
EFD	Enhanced FAT Dump
F_BSY	FC port busy
FA-PWWN	Fabric Assigned Port Word Wide Name
FC	Fibre Channel
FC-SP	Fibre Channel Security Protocol
FCoE	Fibre Channel over Ethernet
FIP	FCoE Initialization Protocol
FLOGI	Fabric login
GFC	gigabit Fibre Channel
GFO	Get Fabric Object
GUI	graphical user interface
HBA	host bus adapter
HTTP	Hypertext Transfer Protocol
I/O	input/output
IEEE	Institute of Electrical and Electronics Engineers
IIS	Internet Information Services
IP	Internet Protocol
IPL	initial program load
JEDEC	Joint Electron Device Engineering Council
JNLP	Java Network Launching Protocol
JRE	Java Runtime Environment
LAN	local area network
LDAP	Lightweight Directory Access Protocol
LED	light-emitting diode
LIP	Loop Initialization Primitive
LLDP	Link Layer Discovery Protocol
LUN	logical unit number
MAC	Media Access Control
MIME	Multipurpose Internet Mail Extension
MTU	maximum transmission unit
NCSI	Network Communication Services Interface
NIC	network interface card
NOS	network operating system
NPIV	N_Port ID Virtualization
NVRAM	non volatile random access memory
OS	operating system
OUI	Organizationally Unique Identifier
PAM	pluggable authentication modules

PCI	Peripheral Component Interconnect (interface)
PFC	priority flow control
PG	priority group
POST	power-on self-test
PXE	Preboot Execution Environment
QoS	quality of service
RAID	redundant array of independent disks
RHEL	Red Hat Enterprise Linux
RMAPI	Remote Management application programming interface
SAN	storage area network
SCSI	Small Computer System Interface
SFP	small form-factor pluggable
SFS	Software Foundation Software
SLES	SUSE Linux Enterprise Server
TCP	Transmission Control Protocol
TCP/IP	TCP over Internet Protocol
Tx	transmit
UEFI	Unified Extensible Firmware Interface
UFP	Universal Fabric Port
ULP	Upper Layer Protocol
VF	virtual function
VLAN	virtual local area network
VLAN ID	VLAN identifier
VM	virtual machine
VPD	vital product data
vPort	virtual port
WLAN	wireless LAN
WWN	World Wide Name
WWNN	World Wide Node Name
WWPN	World Wide Port Name

Chapter 2: Installing and Uninstalling OneCommand Manager Application Components

This section describes installing and uninstalling the OneCommand Manager application.

2.1 Installing the OneCommand Manager Application

2.1.1 In Windows

The OneCommand Manager application can be installed two ways:

- Attended installation using the GUI.
- Unattended installation using the command line.

2.1.1.1 Attended Installation in Windows

To install the OneCommand Manager application in Windows, perform these steps:

1. Download the x64 or x86 OneCommand Manager Enterprise Kit installation file from the Documents and Downloads area of http://www.broadcom.com.
2. Navigate to the directory to which you downloaded the file.
3. Double-click the `elxocm-windows-<version>.exe` file. The **Emulex OCManager Enterprise** window appears. Click **Next**. The **Installation Options** window appears.
4. Select the components that you want to install and click Install. After installing the OneCommand Manager application files, the **OneCommand Manager Management Mode** dialog appears (Figure 1).

Figure 1 Management Mode Dialog

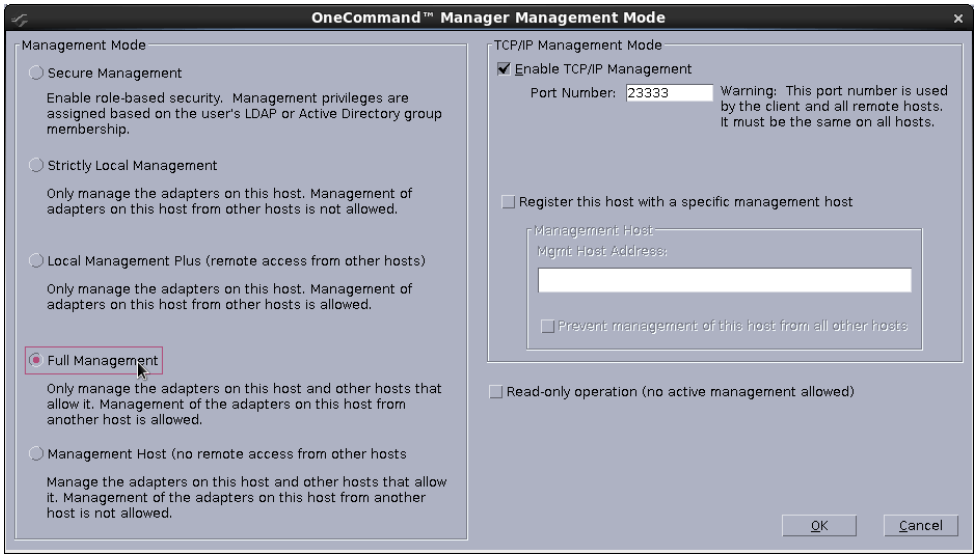

The **OneCommand Manager Management Mode** dialog enables you to select Secure Management to assign the desired user privileges, or you can choose one of the other management modes. See Section 4.2, Using OneCommand Manager Secure Management, or Section 4.3, Changing Management and Read-Only Mode, for more information. Choose the management type you want and click **OK**.

5. Select or clear the **Enable TCP/IP Management** check box to enable or disable remote management over TCP/IP. You can also change the TCP/IP port used (23333 is the IANA registered port for Broadcom).
6. The **Installation Completed** window appears when the installation is finished. Click **Finish**. A shortcut is added to the **Start** menu. You do not need to reboot the system.

2.1.1.2 Unattended Installation in Windows

To install the OneCommand Manager application in Windows, perform these steps:

1. Download the x64 or x86 OneCommand Manager Enterprise Kit installation file to your system from the Documents and Downloads area of http://www.broadcom.com.
2. Activate the kit with switch /q or /q2.
   - The /q switch displays progress reports.
   - The /q2 switch does not display progress reports.
3. You can enable Secure Management mode by adding the `sec=1` argument or disable it by `sec=0`. If the sec argument is not entered, Secure Management mode is disabled by default. See Section 4.2, Using OneCommand Manager Secure Management, for more information.

   To enable Secure Management mode, type the following command at the command prompt:

   ```
   elxocm-windows-x86-<version>.exe  sec=1 /q2
   ```

   To disable Secure Management mode, type the following command at the command prompt:

   ```
   elxocm-windows-x86-<version>.exe  sec=0 /q2
   ```

   NOTE The management mode defaults for unattended installation are:
   - `mmode = 2` (Local Plus Management mode)
   - `achange = 1`
4. You can select a management mode by adding the `mmode` argument and the ability to change that management mode by adding the `achange` argument with selected values as in the following example. Section 4.3, Changing Management and Read-Only Mode, for more information.

   NOTE If you enabled Secure Management mode in step 3 and entered an `mmode` value, it results in a 'conflicting parameters' error.

   For example, type the following command at the command prompt:

   ```
   elxocm-windows-x86-<version>.exe mmode=3 achange=1 /q2
   ```

   The following are the possible mmode values:

   `1` – Local Only Management mode

   `2` – Local Plus Management mode

   `3` – Full Management mode

   `4` – Local Plus Management mode and read only

   `5` – Full Management mode and read only

   `6` – Management host

   The following are the possible `achange` values:

   `0` – Do not allow management mode to change

   `1` – Allow management mode to change

You can also set the following optional parameters:

— `MHost` – This optional switch allows a non-management-host user to select a Management Host with which to register. If this switch is not specified, the default value of 0 is used, and the capability is disabled. If the switch is specified, the value can be a host name or an IP address that is validated by the installer. An error message appears if `/mmode` is set as Local Only or Management Host.

— `excl` – This optional switch allows the non-management-host user to select whether the OneCommand Manager application processes requests exclusively from the Management Host specified by the `MHost` parameter. This option is only accepted if accompanied by a valid `MHost` value; otherwise, an error message appears. If this switch is not specified, the default value of `0` is used. If the parameter is specified, the valid values are:

- `0` – Remotely managed by other hosts.
- `1` – Remotely managed by Management Host *only*.

— `Mtcp` – This optional parameter allows you to enable or disable remote management and to specify the TCP/IP port number over which management occurs. If this parameter is not specified, the default TCP/IP port number `23333` is used.

If the **Management Host** option is selected, you must either select the default port number or enter a valid TCP/IP port number on the command line. A value of `0` is not accepted.

If one of the non-management host options is selected, you can enter the TCP/IP port number on the command line.

2.1.2 In Linux

NOTE The OneCommand Manager application GUI is not supported on Citrix XenServer; however, the OneCommand Manager application CLI is supported. Refer to the *Emulex OneCommand Manager Command Line Interface for LightPulse Adapters User Guide* for Citrix instructions.

The following must be installed before you can install the OneCommand Manager application:

- The appropriate driver version for your operating system. Refer to the Documents and Downloads area of http://www.broadcom.com for the latest drivers.

  NOTE The RHEL 6 Enterprise kit requires the installation of the libstdc++-5.so library. This library is available through the `compat-libstdc++-33-3.2.3-68.<arch>.rpm` or later. The PPC and x86_64 builds require the 64-bit version, which is installed in `/usr/lib64`. The i386 build requires the 32-bit version, which is installed in `/usr/lib`.

- Previous versions of the Linux driver must be uninstalled. You must run the uninstall script that shipped with the version of the Linux driver you want to remove.

2.1.2.1 Attended Installation in Linux

To install the OneCommand Manager application, or to update an existing installation, perform these steps:

1. Log on as root.
2. Download the utilities from the Documents and Downloads area of http://www.broadcom.com.
3. Copy the OneCommand `elxocm-<Platform>-<AppsRev>.tgz` file to a directory on the installation machine.
4. Change to the directory to which you copied the tar file.
5. Untar the file.
   — For RHEL 6 and RHEL 7, type the following:

```
tar zxvf elxocm-rhel6-rhel7-<apps_ver>-<rel>.tgz
```

— For SLES 11and SLES 12, type the following:

```
tar zxvf elxocm-sles11-sles12-<apps_ver>-<rel>.tgz
```

6. Change to the `elxocm` directory created in step 3.

— For RHEL 6 and RHEL 7, type the following:

```
cd elxocm-rhel6-rhel7-<apps_ver>-<rel>
```

— For SLES 11 and SLES 12, type the following:

```
cd elxocm-sles11-sles12-<apps_ver>-<rel>
```

NOTE Prior to installation, OneCommand Manager application groups must be configured on the LDAP network or the local host machine for Secure Management operation. See Section 4.2.1, OneCommand Manager Secure Management Configuration Requirements, for configuration instructions.

7. Run the install script. Type the following:

```
./install.sh
```

8. When prompted, choose whether or not to enable Secure Management for OneCommand:

```
Do you want to enable Secure Management feature for OneCommand? (s/u)
Enter 's' to select secure management. (LDAP/NIS OCM group configuration
required)
Enter 'u' to run without secure management (default).
Enter the letter 's' or 'u'.
```

If you enter `s`, proceed to step 11. You cannot choose a management mode as described in step 9.

9. When prompted, enter the type of management mode you want to use:

```
Enter the type of management you want to use:
1 Local Mode : HBA's on this Platform can be managed by OneCommand clients
on this Platform Only.
2 Managed Mode: HBA's on this Platform can be managed by local or
remote OneCommand clients.
3 Remote Mode : Same as '2' plus OneCommand clients on this Platform can
manage local and remote HBA's.
4 Management Host : Same as '1' plus OneCommand clients on this Platform can
manage remote HBA's.
```

NOTE If you enabled Secure Management in step 8, you cannot configure management mode.

— If you select option `2`, you are asked if you want to enable TCP/IP management from remote hosts.

— If you select option `3`, you are asked if you want to enable TCP/IP management of remote hosts, and enable TCP/IP management from remote hosts. You are prompted to enter the TCP/IP port number to use. (Leaving the field blank defaults to `23333`).

— If you select options `2` or `3`, you are prompted for the management host address. (Leaving the field blanks means none).

— You can enter an IP address or host name. If you enter a management host address, you are prompted to exclude management of this host from any other host.

— If you select option `4`, management of remote hosts is automatically selected and you are prompted to enter the TCP/IP port number to use. (Leaving the field blank defaults to `23333`.)

NOTE Management hosts cannot be managed by remote hosts.

10. If you answered `2`, `3`, or `4` in step 9, you must decide whether you want the OneCommand Manager application to operate in read-only mode. Read-only mode prevents users from performing some operations, such as

resetting adapters, updating an adapter's firmware, or changing adapter driver properties and bindings. It only affects the local OneCommand Manager application interface. These operations can still be performed using remote management. Enter either `y` for yes to allow users to perform these operations, enter `n` for no if read-only mode is desired.

11. You are prompted about allowing users to change the management mode after installation. Enter either `y` for yes, or `n` for no.

2.1.2.2 Unattended Installation in Linux

For unattended or silent installation of the OneCommand Manager application for Linux, installation settings are defined using the installation script command line.

NOTE Prior to installation, OneCommand groups must be configured on the LDAP network or the local host machine for Secure Management operation. See Section 4.2.1, OneCommand Manager Secure Management Configuration Requirements, for configuration instructions.

To view the options for unattended installation, type the following:

```
./install.sh --help
```

To perform an unattended, silent installation, type the following command:

```
#./install.sh -q2
```

NOTE The management mode default for unattended installation is Local Management Plus.

2.1.2.3 Updating an Installation in Linux

The OneCommand Manager application supports the following update paths:

- You can update from an earlier Core Kit to a later Enterprise Kit.
- You can update from an earlier Enterprise Kit to a later Enterprise Kit.

See Section 2.1.2.1, Attended Installation in Linux, or Section 2.1.2.2, Unattended Installation in Linux, for instructions.

2.1.3 In Solaris

The following Solaris drivers must be installed for the utilities to function properly:

- For LPe16202/OCe15100 adapters in NIC mode, the NIC out-of-box driver.
- For LPe16202/OCe15100 adapters in FCoE mode, the FCoE out-of-box driver.
- For all other LightPulse® adapters, the inbox FC driver.

NOTE If an LPe16202/OCe15100 adapter in NIC+FCoE mode is installed in the system, the NIC driver must be installed and reporting all NIC ports. Otherwise, the OneCommand Manager application cannot manage the adapter.

To install the OneCommand Manager application in Solaris, perform these steps:

1. Copy the Solaris utility kit to a temporary directory on your system.
2. Untar the utility kit:

```
tar xvf elxocm-solaris-<version>.tar
```

3. Change to the newly created `elxocm-solaris-<version>` directory:

   ```
   cd ./elxocm-solaris-<version>/
   ```

   NOTE Prior to installation, OneCommand groups must be configured on the LDAP network or the local host machine for Secure Management operation. See Section 4.2.1, OneCommand Manager Secure Management Configuration Requirements, for configuration instructions.

4. Run the installation script to begin installation. If the HBAnyware utility, OneCommand Manager Core, or OneCommand Manager Enterprise applications or the Solaris driver utilities are already present on the system, the installation script attempts to remove them first:

   ```
   ./install
   ```

5. When prompted, choose whether or not to enable Secure Management for OneCommand:

   ```
   Do you want to enable Secure Management feature for OneCommand? (s/u)
   Enter 's' to select secure management. (LDAP/NIS OCM group configuration
   required)
   Enter 'u' to run without secure management (default).
   Enter the letter 's' or 'u'.
   ```

   If you enter `s`, proceed to step 7. You cannot choose a management mode as described in step 6.

6. When prompted, enter the type of management you want to use:

   ```
   Enter the type of management you want to use:
   1 Local Mode:     HBA's on this Platform can be managed by OneCommand
   clients on this Platform Only.
   2 Managed Mode:   HBA's on this Platform can be managed by local or
   remote OneCommand clients.
   3 Remote Mode:    Same as '2' plus OneCommand clients on this Platform can
   manage local and remote HBA's.
   4 Management Host:Same as '1' plus OneCommand clients on this Platform can
   manage remote HBA's.
   ```

   NOTE If you enabled Secure Management in step 5, you cannot configure management mode.

   — If you select option `2`, you are asked if you want to enable TCP/IP management from remote hosts.
   — If you select option `3`, you are asked if you want to enable TCP/IP management of remote hosts, and enable TCP/IP management from remote hosts. You are prompted to enter the TCP/IP port number to use. Leaving the field blank defaults to `23333`.
   — If you select options `2` or `3`, you are prompted for the management host address. Leaving the field blank means none.
   — You can enter an IP address or host name. If you enter a management host address, you are prompted to exclude management of this host from any other host.
   — If you select option `4`, management of remote hosts is automatically selected and you are prompted to enter the TCP/IP port number to use. Leaving the field blank defaults to `23333`.

   NOTE Management hosts cannot be managed by remote hosts.

7. If you answered `2`, `3`, or `4` in step 6, you must decide whether you want the OneCommand Manager application to operate in read-only mode. Read-only mode prevents users from performing some operations such as resetting adapters, updating an adapter's firmware, or changing adapter driver properties and bindings. It only affects the local OneCommand Manager application interface. These operations can still be performed using remote management. Enter either `y` for yes to allow users to perform these operations, or enter `n` for no if read-only mode is desired.

8. You are prompted whether to allow users to change the management mode after installation. Enter either *y* for yes, or *n* for no.

2.1.4 In VMware

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

The Emulex CIM Provider is available as an offline bundle in ESXi platforms. Use the offline bundle to update software on VMware platforms. For more information about the ESXi Patch Management activities, refer to the VMware website.

For the best real-time management of Emulex adapters in VMware ESXi environments, use the OneCommand Manager application for VMware vCenter. For more information, refer to the *Emulex OneCommand Manager for VMware vCenter for LightPulse Adapters User Guide*.

To install the Emulex CIM Provider in a VMware ESXi hypervisor environment, use the esxcli command line utility and perform theses steps:

1. Copy the CIM Provider zip file to `/var/log/vmware`.
2. Log on to the VMware hypervisor host, and execute the following command all on one line:

   ```
   esxcli software vib install -d
   vmware-esx-provider-emulex-cim-provider-<version>.zip --no-sig-check
   ```
3. Reboot the system.

2.2 Uninstalling the OneCommand Manager Application

To uninstall the OneCommand Manager application, perform these steps:

- In Windows:
  a. Select **Start > Control Panel > Programs > Uninstall a Program**.
  b. Select **Emulex OCManager Enterprise <*version*>**, and click **Remove** or **Uninstall**.
- In Linux:
  a. Log on as root.
  b. Change to the `elxocm-<platform>-<version>` installation directory.
  c. Type the following:
     `./uninstall`
- In Solaris:
  a. Log on as root.
  b. Run the OneCommand Manager application uninstallation script:
     `/opt/ELXocm/scripts/uninstall`

Chapter 3: Starting and Stopping the OneCommand Manager Application

This section describes how to start and stop the OneCommand Manager application.

3.1 In Windows

To start the OneCommand Manager application, from the Windows desktop, select **Start > All Programs > Emulex > OCManager**. If Secure Management is enabled, you are prompted for your user name and password. See Section 4.2, Using OneCommand Manager Secure Management, for more information.

To stop the OneCommand Manager application, from the OneCommand Manager application menu bar, select **File > Exit**.

3.2 In Linux and Solaris

OneCommand Manager application Linux and Solaris installations include two basic daemon processes that are affected by the start and stop scripts:

- `elxhbamgrd` – Remote management daemon that services requests from OneCommand Manager application clients running on remote host machines.
- `elxdiscoveryd` – Discovery daemon that maintains all discovery data (remote and local) for OneCommand Manager application clients running on the local machine.

`elxhbamgrd` starts at system boot time. `elxdiscoveryd` starts whenever the OneCommand Manager application GUI process first runs on the host machine.

To start the OneCommand Manager application, use the `ocmanager` script. The script is located in the following OneCommand Manager installation directory:

- Linux: `/usr/sbin/ocmanager`

  Example usage on Linux:

  ```
  # /usr/sbin/ocmanager/ocmanager
  ```

- Solaris: `/opt/ELXocm`

  Example usage on Solaris:

  ```
  # /opt/ELXocm/start_ocmanager
  ```

If Secure Management is enabled, you are prompted for your user name and password. See Section 4.2, Using OneCommand Manager Secure Management, for more information.

To stop the OneCommand Manager application, use one of the following methods:

- From the menu bar, select **File > Exit**.
- From the shell, use the `stop_ocmanager` script located in the OneCommand Manager installation directory.

Example usage on Linux:

```
# /usr/sbin/ocmanager/stop_ocmanager -n
```

Example usage on Solaris:

```
# /opt/ELXocm/stop_ocmanager -n
```

NOTE The `-n` parameter stops only the OneCommand Manager application and associated `elxdiscoveryd` daemon. When the `stop_ocmanager` script is run without parameters, the script stops the OneCommand Manager application and all associated daemons.

Chapter 4: Using the OneCommand Manager Application

To properly view the OneCommand Manager application, make sure that your system meets the following display requirements:

- For Windows, Linux, and Solaris systems, the display resolution must be set to 1024 x 768 or higher. For Windows systems, use the default font size.
- The display must run in 256-color mode or higher. OneCommand Manager application icons use 256 colors. If the display is set for 16 color mode, OneCommand Manager application icons are not displayed.

4.1 The OneCommand Manager Application Window Element Definitions

The **OneCommand Manager application** window (Figure 2) contains five basic components: the menu bar, the toolbar, the discovery-tree, the property tabs, and the status bar.

Figure 2 OneCommand Manager Application Window

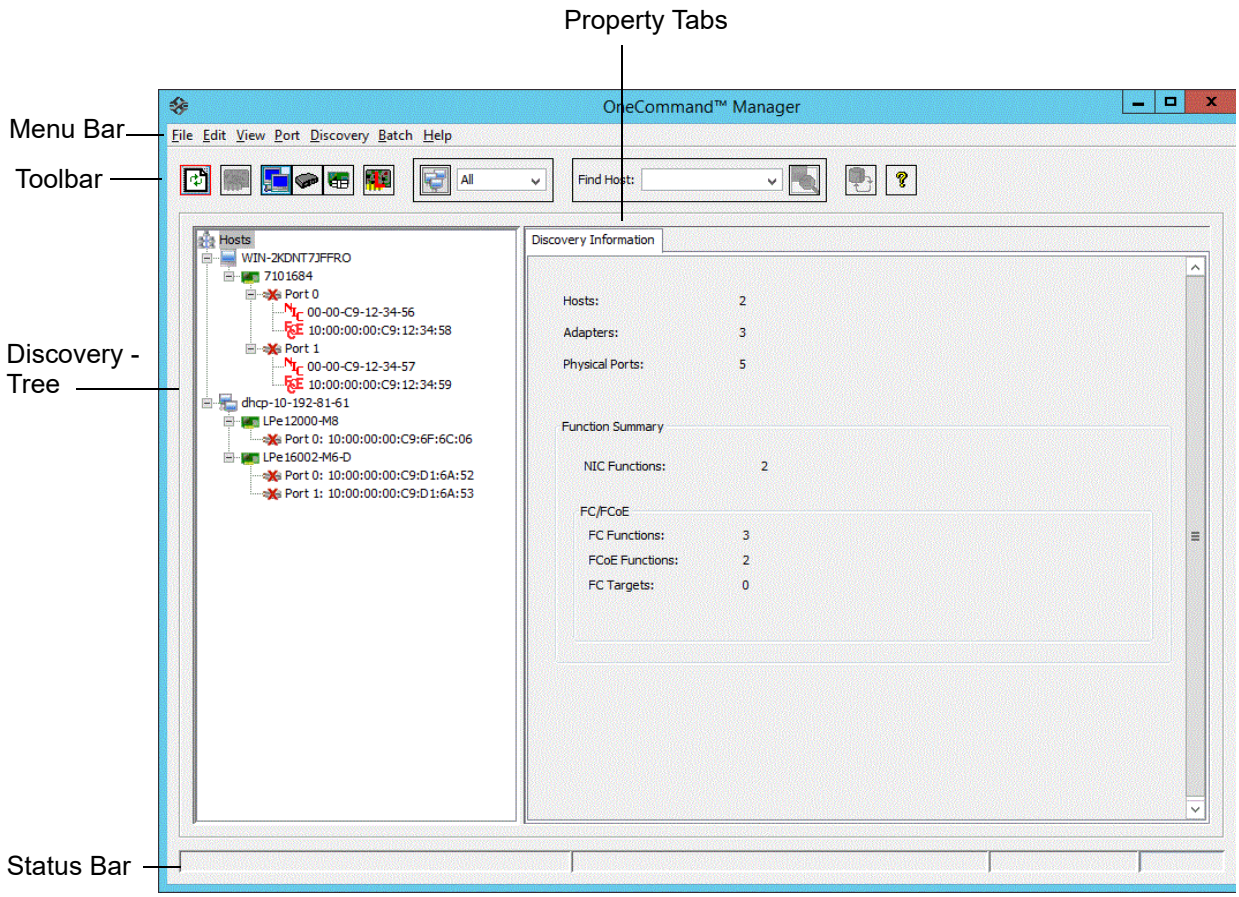

NOTE The element you select in the discovery-tree determines whether a menu item or toolbar icon is active. For example, if you select the local host or other system host, the **Reset Port** item on the **Adapter** menu is unavailable. The **Reset Port** toolbar button is unavailable as well.

The capabilities displayed by your local OneCommand Manager application interface matches those of the remote server. When accessing a remote server running an older version of the OneCommand Manager application, capabilities that are not supported by the server's older version of the OneCommand Manager application are unavailable.

In some instances, the type of information displayed and available functionality is determined by the operating system in use.

4.1.1 Menu Bar

The menu bar contains commands that enable you to perform a variety of tasks, such as exiting the OneCommand Manager application, resetting adapters, and sorting items in the discovery-tree view. Many of the menu bar commands are also available from the toolbar.

4.1.2 Toolbar

The toolbar (Figure 3) contains buttons that enable you to refresh the discovery-tree, reset the selected adapter, and choose how you want to view discovered SAN elements in the discovery-tree. Many of the toolbar functions are also available from the menu bar.

Figure 3 Toolbar

The toolbar is visible by default. Use the **Toolbar** item in the **View** menu to hide the toolbar. If the item is checked, the toolbar is visible.

4.1.2.1 Toolbar Buttons

The toolbar buttons perform the following tasks.

Discovery Refresh button

- Initiates a discovery refresh cycle.

Reset Port button

- Resets the selected FC or FCoE port.

View Buttons on the Toolbar

The View buttons on the toolbar enable you to view SAN elements from the host, fabric, virtual ports, or by local or remote adapter perspective. By default, both local and remote adapters are displayed in the Host view. The OneCommand Manager application displays elements in ascending order.

Host View button (default)
Displays the host system.

NOTE You cannot change host names using the OneCommand Manager application; names must be changed locally on that system.

- Displays the installed adapters within each host system.
- Displays adapter ports and the port numbers if available.
- Displays adapters by the WWNN if multiple adapters have the same model number.
- Displays the WWPN if targets are present. Multiple adapters can refer to the same target.
- Displays the LUN number if LUNs are present.

Fabric View button

- Displays the FC or FCoE fabrics in the SAN with their fabric IDs.
- Displays the ports under each switch.
- If targets are present, displays each WWPN. Multiple adapters can refer to the same target.
- If LUNs are present, displays each LUN number.
- If the fabric ID is all zeros, no fabric is attached.

NOTE NIC ports are not displayed in the Fabric View.

Virtual Ports View button

- Displays virtual ports in the SAN.

NOTE NIC ports are not displayed in the Virtual Ports View.

Local HBAs Only button

- Displays only local adapters.

Show Host Groups button and menu

- Displays hosts by their associated groups.
- Displays available host groups.

Find Host button and search field

- Enables you to search by host name for a particular host in the discovery-tree.

Refresh LUNS button

- Initiates a LUN discovery refresh cycle.

Help button

- Displays the OneCommand Manager application's online help.
- Displays the **About OneCommand Manager** dialog. The dialog displays version information including RMAPI, Discovery, DFClib, MILI Library Version (Windows), and Remote Management Agent Version (Windows). It also enables you to contact Broadcom Technical Support.

4.1.3 Discovery-Tree

The discovery-tree (Figure 4) has icons that represent discovered hosts, adapters, ports, virtual ports, fabrics, targets, and LUNs.

Figure 4 Discovery-Tree

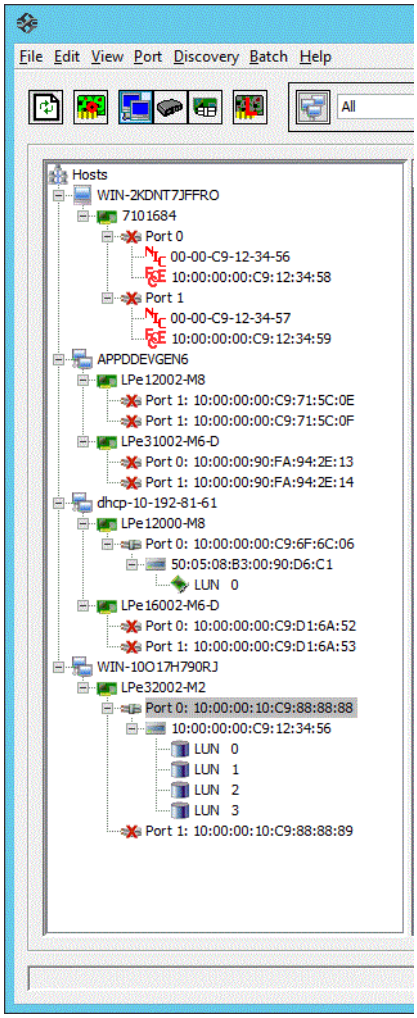

4.1.3.1 Discovery-Tree Icons

Discovery-tree icons represent the following:

The local host.

Other hosts connected to the system.

A green **Adapter** icon with black descriptive text represents an online adapter. Blue text represents an adapter port that had previously been discovered, but currently is not being seen by the discovery engine (service). The adapter is removed from the discovery-tree if it still is not seen after the undiscovered adapter expiration time has elapsed (default is 1800 seconds, or 30 minutes). If the adapter is discovered again before the expiration time has elapsed, it reverts back to normal black text. See Section 5.2, Configuring Discovery and Default CIM Credentials, for more information about discovery settings.

The **Port** icon represents an adapter port. A **Port** icon with a red X indicates the link is down.

NOTE Multiport adapters are represented in the discovery-tree with separate port icons for each port with the port number displayed next to the icon.

A green **FCoE** icon represents an FCoE PCI function online instance. A black **FCoE** icon represents an FCoE PCI function port-disabled instance. A red **FCoE** icon represents an FCoE PCI function link down instance (LPe16202/OCe15100 adapters only).

The **NIC** icon represents a NIC-only PCI function instance. A green icon indicates this function instance is online, black indicates it is disabled, and red indicates a link down instance (LPe16202/OCe15100 adapters only).

The **ASIC Node** icon, only displayed for dual ASIC adapters, represents each ASIC on the adapter. Each ASIC is managed independently. The ASIC node format **ASIC bus#-sub-adapter#** represents the PCI bus number and the sub-adapter number, which is a concatenation of the discovered port numbers for the ASIC. For example, **ASIC 64-12** represents PCI bus number 64, and 12 represents ports 1 and 2. If there were no discovered functions for a port on that ASIC, the label would be **ASIC 64-2** (port 1 is missing).

The **Virtual Port** icon represents a virtual port.

The **Target** icon represents connections to individual storage devices.

The **LUN** icon represents connections to individual disk LUNs.

The **Masked LUN** icon represents a LUN not presented to the host.

The **ExpressLane LUN** icon represents a LUN with ExpressLane™ priority queuing enabled.

The **Media Exchanger** icon represents connections to individual media exchangers. A media exchanger is a jukebox-type device that is capable of swapping various media device instances (such as records or CDs) in and out.

The **Tape LUN** icon represents LUNs that are tape devices.

The **Target Controller LUN** icon represents LUNs that are storage controllers.

The **Switch** icon represents connections to the switch.

4.1.3.2 Expanding or Collapsing the Discovery-Tree View

You can use the Expand/Collapse capability on the **View** menu to change the way discovered elements are displayed. By selecting one of the five levels, the discovery-tree (Figure 4) is expanded or collapsed to that level. You can choose Hosts/Fabrics (depending on the view), adapters, ports, PCI functions, and targets.

4.1.4 Property Tabs

The property tabs display configuration, statistical, and status information for network elements (Figure 2). The set of available tabs is context-sensitive, depending on the type of network element or adapter port currently selected in the discovery-tree (Figure 4).

4.1.5 Status Bar

The status bar is located near the bottom of the **OneCommand Manager application** window (Figure 2). The status bar displays messages about OneCommand Manager application functions, such as `Discovery in progress` or the progress when performing an Export SAN Info operation.

The status bar is visible by default. Use the **Status Bar** item in the **View** menu to hide the status bar. If checked, the status bar is visible.

4.2 Using OneCommand Manager Secure Management

OneCommand Manager Secure Management gives system administrators the ability to further enhance the active management security of their networks. Using Secure Management, administrators can define each user's privileges for managing both local and remote adapters. When running in Secure Management mode, users must log on with their user name and password to run the OneCommand Manager application. If users are authenticated, they can only perform the functions allowed by the OneCommand Manager user group to which they belong. If the systems are running in an LDAP or Active Directory domain, the OneCommand Manager application authenticates users with those defined in that domain. For Linux and Solaris systems, this is accomplished using PAM.

NOTE OneCommand Manager Secure Management is not supported on VMware hosts.

Administrators set up user accounts such that users belong to one of the OneCommand Manager application user groups. The user groups define the management capabilities for each user.

Table 1 defines the OneCommand Manager application user groups and each group's management capabilities.

Table 1 Secure Management User Privileges

Group Name	OneCommand Manager Capability
ocmadmin	Allows full active management of local and remote adapters.
ocmlocaladmin	Permits full active management of local adapters only.
ocmuser	Permits read-only access of local and remote adapters.
ocmlocaluser	Permits read-only access of local adapters.

On Linux or Solaris systems, the UNIX `getent group` utility can be run on the target host system's command shell to verify the correct configuration of the groups. The groups, and users within the groups, appear in the output of this command.

NOTE Although users may belong to the administrator group or be root users, they do not have full privileges to run the OneCommand Manager application unless they are also members of the ocmadmin group. Otherwise, if secure management is enabled, root users or administrators can only manage local adapters (similar to the ocmlocaladmin users).

Remote management operations between two machines are allowed or denied depending on the OneCommand Manager secure management status of the machines, and the domains to which the machines belong. The following tables list the behavior (assuming appropriate user credentials are used).

Table 2 Active Commands: Machines on Same Domain

	Remote Server (Secure)	Remote Server (Not Secure)
Client (Secure)	Allowed	Denied [1]
Client (Not Secure)	Denied	Allowed

1. To inform you of an unsecured server that you may want to secure.

Table 3 Active Commands: Machines on Different Domain

	Remote Server (Secure)	Remote Server (Not Secure)
Client (Secure)	Denied[1]	Denied [2]
Client (Not Secure)	Denied	Allowed

1. Allowed if the username and password are the same on both domains.
2. To inform you of an unsecured server that you may want to secure.

Table 4 Passive Commands: Machines on Any Domain

	Remote Server (Secure)	Remote Server (Not Secure)
Client (Secure)	Allowed	Allowed
Client (Not Secure)	Allowed	Allowed

4.2.1 OneCommand Manager Secure Management Configuration Requirements

For systems to run OneCommand Manager Secure Management, they must be configured to provide the following two capabilities:

1. Authentication – On Linux and Solaris systems, this is accomplished using the PAM interface and must be configured as follows:
   - For Solaris systems, place the correct setting in the `auth` section of `/etc/pam.d/other` file or its earlier equivalent `/etc/pam.conf`.
   - For Linux systems, this is the `/etc/pam.d/passwd` file `auth` section or equivalent.
2. User Group Membership – From the host machine, OneCommand Manager Secure Management must be able to access the OneCommand Manager group to which the user belongs. For Linux and Solaris systems, it uses the `getgrnam` and `getgrid` C-library API calls. The equivalent to the API calls can be obtained by typing `getent group` from the shell command line. If the four OneCommand Manager group names are listed with their member users, the machine is ready to use OneCommand Manager secure management.
3. For Solaris systems, you must use `useradd -G <groupname>` for authentication to work. You cannot use a lowercase `g`.

4.3 Changing Management and Read-Only Mode

NOTE This functionality is only available to root users and administrators even when running in Secure Management mode.

During installation, a management and a read-only mode are selected. If modification of these settings after installation was selected, you can change the management mode:

- Secure Management – The setting enables roles-based security. See Section 4.2, Using OneCommand Manager Secure Management, for details.
- Strictly Local Management – This setting allows management of adapters on this host. Management of adapters on this host from other hosts is not allowed.
- Local Management Plus – This setting only allows management of adapters on this host, but management of adapters on this host from another host is possible.
- Full Management – This setting enables you to manage adapters on this host and other hosts that allow it.

- Management Host – This setting allows this host to manage other hosts, but prevents it from being managed by other hosts.
- Enable TCP/IP Management (of/from remote host) – This setting enables you to manage remote hosts or to manage this host remotely. If enabled, you must supply the port number (between 1024 and 65535). The default port number is `23333`. If the port number or the **Enable TCP/IP Management** check box is changed, a set of warning messages may appear before changes are made. Click **Yes** to continue with the change.

  If the IP port number is changed, the utility restarts the OneCommand Manager application discovery server and management agent to use the new settings. If the servers cannot be stopped and restarted, you are prompted to reboot the host for the new TCP/IP management settings to take effect.

  CAUTION The IP port number must be the same for all hosts that are to be managed. Setting an IP port number for one host to a different value than the other hosts makes the host unable to manage other hosts over TCP/IP using a different port. It also makes the host unmanageable over TCP/IP from other hosts using a different port.
- Register this host with specific management host – This setting enables you to register this host with a specific host for management. If enabled, you must supply the IP address or host name of the management host. You can also choose to prevent management of this host from any other host but the management host. See Section 4.3.1, Management Host, for more information.

If Local Management Plus or Full Management mode is selected, you can also set read-only mode.

- Read-only operation – This setting prevents some operations from being performed, such as resetting adapters, updating the adapter firmware image, and changing adapter settings and driver properties. Dialog controls that pertain to these tasks are completely hidden or disabled.

4.3.1 Management Host

The OneCommand Manager application management host provides enhanced discovery and security by enabling a managed host to register with a management host. The management host receives these registrations when the remote host is started and updates its hosts file so the discovery server discovers the remotely managed host. You do not need to manually add remote hosts to be managed.

If you choose to exclude management from all hosts except the management host, the managed host only responds to requests from the management host. All requests from other hosts are rejected. This TCP/IP management security solution only allows the management host to manage the remote host.

To change management mode and read-only type, perform these steps:

NOTE After making changes, you must restart the OneCommand Manager application to see the new management mode settings.

- In Windows
  a. From the **File** menu, select **Management Mode**. The **Management Mode** dialog appears (Figure 5).

Figure 5 Management Mode Dialog

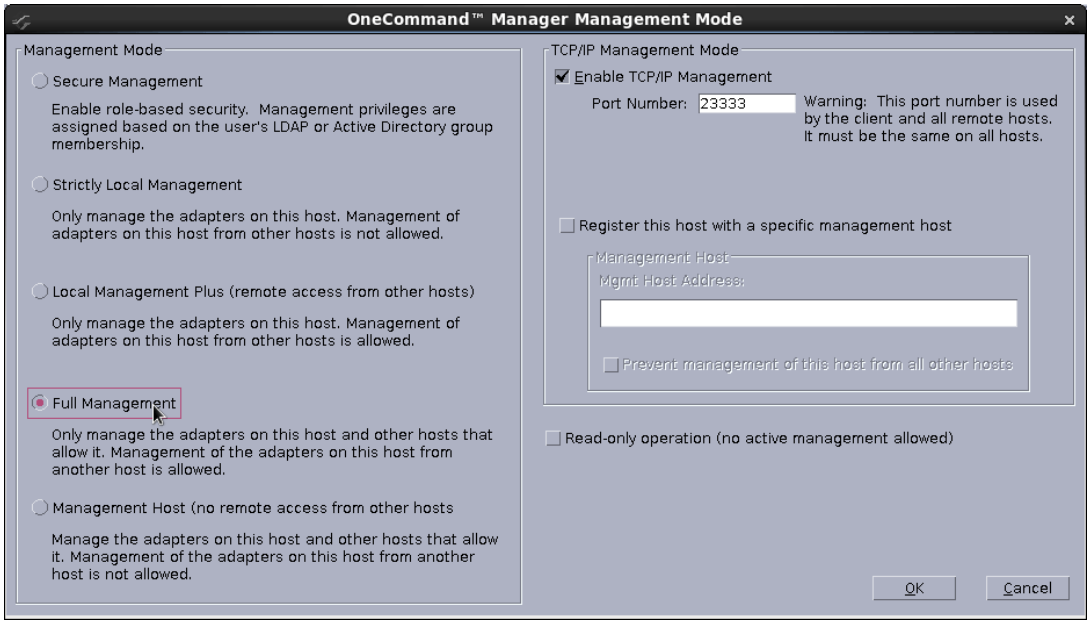

b. Choose the management type and read-only mode you want.
c. Click **OK**.

- In Solaris
  a. Run the following script:
     `/opt/ELXocm/set_operating_mode`
  b. Choose the management type and read-only mode you want.
- In Linux
  a. Stop the OneCommand Manager application.
  b. Run the following script:
     `/usr/sbin/ocmanager/set_operating_mode`
  c. Choose the management type and read-only mode you want.

4.4 Using CIM (Windows Only)

VMware uses CIM as the only standard mechanism for device management. The OneCommand Manager application uses the standard CIM interfaces to manage the adapters in the Visor environment and supports CIM-based devices and HBA management.

To manage the adapters on a VMware host using the OneCommand Manager application, you must install the Emulex CIM Provider on the VMware host.

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

NOTE For VMware hosts, if advanced adapter management capabilities are required (for example, port disable), use the OneCommand Manager application for VMware vCenter. For more details, refer to the *Emulex OneCommand Manager for VMware vCenter for LightPulse Adapters User Guide*.

Chapter 5: Configuring Discovery

This section describes how to configure discovery to find Emulex adapters on remote hosts.

5.1 Discovery Using the TCP/IP Access Protocol

You can discover adapters on IPv4 and IPv6 TCP/IP hosts and on hosts configured to support the CIM interface that have the OneCommand Manager application installed (Figure 6). Remote SAN management over TCP/IP sends remote management requests using the TCP/IP access protocol to remote hosts. TCP/IP access enables you to access adapters by using their host IP address or by the name of the host on which they reside.

Figure 6 Discovery Information

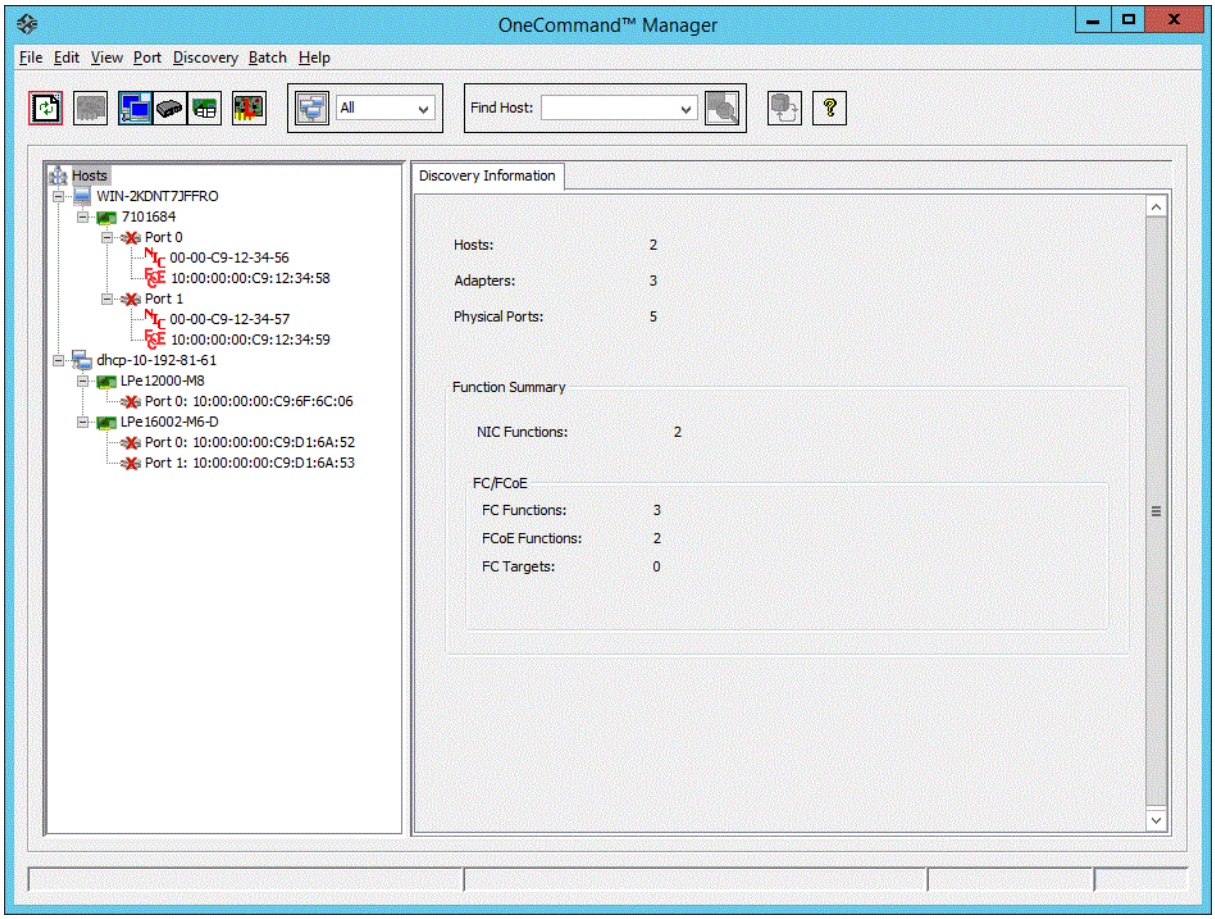

NOTE In Windows, if you are running a firewall, you may have to add the OneCommand Manager application remote server to the firewall's exception list. This remote server's path is:

```
\Program Files\Emulex\Util\Common\rmserver.exe
```

5.1.1 Hosts File

The TCP/IP discovery function of the OneCommand Manager application discovery server relies on a file called the `hosts` file. This plain-text file contains a list of hosts that the utility attempts to discover. The discovery server does not attempt to discover hosts over TCP/IP through any other mechanisms (such as ping sweeps and broadcasts).

The `hosts` file is automatically created or modified when you perform any of the following operations:

- Adding a single host from the **Add Remote Host** window (Figure 7). If the host is discovered, the OneCommand Manager application adds its IP address and name to the `hosts` file.
- Scanning a range of IP addresses for hosts that can be managed. This function is performed in the **Add Remote Hosts** window (Figure 7). For each discovered host, the OneCommand Manager application adds its IP address and name to the `hosts` file.
- Removing a host from the host file using the **Remove Remote Hosts** window. For each removed host, the OneCommand Manager application removes its IP address and name from the `hosts` file.
- Adding or removing a host using the CLI.

5.1.1.1 Manually Editing the `Hosts` File

You can open the `hosts` file with any text editor, modify the contents, and save the file. The name of the `hosts` file is `hbahosts.lst`. After the file is modified and saved, the updated file is used after the next TCP/IP discovery cycle is complete. If the discovery server is running, it does not need to be restarted.

To manually edit the `hosts` file, perform these steps:

1. Locate and open the hosts file.
   - Windows – The file is located on the system drive in the directory `\Program Files\Emulex\Util`.
   - Solaris – The file is located in the directory `/opt/ELXocm`.
   - Linux – The file is located in the directory `/usr/sbin/ocmanager`.
2. Edit the file. Guidelines for editing the file are as follows:
   - Each line of the file starts with an IPv4 or IPv6 address. Following the IP address can be any number of tabs or spaces. These are followed by a # character, zero, or more tabs or spaces, and the name of the host for that IP address. The host name is not required for discovery. Its purpose is to make the file more readable and is used by the OneCommand Manager application to display the host name in the **Remove Remote Hosts** window when the host is not discovered. However, the discovery server only needs the IP address to discover the host.
   - IPv6 address tuples are delimited by colons and can be added in shortened notation as defined by the IPv6 address specification.
   - An IP port number can be specified after the IPv4 address by appending a colon and port number to the address (such as 10.192.80.24:23333).
   - An IP port number can be specified after an IPv6 address by putting the IPv6 address in brackets and following it with a colon and the port number. For example, `[fe80::50f1:832:3ce4:8d30]:23333`
   - Each line in the file can be up to 1023 characters, although this is longer than is typically needed for a host IP address and host name. A line longer than 1023 characters is truncated, possibly causing discovery to not discover some of the hosts.
   - Blank lines are ignored.
3. Save the file.

5.1.1.2 Copying the File

A hosts file on one host can be copied and used on another host. This is useful when there are multiple hosts on the same network running the OneCommand Manager application. For example, after the remote hosts are added to the hosts file on one host, you can copy it to other hosts so you do not need to create another hosts file.

NOTE Because of the line terminator differences between Windows, Solaris, and Linux hosts, `hosts` files cannot be shared between Windows, Solaris, or Linux hosts.

5.1.2 Adding a Single Host

NOTE This option is not available in read-only mode.

The OneCommand Manager application enables you to specify a single TCP/IP host to manage. You can add an RMAPI host or a CIM host using the host name or IP address. If the host is successfully discovered, it is added to the hosts file. If it has not been discovered over FC already, the host and its adapter ports are added to the discovery-tree (Figure 4).

NOTE The OneCommand Manager application must be installed on the remote host.

To add a single host, perform these steps:

1. From the **Discovery** menu, select **TCP/IP > Add Host**. The **Add Remote TCP/IP Host** dialog appears (Figure 7).

Figure 7 Add Remote TCP/IP Host Dialog

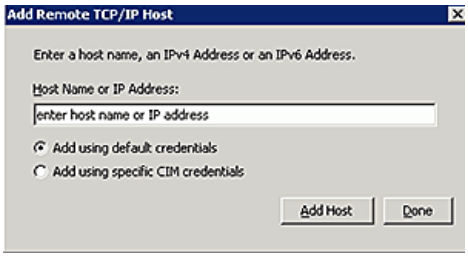

2. Enter the name or the IPv4 or IPv6 address of the host to be added.

NOTE Entering the IP address to identify the host avoids possible name resolution issues. IPv6 address tuples are delimited by colons and can be entered in a shortened form suppressing 0s as defined by the IPv6 address specification.

3. Configure the discovery method:
   - If you want to add the host using the default discovery methods, check **Add using default credentials** and click **Add Host**. A message appears indicating whether the new host was successfully added.
   - If you want to add the new host using specific CIM credentials, check **Add using specific CIM credentials**, modify any additional CIM settings and click **Add Host**. The **Add Remote TCP/IP Host** dialog appears with the default CIM settings (Figure 8).

NOTE Remote CIM hosts can only be managed by Windows client systems.

Figure 8 Add Remote TCP/IP Host Dialog with CIM Credentials

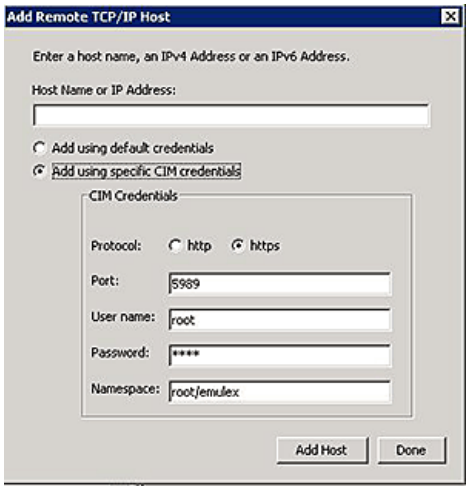

4. Edit the default CIM settings if necessary and click **Add Host**. A message appears indicating the new host was successfully added.

5.1.3 Adding a Range of Hosts (IPv4 Only)

NOTE This option is not available in Strictly Local or Local Plus Management modes.

You can find the TCP/IP-accessed manageable hosts by searching a range of IPv4 addresses. The **Add Range of TCP/IP Hosts** dialog (Figure 9) enables you to build the initial list of TCP/IP accessed manageable hosts.

NOTES

- The ranges of IP addresses are only scanned each time you open the **Add Remote TCP/IP Hosts** dialog and click **Start Discovery**. The ranges are not automatically scanned by the discovery server during its discovery cycles.
- Discovery of VMware (CIM) hosts is only supported on Windows systems. Adding a range of hosts is only supported for IPv4 addresses. It is not supported for IPv6 addresses.
- The OneCommand Manager application must be installed on all remote hosts.

Figure 9 Add Range of TCP/IP Hosts Dialog

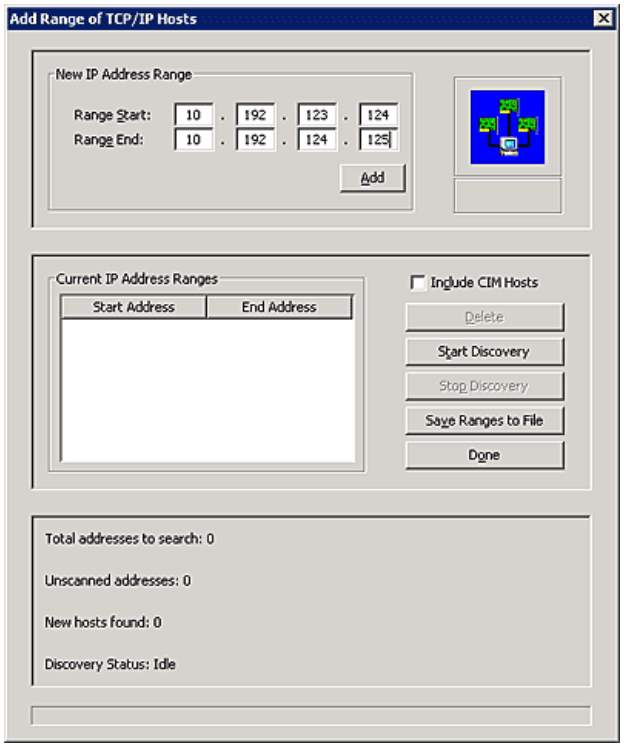

To add a range of remote hosts, perform these steps:

1. From the **Discovery** menu, select **TCP/IP > Add Range of Hosts**. The **Add Range of TCP/IP Hosts** dialog appears (Figure 9).
2. Enter the complete start and end address range (IPv4 only) and click **Add**. The added address range appears in the dialog. Add any additional ranges you want to search.
3. Click **Start Discovery**. If an address is remotely manageable, it is added to the list of addresses that the discovery server attempts to discover. The utility creates a `hosts` file if necessary, and checks each address in the range to determine if the host is available and remotely manageable. The number of addresses (of manageable hosts) discovered is periodically updated on the dialog.

   NOTE The number of hosts found does not correspond directly to the number of hosts added to the discovery-tree (Figure 4). A host can have more than one IP address assigned to it. If multiple IP addresses for a host are discovered during the search, the host is added to the discovery-tree only once.

4. You can save the IP address ranges. Click **Save Ranges to File** to save the specified ranges to a file so that these address ranges appear the next time you use the **Add Range of TCP/IP Hosts** dialog (Figure 9).

5.1.4 Removing Hosts

NOTE This option is not available in read-only mode.

Removing hosts that are no longer discovered improves the operation of the discovery server. For example, you may want to remove a host when it is removed from the network.

To remove hosts, perform these steps:

1. From the **Discovery** menu, select **TCP/IP > Remove Host(s)**. The **Remove Hosts** dialog shows a list of discovered hosts. Any host that is not currently discovered appears in red. Click **Show Undiscovered Hosts Only** to display only currently undiscovered hosts.
2. From the **Remove Hosts** dialog, select the hosts you want to remove. You can select all the displayed hosts by clicking **Select All**.
3. Click **Remove** to remove the selected hosts.

5.2 Configuring Discovery and Default CIM Credentials

Use the **Discovery Settings** dialog (Figure 10) in OneCommand Manager application to configure several discovery server parameters. You can define when to remove previously discovered adapters that are no longer being discovered. You can also define default CIM credentials, such as the protocol, user name, port number, password, and name space.

NOTE Management of CIM hosts is supported only on Windows systems.

A host can have more than one IP address assigned to it. If multiple IP addresses for a host are discovered during the search, the host is added to the discovery-tree (Figure 4) only once. If the same host name appears for more than one host, the adapters of all these hosts are displayed by the OneCommand Manager application as a single host entry.

Figure 10 Discovery Settings Dialog

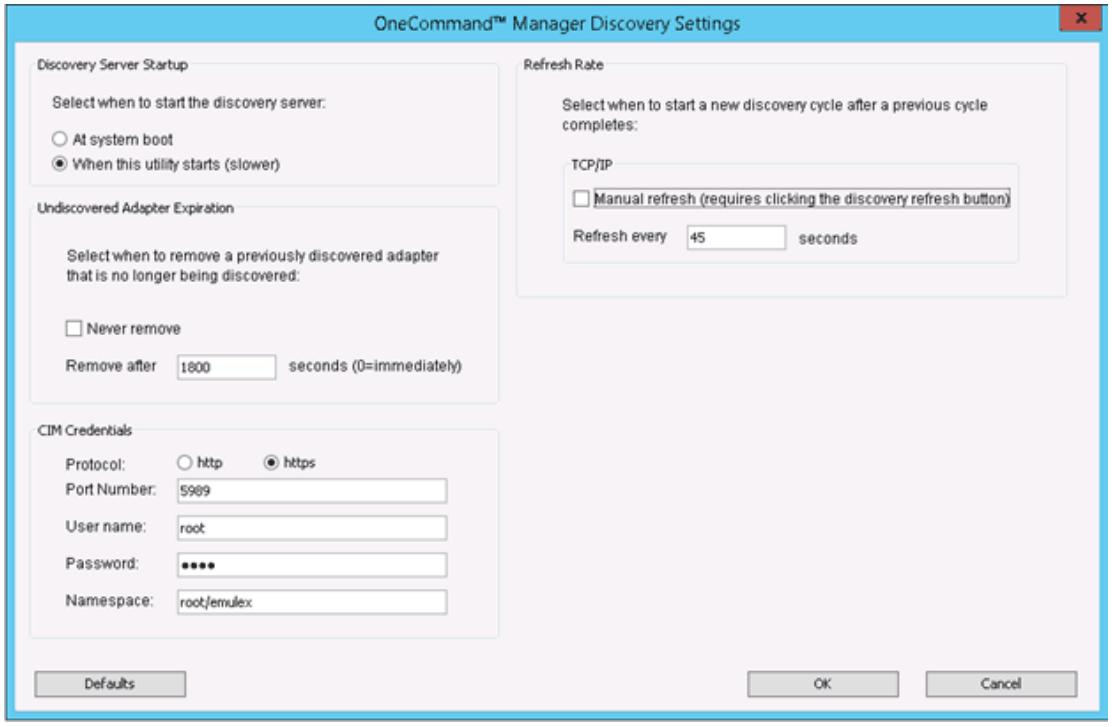

To configure discovery settings, perform these steps:

1. From the **Discovery** menu, select **Modify Settings**. The **OneCommand Manager Discovery Settings** dialog appears (Figure 10).
2. Define the discovery properties you want.

3. The CIM credentials area can be used to set the default CIM credentials that are used to connect to all the ESXi hosts that are managed through the CIM interface.
   - **Protocol:** The HTTP or HTTPS protocol can be used to connect to the VMware hosts.
   - **Port Number:** The default port numbers used for HTTP and HTTPS are `5988` and `5989`, respectively. The port number changes automatically according to the protocol selected. You can also manually change the port number. By default, the HTTP is disabled on `sfcb` in VMware host, so you must use HTTPS to communicate to the VMware host.
   - **User name**: The **User name** field contains the user name with which to connect to the VMware hosts. By default, this is `root`.
   - **Password**: The **Password** field contains the password of the user name that is used to connect to the VMware host.
   - **Namespace**: Namespace is the namespace of the Emulex provider.

     The default namespace is `root/emulex`.
4. Choose the refresh rate settings you want to apply.
5. Click **OK** to apply your changes. Click **Defaults** to return the discovery properties to their default settings.

5.3 Viewing Discovery Information

The Discovery Information page contains a general summary of the discovered elements. The **Host**, **Fabric**, or **Virtual Port** icon, depending upon which view you select, is the root of the discovery-tree, but it does not represent a specific network element. Expanding it reveals all hosts, LUNs, targets, adapter ports, and virtual ports that are visible on the SAN.

To view discovery information, perform these steps:

1. Click the **Hosts**, **Fabrics**, or **Virtual Port** icon at the root of the discovery-tree. Discovered SAN elements appear in the discovery-tree.
2. Select an element from the discovery-tree to learn more about it.

Figure 11 Discovery Information (Host View Selected)

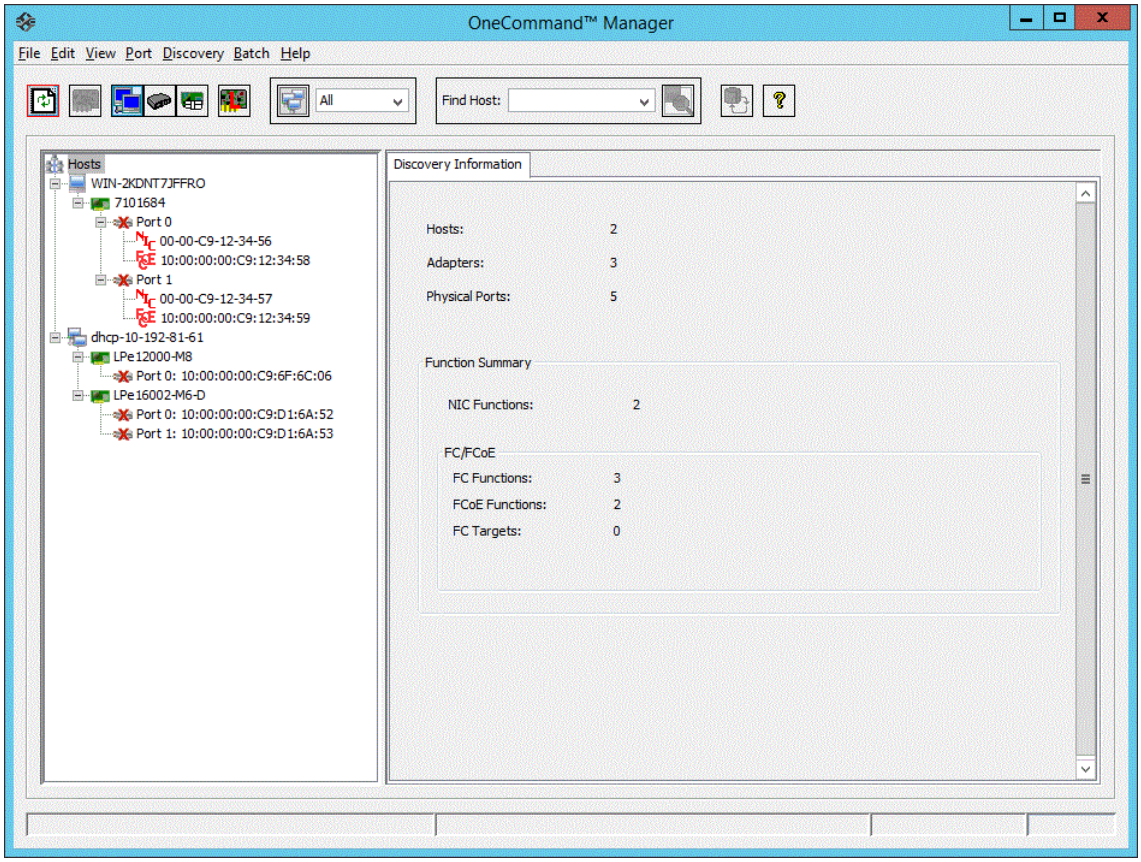

The following **Discovery Information** fields are displayed:

- **Hosts** – The total number of discovered host computers containing manageable Emulex adapters. This number includes servers, workstations, personal computers, multiprocessor systems, and clustered computer complexes.
- **Adapters** – The total number of discovered adapters.
- **Physical Ports** – The number of discovered physical ports that can be managed by this host.
- **Function Summary** – Listed by protocol, the total number of discovered functions and targets.

Chapter 6: Managing Hosts

This section describes viewing host information, managing host groups, and searching for hosts.

6.1 Viewing Host Information

Two tabs show host information: the **Host Information** tab (Figure 12) and the **Host Driver Parameters** tab (Figure 44). The **Host Information** tab is read-only. The **Host Driver Parameters** tab enables you to view and define adapter driver settings for a specific host. See Section 8.1.19.2, Host Driver Parameters Tab, for more information about the **Host Driver Parameters** tab.

To view the **Host Information** tab, perform these steps:

1. Perform one of the following tasks:
   - From **View** menu, click **Group Adapters by Host Name**.
   - From the toolbar, click **Group Adapters by Host Name**.
2. Select a host in the discovery-tree.
3. Select the **Host Information** tab (Figure 12).

Figure 12 Host Information Tab

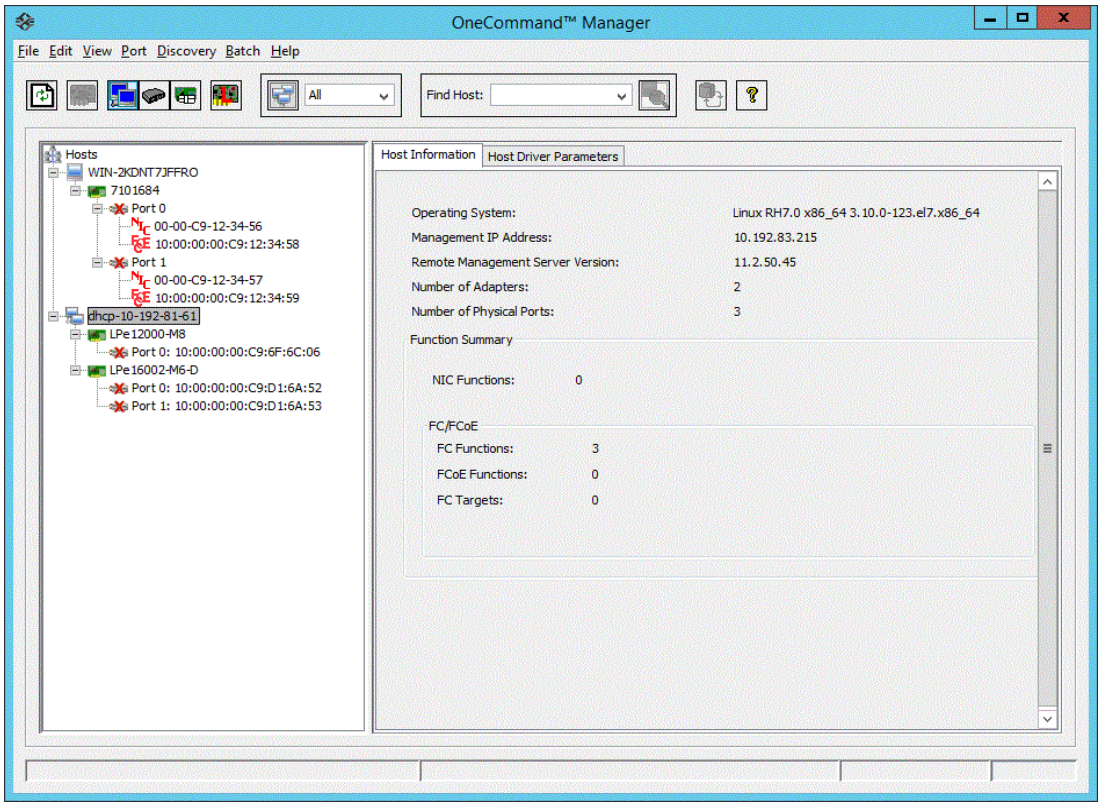

The **Host Information** tab (Figure 12) displays the following fields:

- **Operating System** – Details about the installed operating system.

- **Management IP Address** – The Management IP Address field displays the host's IP address; for example, 138.239.82.131. **Local Host** is displayed if you selected the host from which you are actually running the OneCommand Manager application.
- **Remote Manager Server Version** – The version of the OneCommand Manager application server that is running on the host.
- **Number of Adapters** – The number of adapters installed in the host.
- **Number of Physical Ports** – The number of discovered physical ports that can be managed by this host.
- **CIM Provider Version** – If the host is being managed using the CIM interface, the **CIM Provider Version** field displays the version of the Emulex CIM Provider that is running on the remotely managed system.

NOTE The **CIM Provider Version** field only appears if the host is managed through the CIM interface.

6.1.1 Function Summary Area

The Function Summary area has the following information:

- NIC Functions – The number of NIC functions running on the discovered adapters on this host (LPe16202/OCe15100 adapters only).
- FC Functions – The number of FC functions running on the discovered adapters on this host.
- FCoE Functions – The number of FCoE functions running on the discovered adapters on this host (LPe16202/OCe15100 adapters only).
- FC Targets – The number of FC targets discovered on the FC or FCoE functions on this host.
- VPorts – The number of discovered FC virtual ports that can be managed by this host (not supported on VMware ESXi servers being managed through the CIM interface).

6.2 Viewing Host Grouping Information

The **Host Group Information** tab (Figure 13) displays information about the selected host group, such as the group name and the total number of hosts. See Section 6.3, Grouping Hosts, to learn about creating host groups.

NOTE Host grouping is not supported for VMware.

To view host grouping information, from the discovery-tree (Figure 4), select the host group whose information you want to view.

Figure 13 Host Group Information Tab

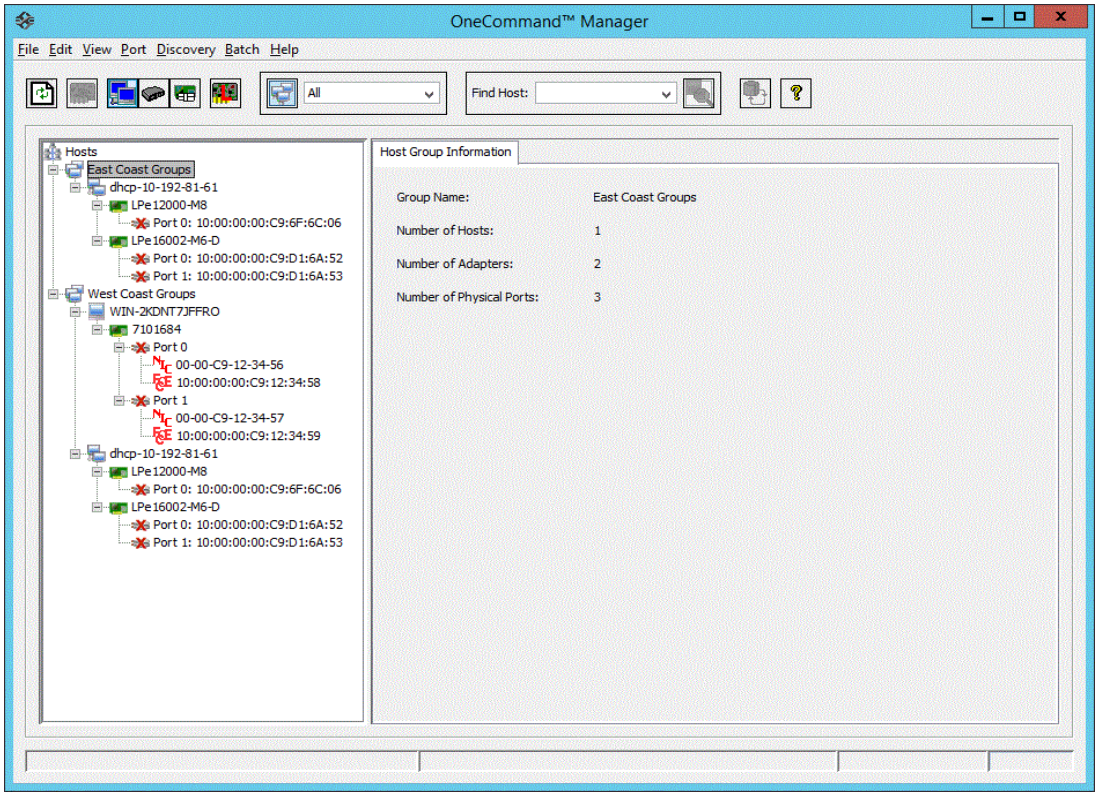

The following Host Group Information fields are displayed:

- **Group Name** – The name of the selected group.
- **Number Hosts** – The total number of hosts assigned to the group.
- **Number of Adapters** – The total number of discovered adapters in the group.
- **Number of Physical Ports** – The total number of ports in the group.

6.3 Grouping Hosts

The OneCommand Manager application enables you to assign related hosts to host groups. Typically, hosts within the same host group share some common functions, or they may simply reside within the same organizational unit within an enterprise, such as a Payroll group or a Shipping/Receiving group.

You can display the hosts in the discovery-tree (Figure 4) in either a group-centric format or in the host-based flat format. The Host grouping capability is available in the Host view, Vport view, or Fabric view mode.

NOTE The same fabric may appear under more than one host group. For example, some ports on the fabric may be attached to ports and hosts in one host group, and other ports on the same fabric may be attached ports and hosts in a different host group.

You can also perform batch operations, such as firmware download and driver parameter updates on a selected set of groups. See Section 9.2, Updating Firmware for Multiple Adapters, for more information.

NOTE Grouping hosts is not supported on VMware.

To display all hosts without grouping, perform one of the following tasks:

- From the **View** menu, clear **Show Groups**.
- From the toolbar unclick **Show Host Groups**.

To display all hosts groups, perform these steps:

1. Perform one of the following tasks:
   - From the **View** menu, check **Show Groups**.
   - From the toolbar click **Show Host Groups**.
2. From the **Available Host Group** list, choose **All**.

To display all hosts assigned to a particular group, perform these steps:

1. Perform one of the following tasks:
   - From the **View** menu, check **Show Groups**.
   - From the toolbar click **Show Host Groups**.
2. From the **Available Host Group** list, choose the group whose hosts you want to view.

6.3.1 Managing Host Groups

Use the **Host Group Management** dialog (Figure 14) to create and delete host groups, add and remove hosts, and restore host groups.

NOTE Managing host groups is not supported on VMware.

Figure 14 Host Group Management Dialog

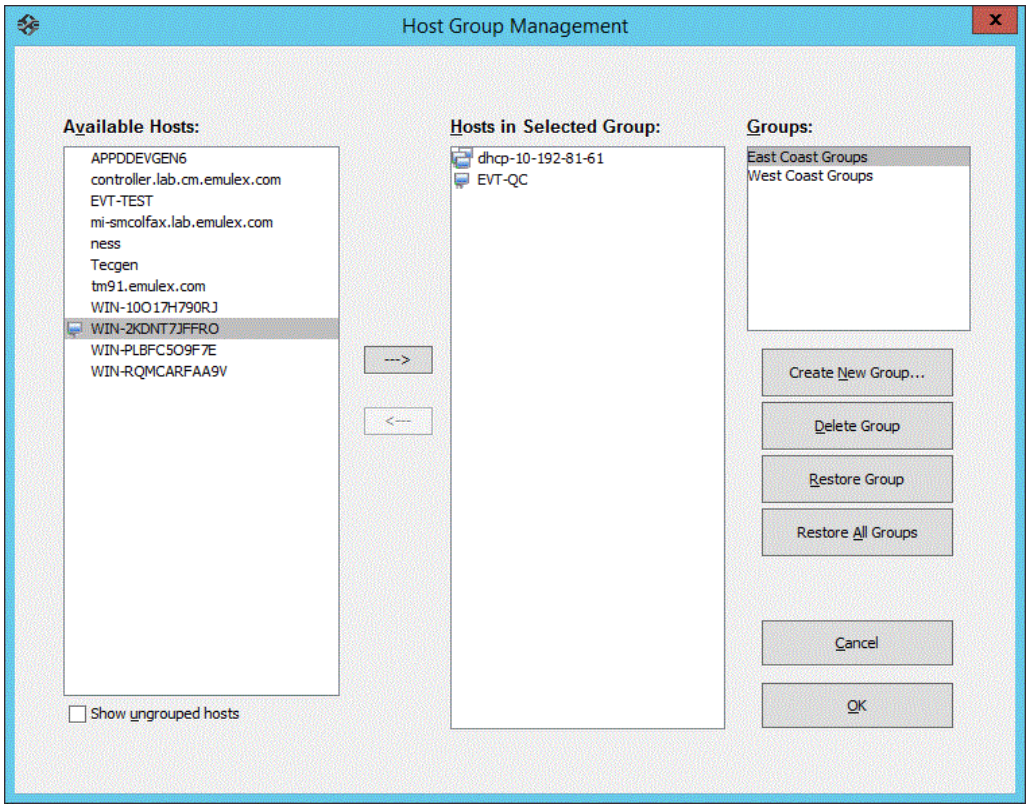

The following Host Group Management fields are displayed:

- **Available Hosts** – The list of hosts that can be added to a host group. You can select a host and right-click to see its group assignments.
- **Show ungrouped hosts** – If selected, displays only hosts that are currently assigned to a host group.
- **Hosts in Selected Group** – The list of hosts assigned to the currently selected host group.
- **Groups** – The list of the currently defined host groups. If you select a group in this list, its host members appear in the Hosts in Selected Group list.

The following icons are used in the **Host Group Management** window:

- Indicates the host is currently assigned to a single host group.
- Indicates the host is currently assigned to multiple host groups.

6.3.2 Creating a Host Group

To create a new host group, perform these steps:

1. From the **View** menu, select **Manage Groups**. The **Host Group Management** dialog appears (Figure 14).
2. Click **Create New Group**. The **Create New Host Group** dialog is displayed (Figure 15).

Figure 15 Create New Host Group Dialog

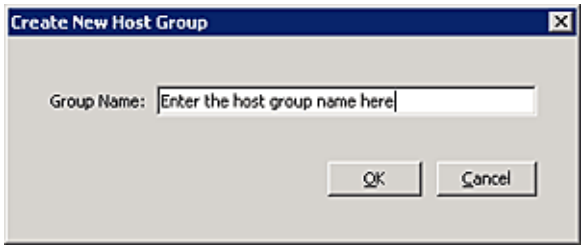

3. Enter the name of the group you want to create and click **OK**. The new group appears in the Groups list on the **Host Group Management** dialog.

6.3.3 Deleting a Host Group

To delete a host group, perform these steps:

1. From the **View** menu, select **Manage Groups**. The **Host Group Management** dialog appears (Figure 14).
2. From the **Groups** list, select the group you want to delete. The **Host Group Management warning** dialog appears (Figure 16).

Figure 16 Host Group Management Warning Dialog

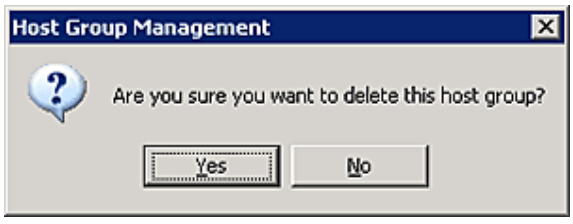

3. Click **Yes** to delete the selected host group.

6.3.4 Adding a Host to a Host Group

To add a host to a group, perform these steps:

1. From the **View** menu, select **Manage Groups**. The **Host Group Management** dialog appears (Figure 14).
2. From the **Groups** list, select the group to which you want to add the host.
3. From the **Available Hosts** list, select the host you want to add (or select multiple hosts by using **Ctrl**-Click or **Shift**-Click), and click the Right Arrow. The selected host is removed from the Available Hosts list and is added to the Hosts in Selected Group list.
4. Click **OK** to commit your changes. The discovery-tree (Figure 4) displays the new configuration.

6.3.5 Removing a Host from a Host Group

To remove a host from a host group, perform these steps:

1. From the **View** menu, select **Manage Groups**. The **Host Group Management** dialog appears (Figure 14).
2. From the **Groups** list, select the group containing the host you want to remove.
3. From the **Hosts in Selected Group** list, select the host you want to remove and click the Left Arrow. The selected host is removed from the Hosts in Selected Group list and is added to the Available Hosts list.
4. Click **OK** to commit your changes. The discovery-tree (Figure 4) displays the new configuration.

6.3.6 Restoring a Host Group

To restore a host group, perform these steps:

1. From the **View** menu, select **Manage Groups**. The **Host Group Management** dialog appears (Figure 14).
2. Click **Restore Group** to return the configuration settings for the currently selected host group to those in use when the dialog was opened.

NOTE If the currently selected group was created during the current configuration session, clicking **Restore Group** deletes the new group name.

6.3.7 Restoring All Host Groups

To restore all host groups, perform these steps:

1. From the **View** menu, select **Manage Groups**. The **Host Group Management** dialog appears (Figure 14).
2. Click **Restore All Groups** to return the entire host group configuration to the state that existed when the dialog was opened. All host group assignments are returned to their original configuration. Newly added host groups yet to be committed are removed, and host groups that were deleted are restored.

6.3.8 Exporting Host Grouping Configurations

To export the host grouping configuration to a remote host, you must copy the various host group configuration files from the host on which the configuration was created to the remote host. Copy the entire contents of the `config/hostgroups` subdirectory under the OneCommand installation directory to the equivalent location on the remote system.

The hostgroups configuration file locations for the supported platforms are:

- Windows: `<InstallationDriveLetter>:\Program Files\Emulex\Util\Config\hostgroups`

- Linux: `/usr/sbin/ocmanager/config/hostgroups`
- Solaris: `/opt/ELXocm/config/hostgroups`

The host group configuration files are completely interchangeable between different operating systems. For example, the host group configuration files created on a Solaris host can be copied directly to a Linux or Windows host, with no conversion required.

6.4 Searching for Hosts in the Discovery-Tree

The OneCommand Manager application enables you to search the discovery-tree (Figure 4) for a particular host by the host's name. If the specified host name is found, the discovery-tree scrolls up or down to bring the desired host name into view.

This capability is especially useful when you are searching for a host in a large installation with hundreds or thousands of hosts. It is also helpful in the Fabric view mode because the ports on a specific host may be dispersed among several fabrics, making the ports on that host difficult to find in the discovery-tree (Figure 4).

To search for a host, perform these steps:

1. Perform one of the following tasks:
   - From the **Edit** menu, select **Find** and enter the name of the host you are searching for into the **Find Host** field.
   - From the toolbar, enter the name of the host you are searching for into the **Find Host** field.
2. From the toolbar, click  **Find Host** or press **Enter** on the keyboard.

   The host you are searching for is highlighted in the discovery-tree (Figure 4).

   The **Find Next** option on the **Edit** menu, or pressing **F3**, enables you to continue searching for more instances of the name you specified.

Chapter 7: Managing Adapters and Ports

This section describes the various adapter and port management functions you can perform using the OneCommand Manager application.

7.1 FC Adapters

This section pertains to viewing FC adapter and port information, including LPe16202/OCe15100 adapters in FC mode.

7.1.1 Viewing FC Adapter Information

When you select an FC adapter from the discovery-tree (Figure 4), the **Adapter Information** tab (Figure 17) contains general attributes associated with the selected FC adapter.

To view FC adapter information, perform these steps:

1. Select the **Host**, **Fabric,** or **Virtual Ports** view.
2. Select an FC adapter in the discovery-tree. The **Adapter Information** tab appears (Figure 17).

Figure 17 FC Adapter Information Tab

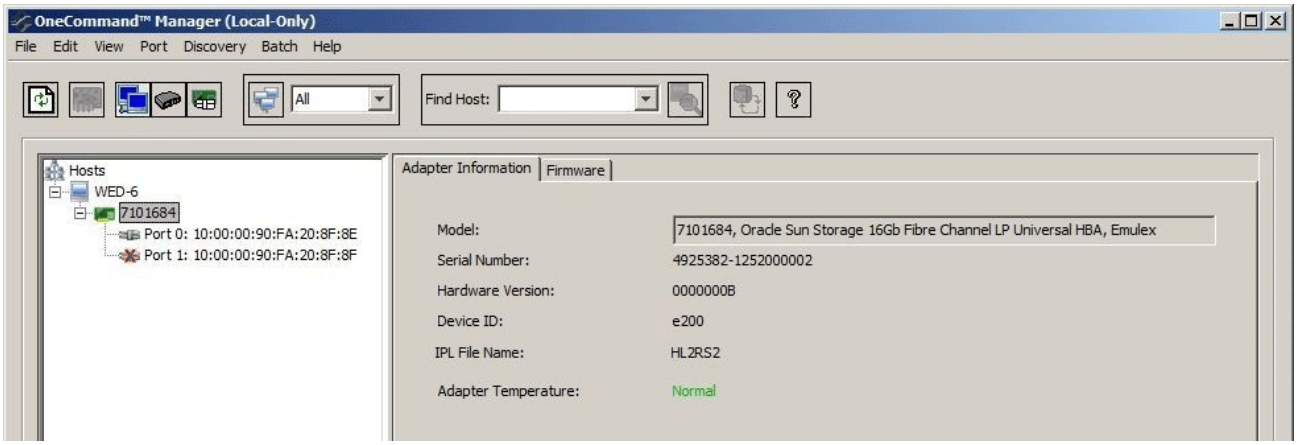

The following Adapter Information fields are displayed:

- **Model** – The complete model name of the adapter.
- **Serial Number** – The manufacturer's serial number for the adapter.
- **Hardware Version** – The JEDEC ID.
- **Device ID** – The default device ID for the selected adapter.
- **IPL File Name** – The IPL file name for the selected adapter.
- **Adapter Temperature** – If the adapter's temperature is not available, **Not Supported** is displayed. If supported by the adapter, this field displays the adapter's temperature and one of the following temperature-related status messages:
  - **Normal:** The adapter's temperature is within normal operational range.
  - **Warning:** The adapter's temperature is beyond normal operational range. If the temperature continues to increase, the adapter shuts down. You must determine the cause of the temperature issue and fix it

immediately. Check for system cooling issues. Common causes of system cooling issues include clogged air filters, inoperative fans, and air conditioning issues that cause high ambient air temperatures.

— **Exceeds operational range – Adapter stopped:** The temperature has reached critical limit, forcing the adapter to shut down. You must determine the cause of the temperature issue and fix it before resuming operation. Check for system cooling issues. Common causes of system cooling issues include clogged air filters, inoperative fans, and air conditioning issues that cause high ambient air temperatures.

After the system overheating issue is resolved and the adapter has cooled down, reboot the system, or if the system supports hot swapping, cycle the power of the adapter slot.

7.1.2 Viewing FC Port Information

When you select an FC port from the discovery-tree, the **Port Information** tab (Figure 18) contains general attributes associated with the selected FC adapter.

To view FC Port information, perform these steps:

1. Select the **Host** or **Fabric** view.
2. Select an FC port in the discovery-tree.
3. Select the **Port Information** tab (Figure 18).

Figure 18 FC Port Information Tab

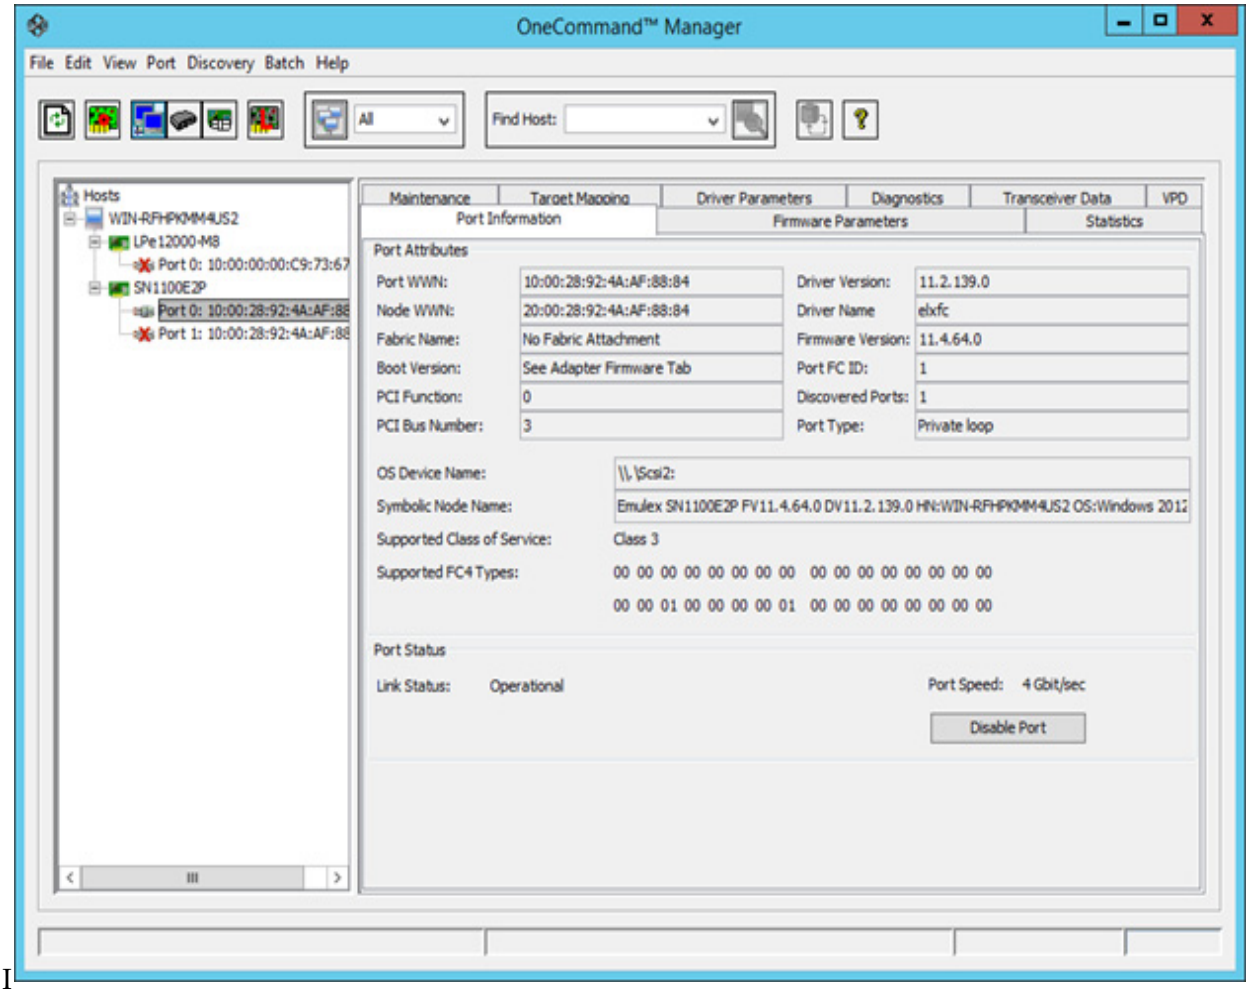

I

The following Port Information fields are displayed:

- Port Attributes area:
  - **Port WWN** – The WWPN of the adapter.
  - **Node WWN** – The WWNN of the adapter.
  - **Fabric Name** or **Host Name** – The **Fabric Name** field is displayed in the Host view. This is a 64-bit worldwide unique identifier assigned to the fabric. The **Host Name** field is displayed in the Fabric view and is the name of the host containing the adapter.
  - **Boot Version** – The version of boot code installed on the selected adapter port. If the boot code is disabled, the field displays **Disabled**.
  - **PCI Function** – The PCI function number assigned by the system.
  - **PCI Bus Number** – The PCI bus number assigned to the FC function.
  - **Driver Version** – The version of the driver installed for the adapter.
  - **Driver Name** – The executable file image name for the driver as it appears in the Emulex driver download package.
  - **Firmware Version** – The version of Emulex firmware currently active on the adapter port.
  - **Port FC ID** – The FC ID for the selected adapter port.
  - **Discovered Ports** – The number of mapped and unmapped ports found during discovery by the Emulex adapter driver. The mapped ports are targets and the unmapped ports are non-targets, such as switches or adapters.
  - **Port Type** – The FC type of the selected adapter's port (not available if the port link is down).
  - OS Device Name – The platform-specific name by which the selected adapter is known to the operating system.
  - **Symbolic Node Name** – The FC name used to register the driver with the name server.
  - **Supported Class of Service** – A frame delivery scheme exhibiting a set of delivery characteristics and attributes. Three classes of service include:
    - **Class 1** – Provides a dedicated connection between a pair of ports with confirmed delivery or notification of non-delivery.
    - **Class 2** – Provides a frame switched service with confirmed delivery or notification of non-delivery.
    - **Class 3** – Provides a frame switched service similar to Class 2 but without notification of frame delivery or non-delivery.
  - **Supported FC4 Types** – A 256-bit (8-word) map of the FC-4 protocol types supported by the port containing the selected adapter.
- Port Status area
  - **Link Status** – The status of the link on the selected adapter port.
  - **Port Speed** – The current port speed of the selected adapter port.

7.1.2.1 Enabling and Disabling an FC Port

You can enable or disable an FC port from the **Port Information** tab. When you disable an FC port, you disable all functions for the port.

CAUTION Do not disable a boot port; this could result in data loss or corruption.

NOTE Ensure that there is no I/O traffic on the port before disabling it.

To enable or disable an FC port, from the **Port Information** tab, click **Enable Port** or **Disable Port**.

7.1.3 Viewing FC Firmware Parameters

To view firmware parameters, select the **Firmware Parameters** tab (Figure 19).

Figure 19 FC Firmware Parameters Tab

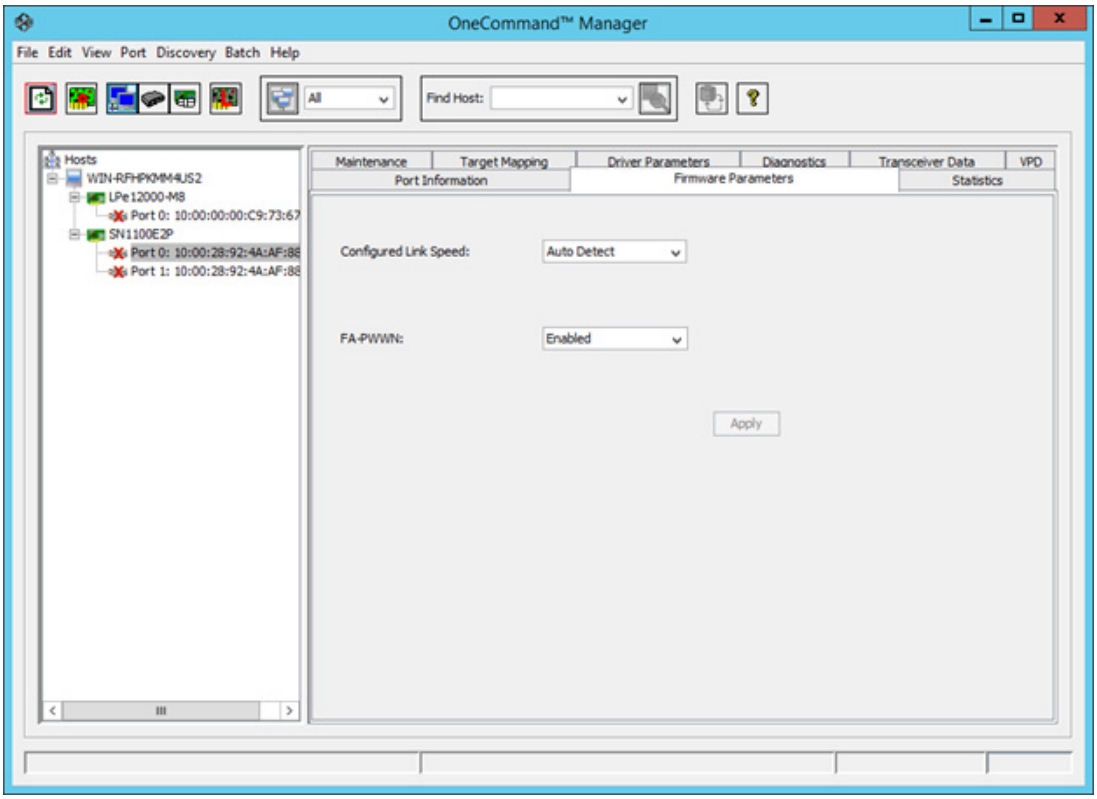

The following fields are displayed:

- **Configured Link Speed** – Link speeds that are supported on the port. The list varies depending on the adapter type. The list also includes an Auto Detect option, which indicates that the link speed should be auto-negotiated.

  NOTE If an installed adapter does not support forced link speeds, the Configured Link Speed settings and the **Apply** button are not shown.

- **FA-PWWN** – Displays the FA-PWWN status. **Disabled** is the default setting.

  NOTES
  - FA-PWWN is not supported on VMware ESXi servers being managed through the CIM interface.
  - The switch must support FA-PWWN. Refer to the documentation that accompanied the switch for instructions on configuring FA-PWWN on the switch.
  - After enabling or disabling FA-PWWN, the port must be reset for changes to take effect.
  - When a new WWPN is assigned using FA-PWWN, persistently stored configuration information associated with the original WWPN, such as driver parameters and LUN frame priority settings, is not applied to the newly assigned WWPN. The configuration information associated with the original WWPN must be reconfigured for the new WWPN.

7.1.3.1 Configuring Link Speed

To configure a link speed, perform these steps:

1. Select the **Firmware Parameters** tab (Figure 19).
2. Select a link speed from the **Configured Link Speed** drop-down list (Figure 20).

Figure 20 Port Status Dialog

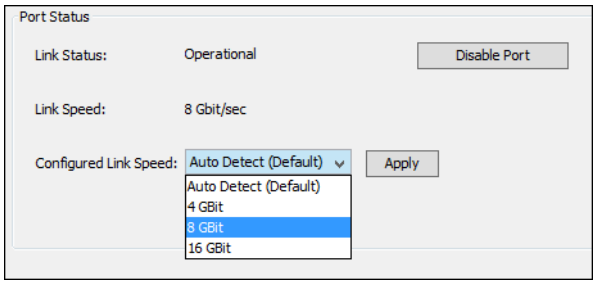

3. Click **Apply** to set the new link speed. The **Apply** button will only be enabled if the currently selected link speed does not match the currently configured speed.

   If the speed has been set successfully, the following message is displayed stating that the port must be reset to activated the new speed setting.

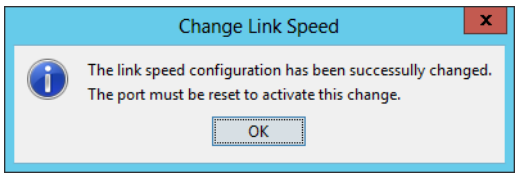

4. Click **OK**.
5. From the toolbar, click  **Reset**. **Port**.

In some situations, the currently configured link speed is not in the supported speed list for the port, which can occur if a new SFP is installed that supports a different set of link speeds than the previously installed SFP. If the currently configured link speed is not in the supported speed list, the following message is displayed:

```
Warning: The currently configured port speed is not a valid supported speed.
Please select a link speed and click Apply.
```

The **Apply** button remains enabled until you select a valid port speed.

If the installed SFP is not supported by the adapter, you cannot configure a link speed. If this is attempted, the following message is displayed:

```
Unsupported optics installed.
```

If an adapter does not support forced link speed, the **Firmware Parameters** tab does not show a Link Speed drop-down list.

7.1.3.2 Enabling and Disabling the FA-PWWN

NOTES

- Fabric Assigned Port Word Wide Name (FA-PWWN) is not supported on VMware ESXi servers being managed through the CIM interface.

- The switch must support FA-PWWN. Refer to the documentation that accompanied the switch for instructions on configuring FA-PWWN on the switch.
- After enabling or disabling FA-PWWN, the port must be reset for changes to take effect.
- When a new WWPN is assigned using FA-PWWN, persistently stored configuration information associated with the original WWPN, such as driver parameters and LUN frame priority settings, is not applied to the newly assigned WWPN. The configuration information associated with the original WWPN must be reconfigured for the new WWPN.

To enable or disable FA-PWWN, perform these steps:

1. Select the **Firmware Parameters** tab (Figure 19).
2. Select **Enable** or **Disable** from the **FA-PWWN** dropdown list.
3. Click **Apply**. The **Change Firmware Parameter** dialog appears (Figure 21).
4. Click **OK**.

Figure 21 Change Firmware Parameter Dialog

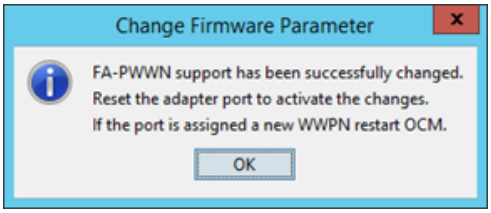

5. From the toolbar, click **Reset**. **Port**.
6. Restart the OneCommand Manager application.

NOTE An **error** dialog notifies you if the FA-PWWN change was unsuccessful (Figure 22).

Figure 22 FA-PWWN Failure Dialog

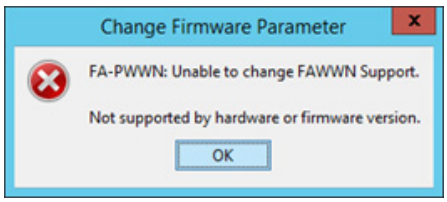

7.2 LPe16202/OCe15100 Adapters in NIC+FCoE Mode

This section describes viewing port and adapter information and managing LPe16202/OCe15100 adapters in NIC+FCoE mode.

7.2.1 Viewing LPe16202/OCe15100 Adapter Information

When you select an LPe16202/OCe15100 adapter in NIC+FCoE mode from the discovery-tree, the **Adapter Information** tab contains general attributes associated with the selected adapter.

To view general adapter information, perform these steps:

1. Select the **Host** view.
2. Select an LPe16202/OCe15100 adapter in the discovery-tree. The **Adapter Information** tab appears (Figure 23).

Figure 23 Adapter Information Tab

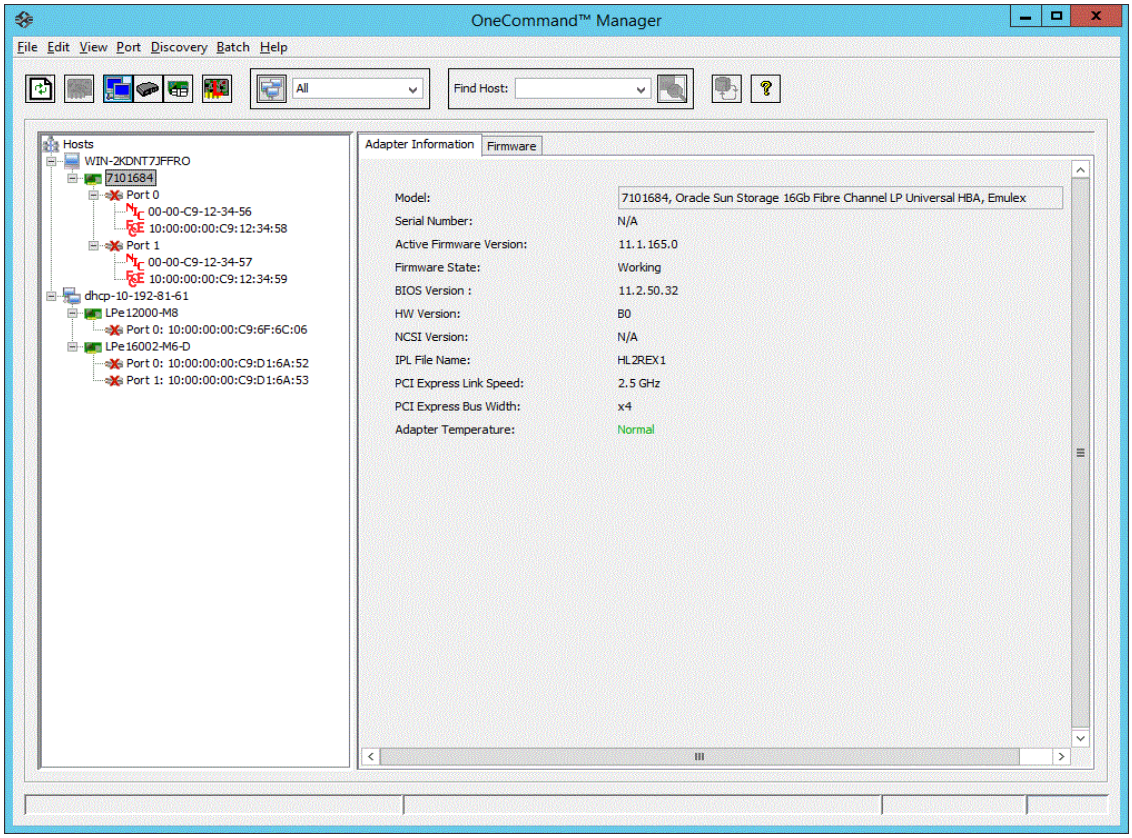

The following Adapter Information fields are displayed:

- **Model** – The model of the selected adapter.
- **Serial Number** – The serial number of the selected adapter.
- **Active Firmware Version** – The version of the firmware running on the selected adapter.
- **Firmware State** – The condition of the firmware.
- **BIOS Version** – The version of the BIOS in use.
- **HW Version** – The hardware version of the selected adapter.
- **NCSI Version** – The NCSI version.
- **IPL File Name** – The name of the IPL file currently loaded.
- **PCI Express Link Speed** – The speed of the PCI bus in which the adapter running.
- **PCI Express Bus Width** – The number of lanes for the slot in which the adapter is running.
- **Adapter Temperature** – If the adapter's temperature is not available, **Not Supported** is displayed. If supported by the adapter, this field displays the adapter's temperature and one of the following temperature-related status messages:
  - **Normal**: The adapter's temperature is within normal operational range.
  - **Warning**: The adapter's temperature is beyond normal operational range. If the temperature continues to increase, the adapter shuts down. You must determine the cause of the temperature issue and fix it

immediately. Check for system cooling issues. Common causes of system cooling issues include clogged air filters, inoperative fans, and air conditioning issues that cause high ambient air temperatures.

— **Exceeds operational range**: The temperature has reached critical limit. You must determine the cause of the temperature issue and fix it before resuming operation. Check for system cooling issues. Common causes of system cooling issues include clogged air filters, inoperative fans, and air conditioning issues that cause high ambient air temperatures.

After the system overheating issue is resolved and the adapter has cooled down, reboot the system, or if the system supports hot swapping, cycle the power of the adapter slot.

7.2.2 Viewing Firmware Information

Use the **Firmware** tab (Figure 24) to download firmware and create diagnostic dumps for LPe16000-series, LPe31000-series, and LPe32000-series adapters.

NOTE For LPe12000-series adapters, see Section 8.1.8, Viewing FC Maintenance Information.

To view firmware information, perform these steps:

1. Select the **Host** view.
2. Select an adapter in the discovery-tree (Figure 4).
3. Select the **Firmware** tab (Figure 24).

Figure 24 Firmware Tab

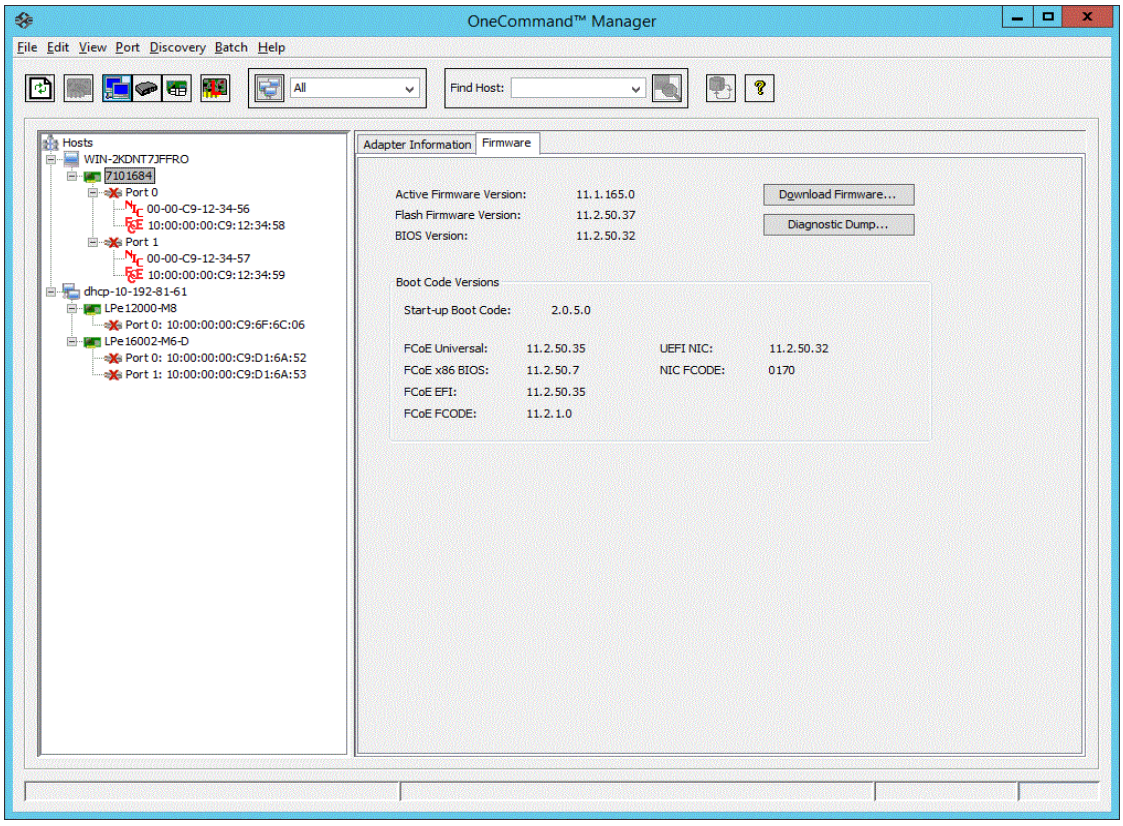

The following Firmware fields are displayed:

- **Active Firmware Version** – The firmware version currently being used by the adapter.
- **Flash Firmware Version** – The flash firmware version currently being used by the adapter.
- **BIOS Version** – The version of the BIOS currently being used by the adapter.
- Boot Code Versions area:
  - **Startup-up Boot Code** – The boot code version currently being used by the adapter.
    - This is the version of the code that boots the adapter. It has no relation to the FC or PXE boot code versions.
  - **FCoE Universal** – The combined flash image that includes three system-specific FCoE Boot images (Open Boot, x86, and EFI 2.0).
  - **FCoE x86 BIOS** – The single flash image containing x86 Boot for FCoE only.
  - **FCoE EFI** – The single flash image containing EFI for FCoE only.
  - **FCoE FCODE** – The single flash image containing Open Boot FCode for FCoE only.
  - **UEFI NIC** – The single flash image containing UEFI for NIC and PXE Boot.
  - **NIC FCODE** – The single flash image containing FCode for NIC only.

NOTE The **Firmware** tab buttons are not available in read-only mode.

See Section 9, Updating Adapter Firmware, for information on updating firmware.

See Section 12.1.6, Creating Diagnostic Dumps, for information about performing a diagnostic dump.

7.2.3 Viewing Physical Port Information

The **Physical Port Info** tab (Figure 25) contains a general summarization of the PCI functions under that physical port and the current physical port status. The tab is available only for LPe16202/OCe15100 adapters in NIC+FCoE mode.

To view physical port information, perform these steps:

1. Select the **Host** view.
2. Select a port on an LPE16202/OCe15100 adapter in NIC+FCoE mode in the discovery-tree.
3. Select the **Physical Port Info** tab (Figure 25).

Figure 25 Physical Port Info Tab

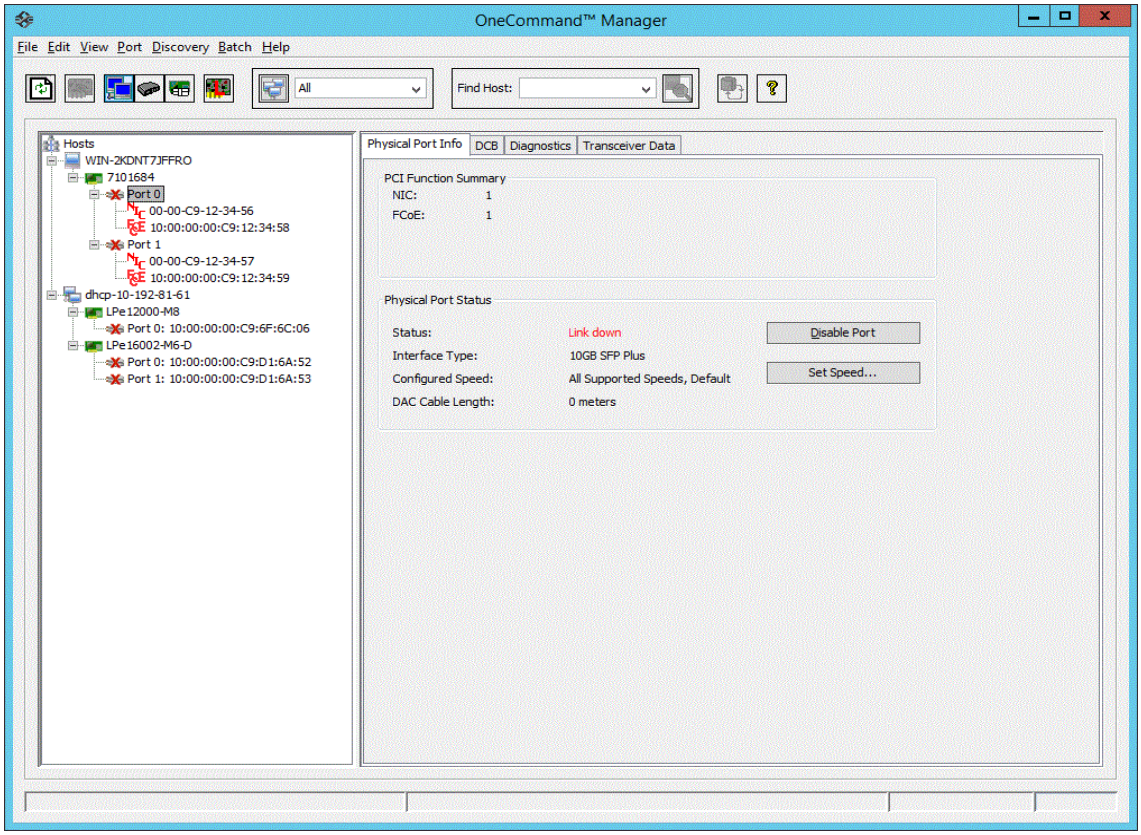

7.2.3.1 Enabling and Disabling Physical Ports

From the **Physical Port Info** tab (Figure 25), you can enable or disable the physical port. If you disable a physical port, you disable all functions, such as NIC and FCoE, for the port. Disabled ports appear in the discovery-tree as a black **Port** icon.

NOTE You cannot disable a port if PXE Boot is enabled.

To enable or disable a physical port, perform these steps:

1. In the discovery-tree (Figure 4), select the physical port you want to enable or disable.
2. Select the **Physical Port Info** tab (Figure 25).
3. Click either **Enable Port** or **Disable Port**.

NOTE Disabling a port does not remove the functions from the PCIe bus. It only disables the operation of the port on the adapter.

7.2.3.2 Setting Port Speed and DAC Cable Length

The **Physical Port Info** tab enables you to set port speed and DAC cable lengths.

To set the port speed, perform these steps:

1. From the discovery-tree (Figure 4), select the adapter port whose speed you want to change.
2. Click **Set Speed** on the **Physical Port Info** tab (Figure 25). The **Change Port Speed** dialog appears (Figure 26).

Figure 26 Change Port Speed Dialog

3. Set the desired mode and port speed. The (port speed) mode setting can one of the following:
   - **Default** – Sets the port speed to the factory default configured speed of the adapter (from IPL).
   - **Force** – Sets the speed to a single speed value.
   - **Auto-negotiate** – Requires a speed setting to single speed or multiple speed choices that the port can use to auto-negotiate the port speed with the switch port.

   NOTE If the adapter's port speed setting and the switch's port speed setting conflict, the link does not come up.
4. If you set the Mode to Force and the Speed to 10 GB SFP+ or QSFP+, you must set the DAC cable length in the range of 0–10 meters.
   - For 10GB SFP+, the length is the actual DAC cable length, and 0 indicates an optical cable.
   - For the QSFP+ module type, a length of 0 indicates an optical cable, and any non-zero length indicates a DAC cable (that is, the length does need to be the actual DAC cable length).
5. Click **OK**.

7.2.4 Viewing Transceiver Information

When you select an LPe16202/OCe15100 adapter port from the discovery-tree (Figure 4), the **Transceiver Data** tab (Figure 27) enables you to view transceiver information such as vendor name, serial number, part number, and so on. If the adapter/transceiver does not support some or all of the transceiver data, the fields display N/A.

To view transceiver information, perform these steps:

1. Select the **Host** or **Fabric** view.
2. In the discovery-tree (Figure 4), select the LPe16202/OCe15100 adapter port whose transceiver information you want to view.
3. Select the **Transceiver Data** tab (Figure 27).

Figure 27 Transceiver Data Tab

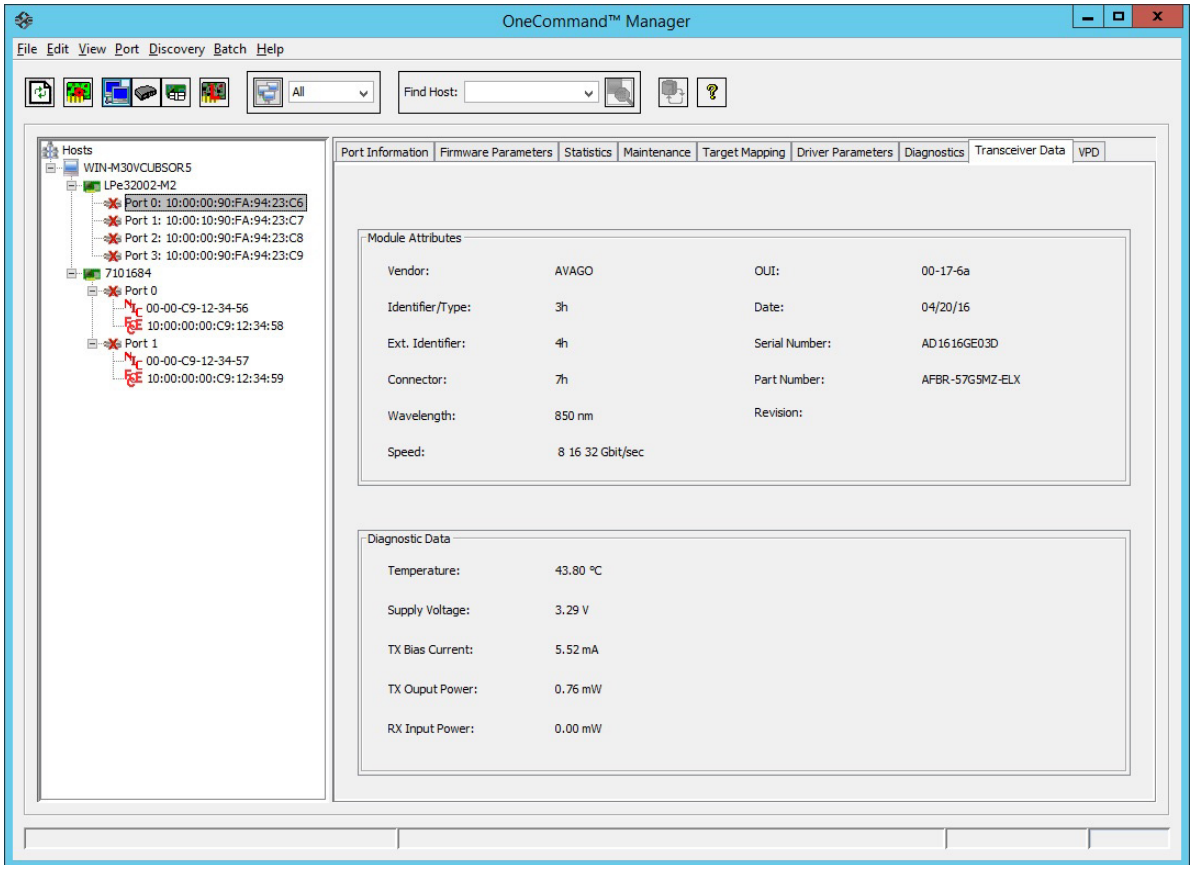

The following Transceiver Data fields are displayed:

- Module Attributes area:
  - **Vendor** – The name of the vendor.
  - **Identifier/Type** – The identifier value that specifies the physical device described by the serial information.
  - **Ext. Identifier** – Additional information about the transceiver.
  - **Connector** – The external optical or electrical cable connector provided as the media interface.
  - **Wavelength** – The nominal transmitter output wavelength at room temperature.
  - **Speed** – The speed, or speeds, at which the selected port can run.
  - **OUI** – The vendor's OUI. It is also known as the IEEE Company Identifier for the vendor.
  - **Date** – The vendor's date code in the MM/DD/YY format.
  - **Serial Number** – The serial number provided by the vendor.
  - **Part Number** – The part number provided by the SFP vendor.
  - **Revision** – The vendor revision level.
- Diagnostic Data area:
  - **Temperature** – The internally measured module temperature.
  - **Supply Voltage** – The internally measured supply voltage in the transceiver.
  - **TX Bias Current** – The internally measured TX bias current.
  - **TX Output Power** – The measured TX output power.
  - **RX Input Power** – The measured RX input power.

Chapter 8: Managing Protocols

This section describes how to manage the available protocols.

8.1 FC and FCoE Functions

This section describes managing FC/FCoE ports.

8.1.1 Viewing FCoE Function Information (LPe16202/OCe15100 Adapters)

When you select an FCoE function from the discovery-tree, the **Port Information** tab (Figure 28) contains general attributes associated with the selected FCoE port.

To view FCoE Port information, perform these steps:

1. Select the **Host** or **Fabric** view.
2. Select an FCoE function in the discovery-tree.
3. Select the **Port Information** tab (Figure 28).

Figure 28 FCoE Port Information Tab

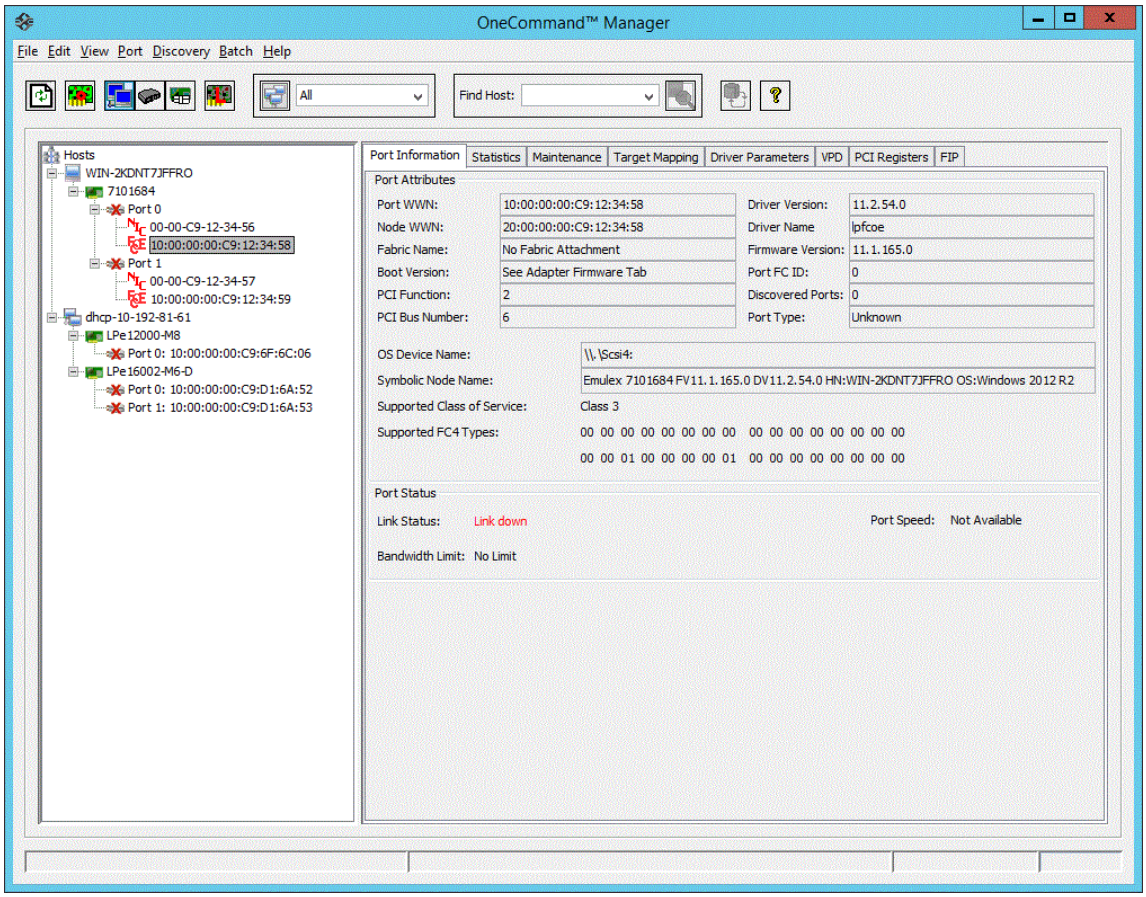

The following FCoE Port Information fields are displayed:

- Port Attributes area:
  - **Port WWN** – The WWPN of the FCoE function.
  - **Node WWN** – The WWNN of the FCoE function.
  - **Fabric Name** or **Host Name** – The **Fabric Name** field is displayed in the Host view, which is a 64-bit worldwide unique identifier assigned to the fabric. The **Host Name** field is displayed in the Fabric view, which is the name of the host containing the FCoE function.
  - **Boot Version** – The version of boot code installed on the selected FCoE function. If the boot code is disabled, the field displays **Disabled**.
  - **Port FC ID** – The FCoE ID for the selected FCoE function.
  - **PCI Function** – The PCI function number assigned by the system to the FCoE function.
  - **PCI Bus number** – The PCI BUS number assigned by the system to the FCoE function.
  - **Driver Version** – The version of the driver installed for the FCoE function.
  - **Driver Name** – The executable file image name for the driver as it appears in the Emulex driver download package.
  - **Firmware Version** – The version of Emulex firmware currently active on the FCoE function.
  - **Discovered Ports** – The number of mapped and unmapped FC/FCoE ports found during discovery by the Emulex adapter driver. The mapped ports are targets and the unmapped ports are non-targets, such as switches or adapters.
  - **Port Type** – The current operational mode of the selected adapter's port.
  - **OS Device Name** – The platform-specific name by which the selected FCoE function is known to the operating system.
  - **Symbolic Node Name** – The FC name used to register the driver with the name server.
  - **Supported Class of Service** – A frame delivery scheme exhibiting a set of delivery characteristics and attributes. Three classes of service include:
    - **Class 1** – Provides a dedicated connection between a pair of ports with confirmed delivery or notification of non-delivery.
    - **Class 2** – Provides a frame switched service with confirmed delivery or notification of non-delivery.
    - **Class 3** – Provides a frame switched service similar to Class 2 but without notification of frame delivery or non-delivery.
  - **Supported FC4 Types** – A 256-bit (8-word) map of the FC-4 protocol types supported by the FCoE function containing the selected adapter.
- Port Status area:
  - **Link Status** – The status of the link on the selected FCoE function.
  - **Port Speed** – The current speed of the selected FCoE function.
  - **Bandwidth Limit** – The QoS bandwidth restriction on the FCoE function.

8.1.2 Viewing and Clearing FC and FCoE Statistics

When you select an FC port or FCoE function from the discovery-tree, the **Statistics** tab (Figure 29) provides cumulative totals for various error events and statistics on the port. When supported by the adapter, you can also clear all the values displayed in the tab.

NOTE Some statistics are cleared when the adapter is reset.

To view FC or FCoE statistics, perform these steps:

1. Select the **Host** or **Fabric** view.
2. Select an FC or FCoE function in the discovery-tree.

3. Select the **Statistics** tab.

Figure 29 Statistics Tab

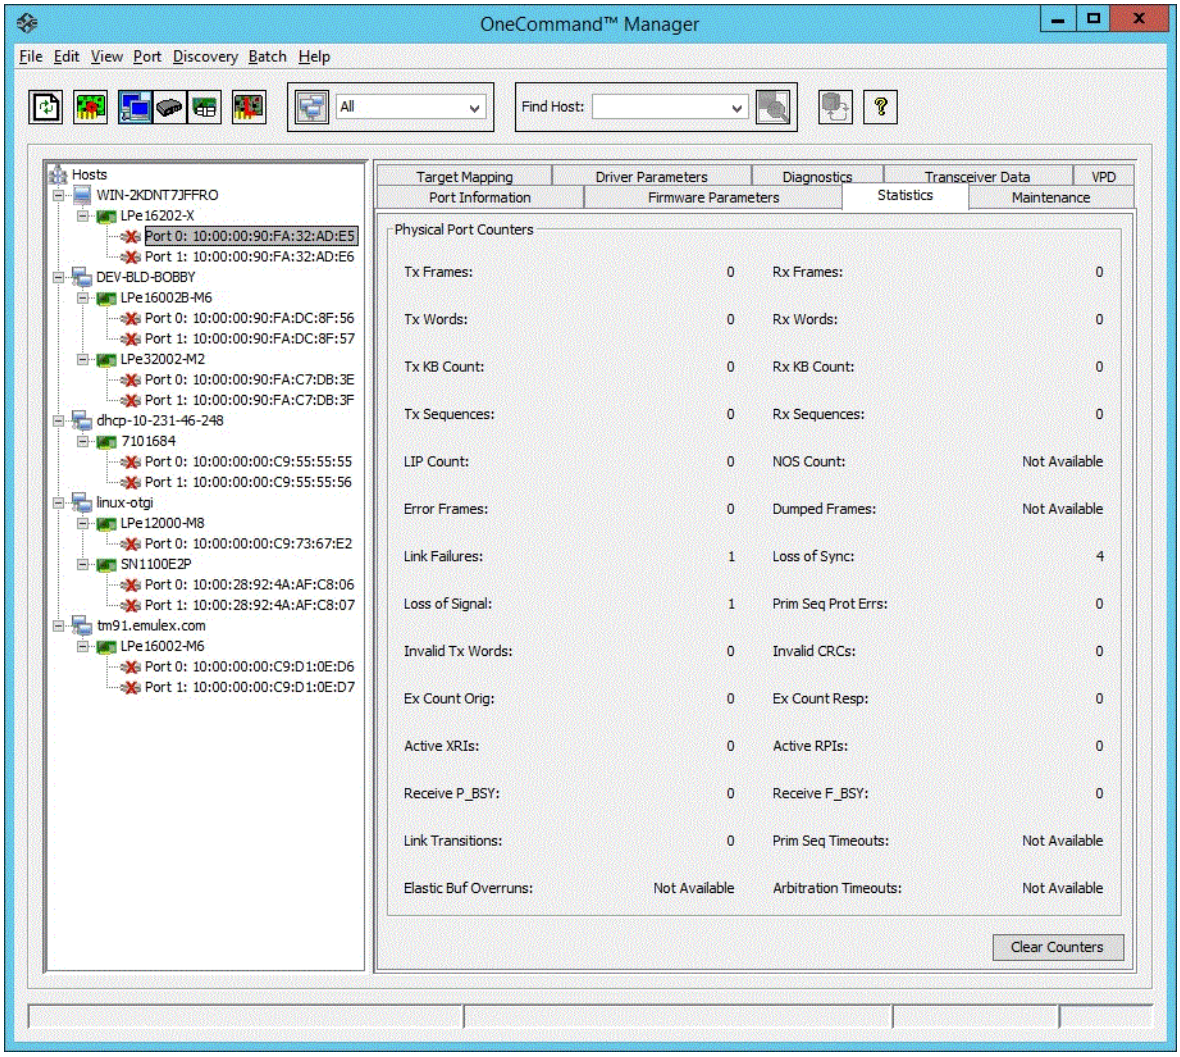

The following Port Statistics fields are displayed:

- **Tx Frames** – FC frames transmitted by this FC or FCoE function.
- **Tx Words** – FC words transmitted by this FC or FCoE function.
- **Tx KB Count** – FC kilobytes transmitted by this FC or FCoE function.
- **Tx Sequences** – FC sequences transmitted by this FC or FCoE function.
- **LIP count** – The number of loop initialization primitive (LIP) events that have occurred for the FC or FCoE function. This field is not supported if the topology is not arbitrated loop. Loop initialization consists of the following:
  - Temporarily suspending loop operations.
  - Determining whether loop capable ports are connected to the loop.
  - Assigning AL_PA IDs.
  - Providing notification of configuration changes and loop failures.
  - Placing loop ports in the monitoring state.
- **Error Frames** – The number of frames received with CRC errors.
- **Link Failures** – The number of times the link has failed. A link failure is a possible cause of a timeout.

- Loss of Signal – The number of times the signal was lost.
- Invalid Tx Words – The total number of invalid words transmitted by this FC or FCoE function.
- Ex Count Orig – The number of FC exchanges originating on this FC or FCoE function (not supported on VMware ESXi servers being managed through the CIM interface).
- Active XRIs – The number of active exchange resource indicators (not supported on VMware based ESXi platforms using the CIM interface).
- Received P_BSY – The number of FC port-busy link response frames received.
- Link Transitions – The number of times the SLI port sent a link attention condition.
- Elastic Buf Overruns – The number of times the link interface has had its elastic buffer overrun.
- Rx Frames – The number of FC frames received by this FC or FCoE function.
- Rx Words – The number of FC words received by this FC or FCoE function.
- Rx KB Count – The received kilobyte count by this FC or FCoE function.
- Rx Sequences – The number of FC sequences received by this FC or FCoE function (not supported on VMware ESXi servers being managed through the CIM interface).
- NOS count – The number of NOS events that have occurred on the switched fabric. The NOS count is not currently supported for Emulex Windows drivers or for arbitrated loop.
- Dumped Frames – The number of frames that were lost because of a lack of host buffers available. This option is not currently supported for the Storport Miniport driver or the driver for Solaris.
- Loss of Sync – The number of times loss of synchronization has occurred.
- Prim Seq Prot Errs – The primitive sequence protocol error count. This counter is incremented whenever there is any type of protocol error.
- Invalid CRCs – The number of frames received that contain CRC failures.
- Ex Count Resp – The number of FC exchange responses made by this FC or FCoE function (not supported on VMware ESXi servers being managed through the CIM interface).
- Active RPIs – The number of RPIs (not supported on VMware ESXi servers being managed through the CIM interface).
- Receive F_BSY – The number of FC port-busy link response frames received.
- Primitive Seq Timeouts – The number of times a primitive sequence event timed out (not supported on VMware ESXi servers being managed through the CIM interface).
- Arbitration Timeouts – The number of times the arbitration loop has timed out. Large counts could indicate a malfunction somewhere in the loop or heavy usage of the loop (not supported on VMware ESXi servers being managed through the CIM interface).

If supported by the adapter, click **Clear Counters** to clear all the values displayed on the tab.

8.1.3 Viewing FC and FCoE Virtual Port Information

The **Discovery Information** tab (Figure 30) displays information about FC and FCoE virtual ports and their associated targets and LUNs.

To view virtual port information, perform these steps:

1. Perform one of the following tasks:
   - From the **View** menu, select **Group Adapters by Virtual Port**.
   - From the toolbar, click **Group Adapters by Virtual Port**.

   The **Discovery Information** tab appears (Figure 30).

Figure 30 Discovery Information Tab

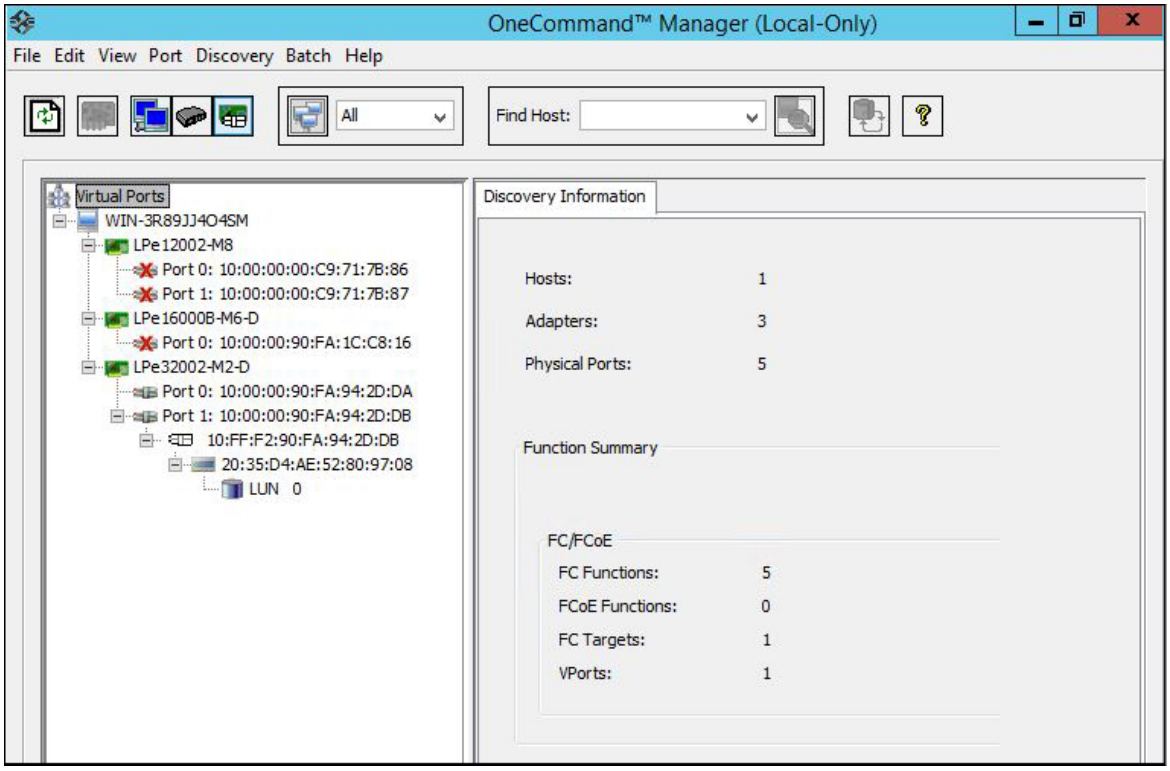

The following Discovery Information fields are displayed:

- **Hosts** – The total number of hosts discovered in the SAN.
- **Adapters** – The total number of adapters discovered in the SAN.
- **Physical Ports** – The total number of physical ports discovered in the SAN.
- Function Summary area:
  - FC/FCoE:
    - **FC Functions** – The total number of FC functions discovered in the SAN.
    - **FCoE Functions** – The total number of FCoE functions discovered in the SAN.
    - **FC Targets** – The total number of FC targets discovered in the SAN.
    - **VPorts** – The total number of virtual FC/FCoE ports discovered in the SAN.

8.1.4 Creating and Deleting FC and FCoE Virtual Ports

This section describes how to create and delete virtual ports.

8.1.4.1 Creating Virtual Ports

Using the **Virtual Ports** tab (Figure 31), you can automatically generate the WWPN for the virtual port based on the WWPN for the physical port or you can manually type the WWPN.

NOTES

- The OneCommand Manager application cannot create or delete virtual ports on VMware ESXi server systems. Although VMware ESXi server supports NPIV, only VMware management tools can create or delete virtual ports.
- In Linux, virtual ports do not persist across system reboots.

The NPIV driver parameter must be enabled before you attempt to create a virtual port. The driver parameter name varies slightly depending upon your operating system:

- For Windows: From the **Driver Parameters** tab, highlight **enableNPIV**, then select **Enable**. On the Storport Miniport system, the `SLIMode` driver parameter must also be set to `0` or `3`.
- For Solaris: `enable-npiv`
- For Linux: `lpfc_enable_npiv`

See Section 8.1.19, Configuring the FC and FCoE Driver Parameters, for more information on enabling driver parameters.

To create a virtual port, perform these steps:

1. Perform one of the following tasks:
   — From the **View** menu, select **Group Adapters by Virtual Ports**.
   — From the toolbar, click **Group Adapters by Virtual Ports**.
2. From the discovery-tree, select the FC or FCoE function on which you want to create a virtual port. The **Virtual Ports** tab appears (Figure 31).

Figure 31 Virtual Ports Tab

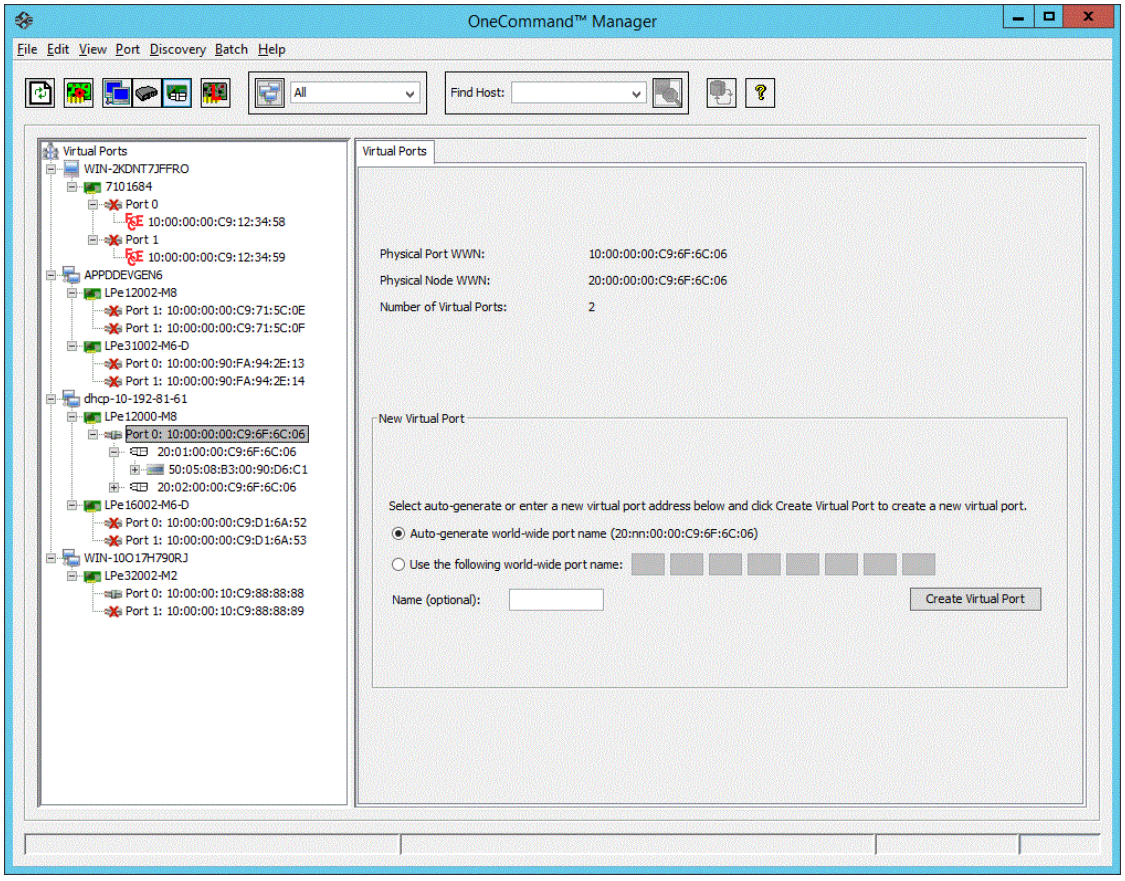

3. Perform one of the following tasks:
   — Select **Auto-generate world wide port name**. The OneCommand Manager application creates the unique WWPN for the new virtual port based on the WWPN of the FC or FCoE function. This option allows you to automatically create up to 255 unique virtual ports for each physical port. It also has the advantage that the new WWPN is unique.

NOTE After auto-generating 255 unique virtual ports, you cannot auto-generate any more virtual ports even if you delete existing auto-generated virtual ports. However, you can still enter your own WWPN to create a virtual port.

— Check **Use the following world-wide port name** and enter a unique WWPN you want to use. You can create as many virtual ports as you want. A valid port name must have one of the following formats:

10:00:xx:xx:xx:xx:xx:xx
2x:xx:xx:xx:xx:xx:xx:xx
3x:xx:xx:xx:xx:xx:xx:xx
5x:xx:xx:xx:xx:xx:xx:xx

where x is a hexadecimal value.

CAUTION Make sure that a manually entered WWPN is unique to your particular SAN. Otherwise, a non-functioning SAN and data loss could occur.

4. Enter an optional name for the virtual port if you want. You can give the new virtual port any name you want up to 99 characters in length. This name is used as part of the Symbolic Node Name for the vPort.
5. Click **Create Virtual Port**. A dialog appears notifying you that the virtual port was created. The dialog also displays the new virtual port's WWPN. Each virtual port has its own WWPN, but its WWNN is the same as the physical port's WWNN.

NOTE If you entered a WWPN that is already in use, you are prompted to enter another WWPN.

6. Click **OK**. The new virtual port is added to the discovery-tree (Figure 4) under the physical port where it was created and the **Number of Virtual Ports** field is updated.

NOTE The OneCommand Manager application automatically refreshes its discovery after a virtual port is created. However, targets for a new virtual port may not be discovered during the refresh. Therefore, you must refresh the discovery until the targets appear under the virtual port in the discovery-tree (Figure 4).

8.1.4.2 Deleting Virtual Ports

NOTE The OneCommand Manager application cannot create or delete virtual ports on VMware ESXi server systems. Although VMware ESXi server supports NPIV, only VMware management tools can create or delete virtual ports.

To delete a virtual port, perform these steps:

1. Perform one of the following tasks:
   — From the **View** menu, select **Group Adapters by Virtual Ports**.
   — From the toolbar, click **Group Adapters by Virtual Ports**.
2. From the discovery-tree, select the virtual port you want to delete. The **Virtual Ports** tab appears (Figure 32).

Figure 32 Virtual Port Tab

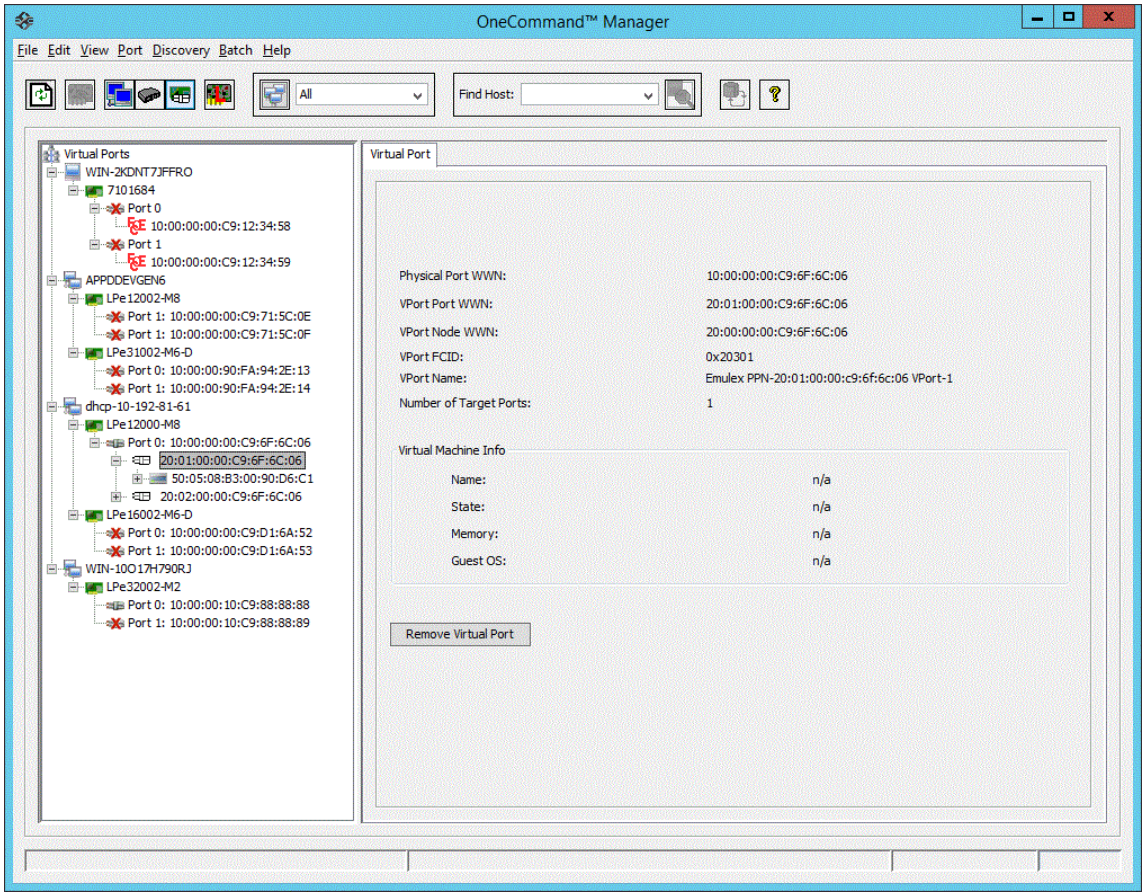

3. Click **Remove Virtual Port.** The **Delete Virtual Port Warning** dialog appears (Figure 33).

Figure 33 Delete Virtual Port Warning Dialog

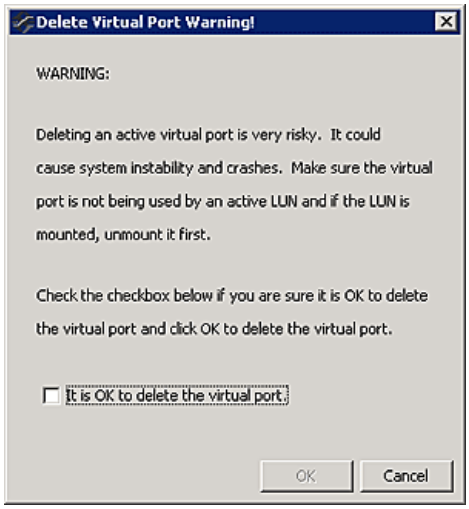

NOTE The link on the physical port must be up to delete a virtual port. The **Remove Virtual Port** button on the **Virtual Port** tab is disabled if the link is down.

4. Select **It is OK to delete the virtual port** and click **OK**. You are notified that the virtual port is no longer available and that it was removed from the discovery-tree (Figure 4).
5. Click **OK**.

8.1.5 Viewing FC and FCoE Fabric Information

The **Discovery Information** tab (Figure 34) contains information about the selected fabric.

To view fabric discovery information, perform one of the following tasks:

— From the **View** menu, select **Group Adapters by Fabric Address**.
— From the toolbar, click **Group Adapters by Fabric Address**.

The **Discovery Information** tab is displayed (Figure 34).

Figure 34 Fabric Discovery Information

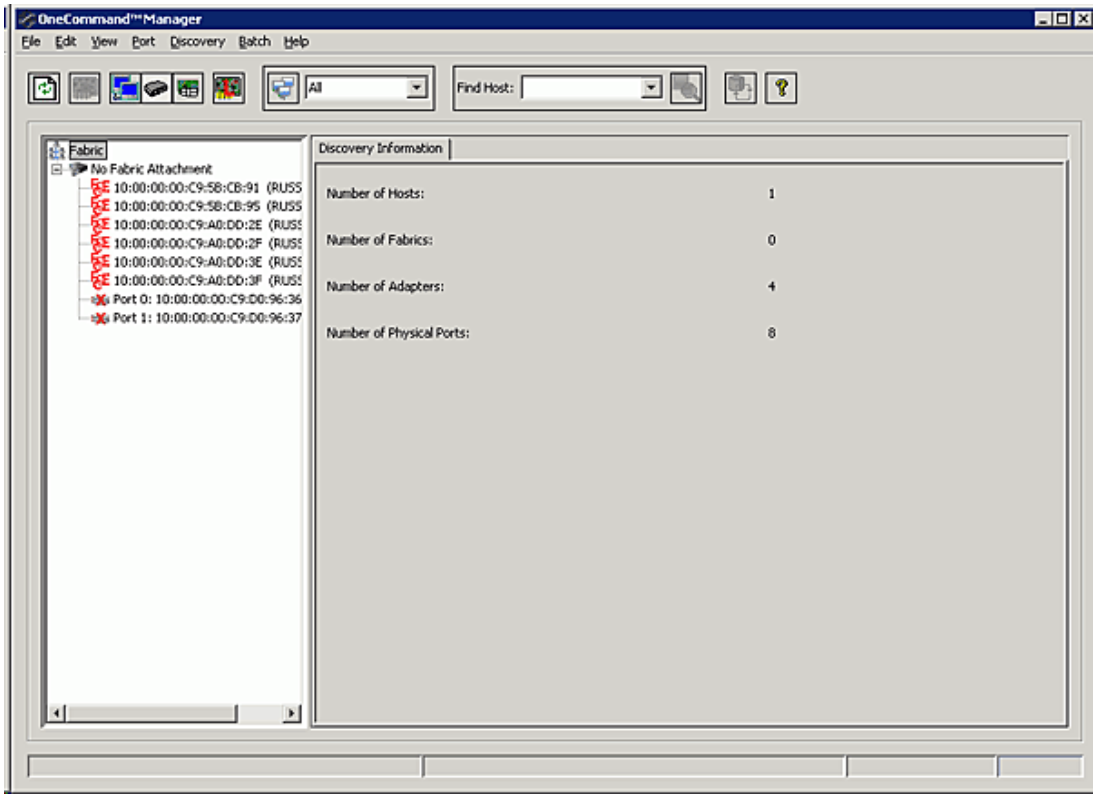

The following Discovery Information fields are displayed:

- **Number of Hosts** – The number of hosts discovered or seen by this host on the selected fabric.
- **Number of Fabrics** – The number fabrics identified during discovery.
- **Number of Adapters** – The number of adapters discovered by this host on the selected fabric.
- **Number of Physical Ports** – The number of discovered physical ports on this host that can be managed by this host.

8.1.6 Viewing FC Port Transceiver Information

NOTE This section does not apply to LPe16202/OCe15100 adapters in NIC+FCoE mode.

When you select an FC port from the discovery-tree (Figure 4), the **Transceiver Data** tab (Figure 35) enables you to view transceiver information such as, vendor name, serial number, and part number. If the adapter/transceiver does not support some or all of the transceiver data, the fields display N/A.

To view FC transceiver information, perform these steps:

1. Select the **Host** or **Fabric** view.
2. In the discovery-tree (Figure 4), select the FC port whose transceiver information you want to view.
3. Select the **Transceiver Data** tab (Figure 35).

Figure 35 FC Transceiver Data Tab

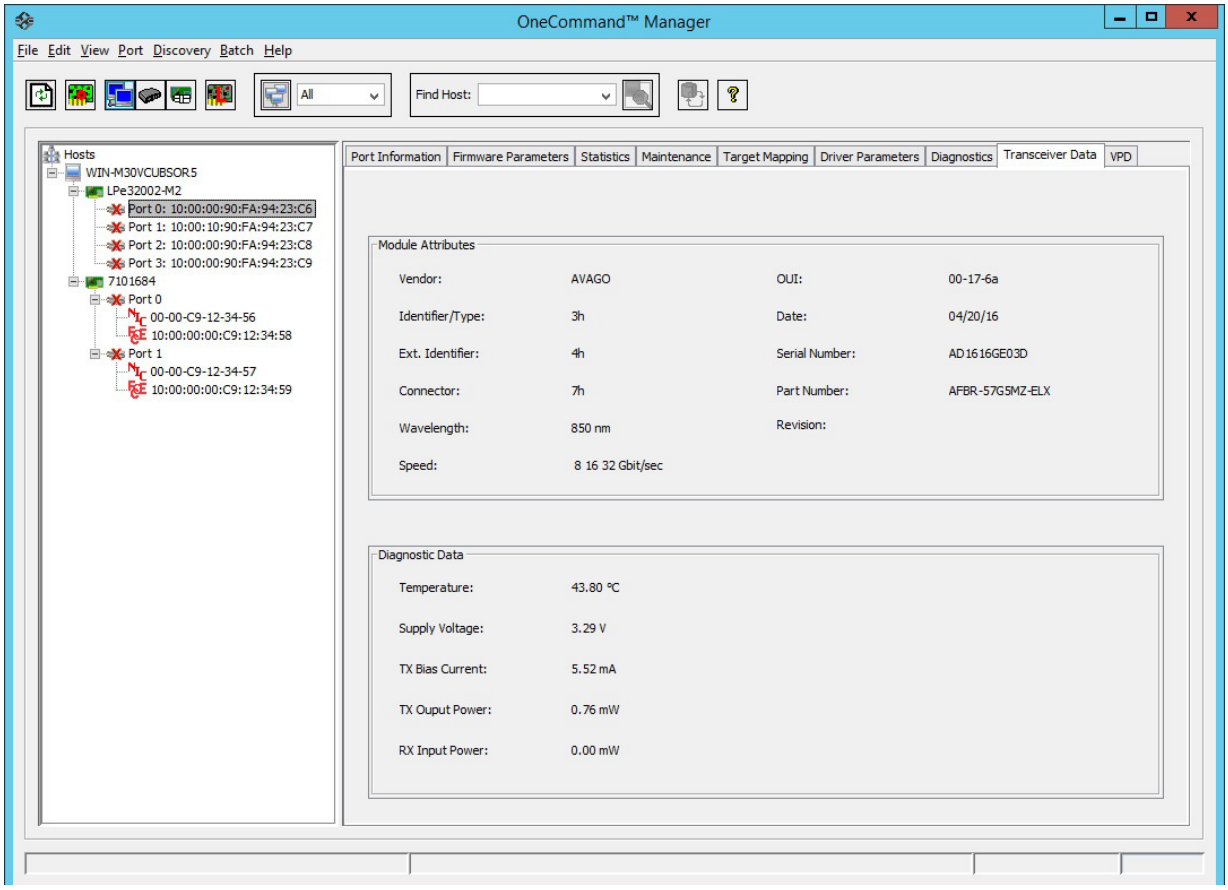

The following Transceiver Data fields are displayed:

- Module Attributes area:
  - **Vendor** – The name of the vendor.
  - **Identifier/Type** – The identifier value that specifies the physical device described by the serial information.
  - **Ext. Identifier** – Additional information about the transceiver.
  - **Connector** – The external optical or electrical cable connector provided as the media interface.
  - **Wavelength** – The nominal transmitter output wavelength at room temperature.

- **Speed** – The speed, or speeds, at which the selected port can run.
- **OUI** – The vendor's OUI. It is also known as the IEEE Company Identifier for the vendor.
- **Date** – The vendor's date code in the MM/DD/YY format.
- **Serial Number** – The serial number provided by the vendor.
- **Part Number** – The part number provided by the SFP vendor.
- **Revision** – The vendor revision level.

- Diagnostic Data area:
  - **Temperature** – The internally measured module temperature.
  - **Supply Voltage** – The internally measured supply voltage in the transceiver.
  - **TX Bias Current** – The internally measured transmitted bias current.
  - **TX Output Power** – The measured transmitted output power.
  - **RX Input Power** – The measured received input power.

8.1.7 Viewing FC and FCoE VPD Information

The **VPD** tab (Figure 36) displays vital product data (if available) for the selected FC or FCoE adapter port such as the product name, part number, and serial number.

To view VPD information, perform these steps:

1. Select the **Host** or **Fabric** view.
2. In the discovery-tree, select the FC or FCoE function whose VPD information you want to view.
3. Select the **VPD** tab.

Figure 36 FC or FCoE VPD Tab

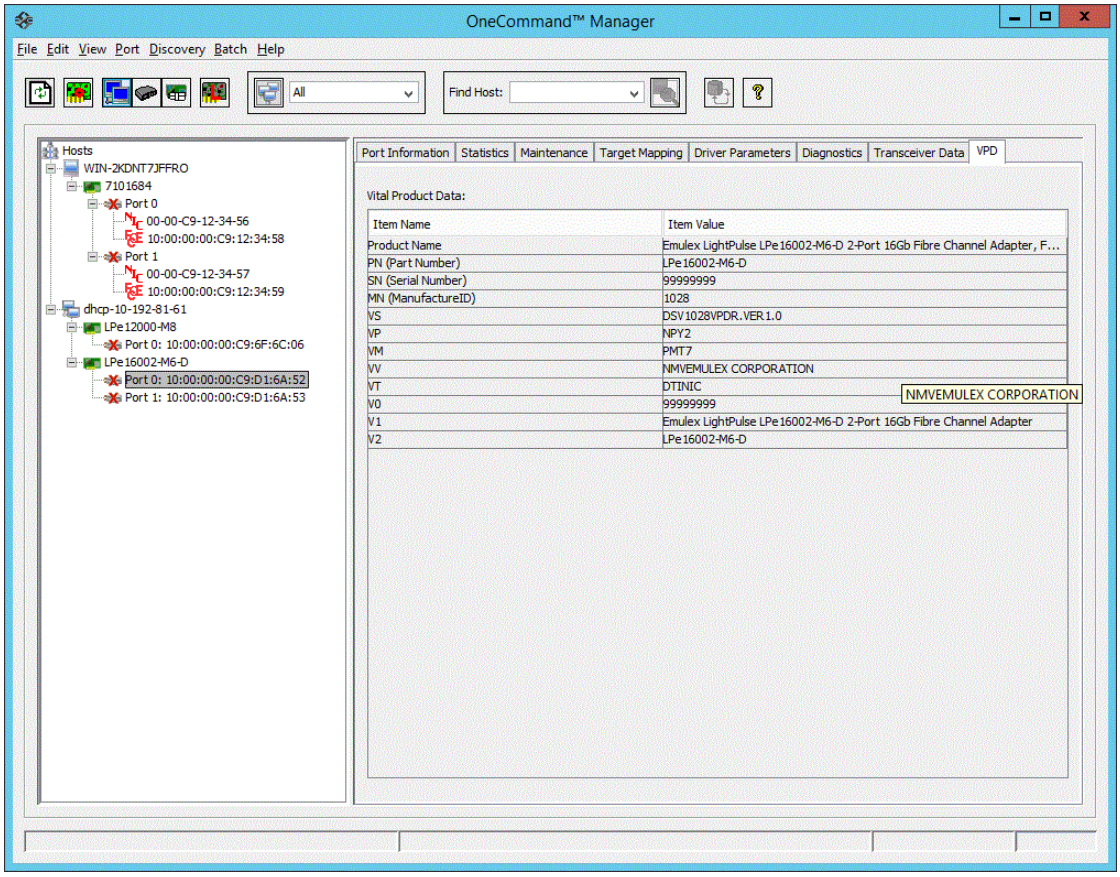

The following Virtual Product Data fields are displayed:

- **Product Name** – Product information about the selected FC or FCoE function.
- **PN (Part Number)** – The adapter's part number.
- **SN (Serial Number)** – The adapter's serial number.
- **MN (Manufacture ID)** – The manufacturer's identification number.
- **VO** – Vendor unique data. **V** indicates a vendor-specific field. An adapter may have none, one, or more of these fields defined. Valid values for this field are VO (the letter O, not the number zero) and V*x* (where *x* is a number).

NOTE Some adapters may show additional VPD information such as, EC (EC level), MN (manufacturer ID), and XY data. Data in the **XY** field is a vendor-specific hexadecimal dump.

8.1.8 Viewing FC Maintenance Information

NOTE This option is not available in read-only mode.

Use the **Maintenance** tab (Figure 37) to view firmware information and update adapter firmware. You can also configure boot from SAN and change WWPN and WWNN information for the selected adapter port.

To view FC firmware information, perform these steps:

1. Select the **Host** or **Fabric** view.
2. Select an FC port in the discovery-tree.

3. Select the **Maintenance** tab (Figure 37).

Figure 37 FC Maintenance Tab

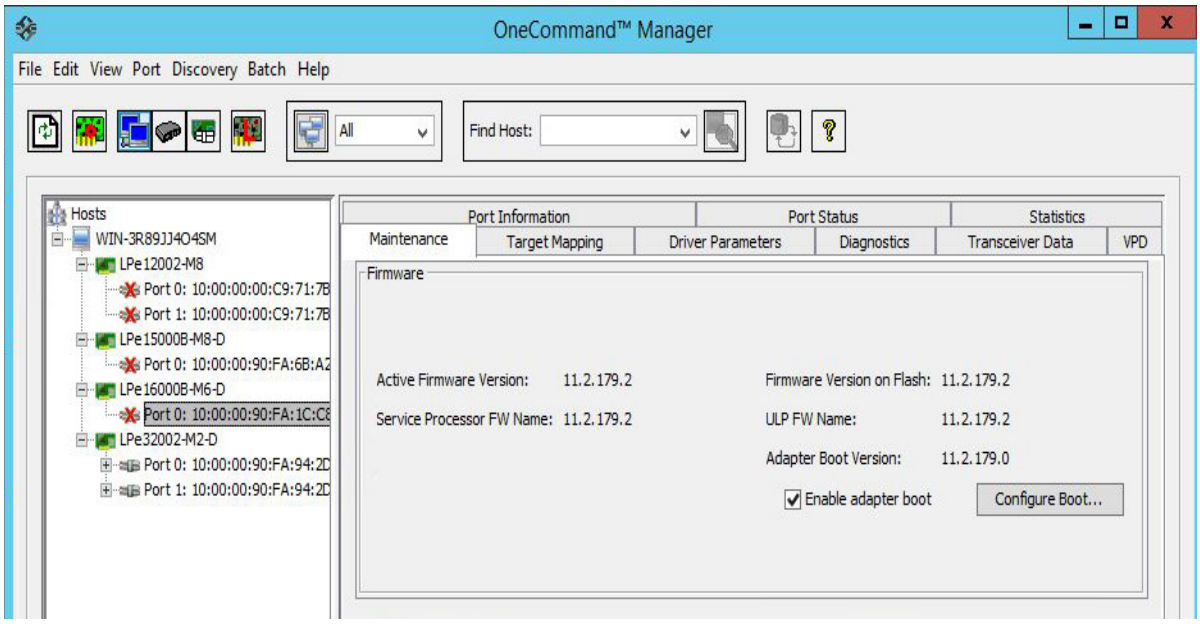

The following **Maintenance** tab fields are displayed:

- Firmware area:
  - **Active Firmware Version** – The Emulex firmware version number for this FC port.
  - **Service Processor FW Name** – The Emulex firmware name for this FC port.
  - **Firmware Version on Flash** – The flash firmware version currently being used by the adapter.
  - **ULP FW Name** – The firmware version running on the ULP processors within the ASIC.
  - **Adapter Boot Version** – Displays one of the following:
    - The selected adapter port's boot code version if boot code is present.
    - Disabled if the boot code is disabled.
    - Not Present if boot code is not loaded. If boot code is not loaded, the **Enable Adapter boot** check box is not visible and you cannot configure the selected port to boot from SAN.
  - **Enable adapter boot** check box –Select this check box if you want the FC port to load and execute boot code during system startup. Click **Configure Boot** to configure boot from SAN (not available in read-only mode).

    NOTE Enabling adapter boot only causes the FC port to load the boot code and run it during system startup. It does not mean that the FC port boots from SAN. To boot from SAN, the boot type must be enabled. Enable this in the **Boot from SAN configuration** window for each boot type.
- WWN Management area:

  NOTE The WWN Management area is disabled when FA-PWWN is enabled on an adapter port.
  - Current:
    - **WWPN** – The World Wide Port Name for the selected FC port.
    - **WWNN** – The World Wide Node Name for the selected FC port.

— Pending Changes:

- **WWPN** – Works with the **Change WWN** button. It displays the WWPN you assigned for the selected FC port, but the system must be rebooted for these changes to take effect and appear under the Current listing. See Section 8.1.16, Changing the FC and FCoE WWPN and WWNN, for more information.
- **WWNN** – Works with the **Change WWN** button. It displays the WWNN you assigned for the selected FC port, but the system must be rebooted for these changes to take effect and appear under the Current listing. See Section 8.1.16, Changing the FC and FCoE WWPN and WWNN, for more information.

For LPe12000-series adapters, the tab includes a **Download Firmware** button. For instructions on updating firmware on a port of an LPe12000-series adapter, see Section 9, Updating Adapter Firmware.

8.1.9 Viewing FCoE Maintenance Information

NOTE This option is not available in read-only mode.

Use the **Maintenance** tab (Figure 38) to view firmware information. You can also configure boot from SAN and change WWPN and WWNN information for the selected FCoE function.

To view FCoE firmware information, perform these steps:

1. Select the **Host** or **Fabric** view.
2. Select an FCoE function in the discovery-tree.
3. Select the **Maintenance** tab.

Figure 38 FCoE Maintenance Tab

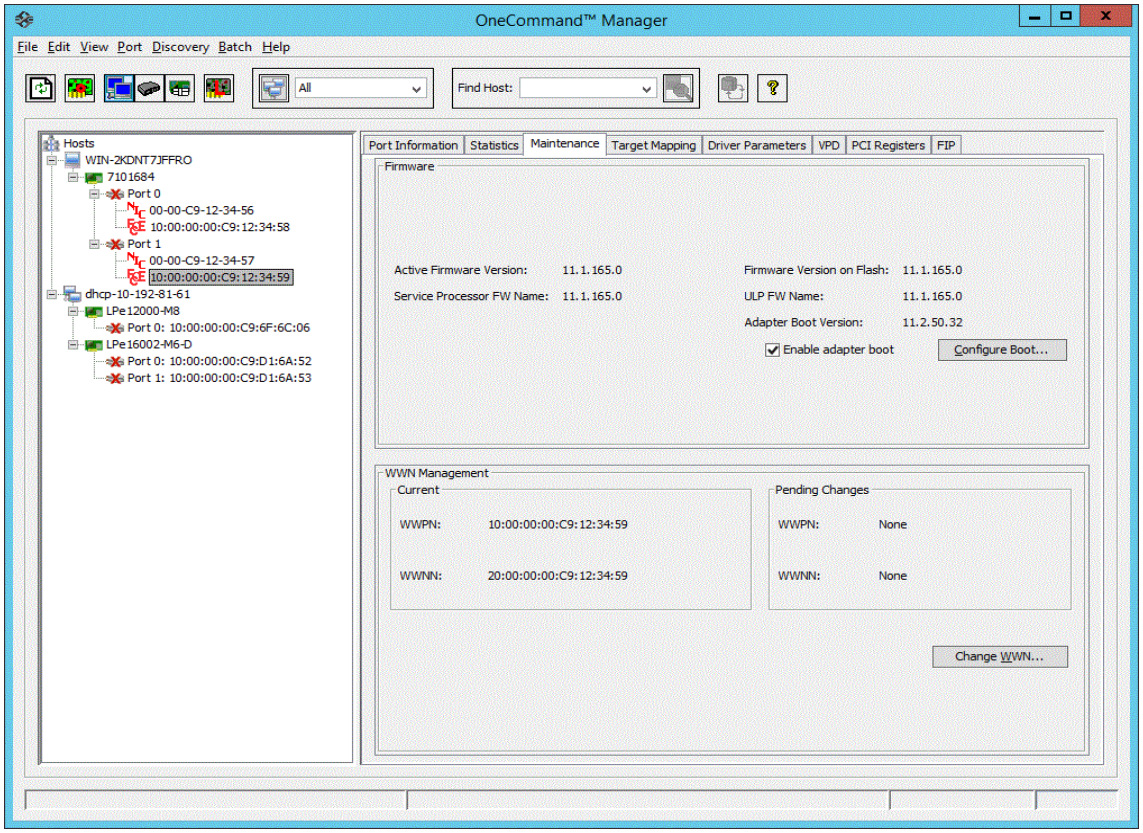

The following **Maintenance** tab fields are displayed:

- Firmware area:
  - **Firmware Version on Flash** – The firmware version stored on the adapter's non-volatile storage. When the system restarts, this version becomes the active firmware version.
  - **Service Processor FW Version** – The firmware version that is currently operational on the adapter.
  - **Active Firmware Version** – The version of firmware running on the adapter.
  - **ULP FW Name** – The firmware version running on the ULP processors within the ASIC.
  - **Adapter Boot Version** – Displays one of the following:
    - The selected adapter port's boot code version if boot code is present.
    - Disabled if the boot code is disabled.
    - Not Present if boot code is not loaded. If boot code is not loaded, the **Enable Adapter boot** check box is not visible and you cannot configure the selected port to boot from SAN.
  - **Enable adapter boot** check box – Select this check box if you want the FCoE port to load and execute boot code during system startup. Click **Configure Boot** to configure boot from SAN (not available in read-only mode).

    NOTE Enabling adapter boot causes only the FCoE port to load the boot code and run it during system startup. It does not mean that the FCoE port boots from SAN. To boot from SAN, the boot type must be enabled. Enable this in the **Boot from SAN Configuration** window for each boot type.
- WWN Management area:
  - Current:
    - **WWPN** – The World Wide Port Name for the selected FCoE function.
    - **WWNN** – The World Wide Node Name for the selected FCoE function.
  - Pending Changes:
    - **WWPN** – Works with the **Change WWN** button. It displays the WWPN you assigned for the selected FCoE function, but the system must be rebooted for these changes to take effect and appear under the Current listing. See Section 8.1.16, Changing the FC and FCoE WWPN and WWNN, for more information.
    - **WWNN** – Works with the **Change WWN** button. It displays the WWNN you assigned for the selected FCoE function, but the system must be rebooted for these changes to take effect and appear under the Current listing. See Section 8.1.16, Changing the FC and FCoE WWPN and WWNN, for more information.

8.1.9.0.1 Maintenance Tab Buttons

NOTE The **Maintenance** tab buttons are not available in read-only mode.

The following buttons are available on the **Maintenance** tab:

- Configure Boot – Click to configure boot from SAN (not available on VMware ESXi servers being managed through the CIM interface.)
- Change WWN – Click to change the selected FCoE function's WWNN or WWPN.

8.1.10 Viewing FC or FCoE Target Information

When you select a target associated with an FC or FCoE adapter from the discovery-tree (Figure 4), the **Target Information** tab (Figure 39) displays information associated with that target.

To view FC or FCoE target information, perform these steps:

1. Select the **Host**, **Fabric,** or **Virtual Port** view.
2. In the discovery-tree (Figure 4), select the FC or FCoE target whose information you want to view. The **Target Information** tab appears (Figure 39).

Figure 39 Target Information Tab

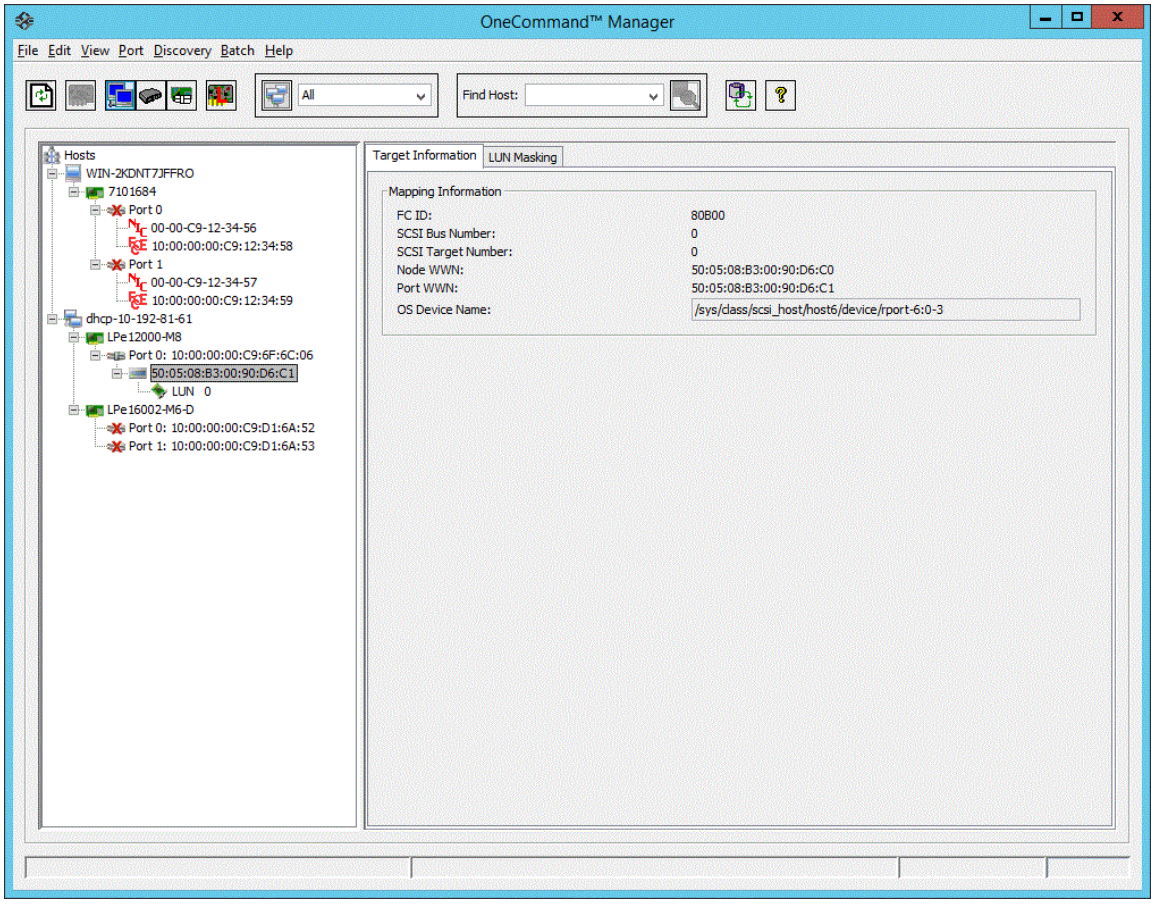

The following Target Information fields are displayed:

- Mapping Information area:
  - **FC ID** – The FC ID for the target; assigned automatically in the firmware.
  - **SCSI Bus Number** – The SCSI bus number to which the target is mapped.
  - **SCSI Target Number** – The target's identifier on the SCSI bus.
  - **Node WWN** – A unique 64-bit number, in hexadecimal, for the target (N_PORT or NL_PORT).
  - **Port WWN** – A unique 64-bit number, in hexadecimal, for the fabric (F_PORT or Switched Fabric Loop Port [FL_PORT]).
  - **OS Device Name** – The operating system device name.

8.1.11 Viewing FC or FCoE LUN Information

When you select a LUN associated with an FC or FCoE target from the discovery-tree (Figure 40), the **LUN** tab displays information associated with that LUN.

NOTE The **Refresh LUNs** button refreshes only the LUN list for the currently selected target.

NOTE On Linux systems, to make LUNs that are newly added to a storage array appear on the host, the following script must run from the command shell:

```
/usr/sbin/lpfc/lun_scan all
```

This script prevents you from having to reboot. If the host machine is rebooted after the LUN is added to the target array, you do not need to run the script.

To view the LUN information, perform these steps:

1. Select the **Host**, **Fabric,** or **Virtual Port** view.
2. From the discovery-tree, select a LUN. The **LUN Information** tab appears (Figure 40).

Figure 40 LUN Information Tab

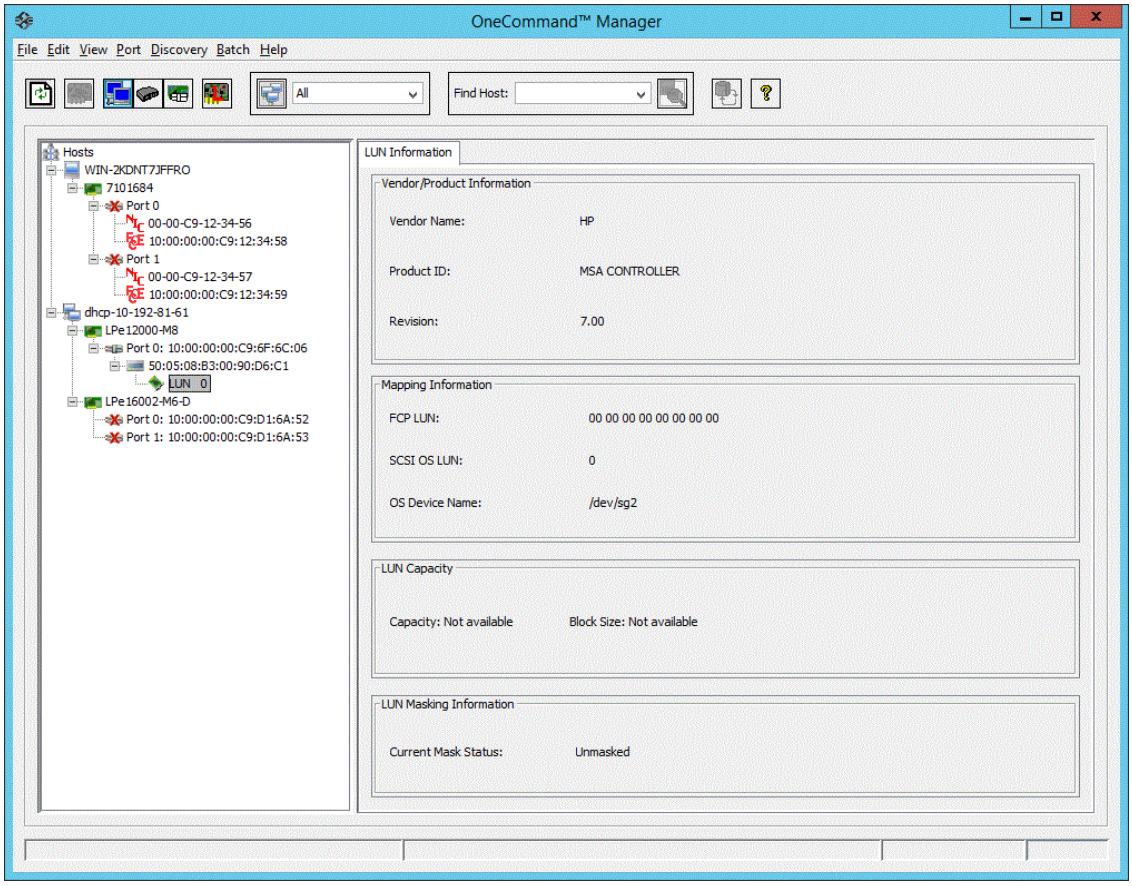

The following LUN Information fields are displayed:

- Vendor Product Information area:
  - **Vendor Name** – The name of the vendor of the LUN.
  - **Product ID** – The vendor-specific ID for the LUN.
  - **Revision** – The vendor-specific revision number for the LUN.
- Mapping Information area:
  - **FCP LUN** – The FC identifier used by the adapter to map to the SCSI OS LUN.
  - **SCSI OS LUN** – The SCSI identifier used by the operating system to map to the specific LUN.

- **OS Device Name** – The name assigned by the operating system to the LUN.

- LUN Capacity area:

  NOTE LUN capacity information is only provided if the LUN is a mass-storage (disk) device. Other devices, such as tapes and scanners, do not display capacity.

  - **Capacity** – The capacity of the LUN, in megabytes.
  - **Block Size** – The length of a logical unit block in bytes.

- LUN Masking area:
  - **Current Mask Status** – Possible states are masked or unmasked. See Section 8.1.14, Masking and Unmasking LUNs (Windows) for more information on LUN Masking.

8.1.12 Viewing FC and FCoE Target Mapping

The **Target Mapping** tab (Figure 41) enables you to view current target mapping and to set up persistent binding.

NOTE Persistent binding is not supported on Solaris systems.

To view target mapping, perform these steps:

1. Select the **Host** or **Fabric** view.
2. In the discovery-tree, select the FC or FCoE function whose target mapping information you want to view.
3. Select the **Target Mapping** tab (Figure 41).

Figure 41 Target Mapping Tab

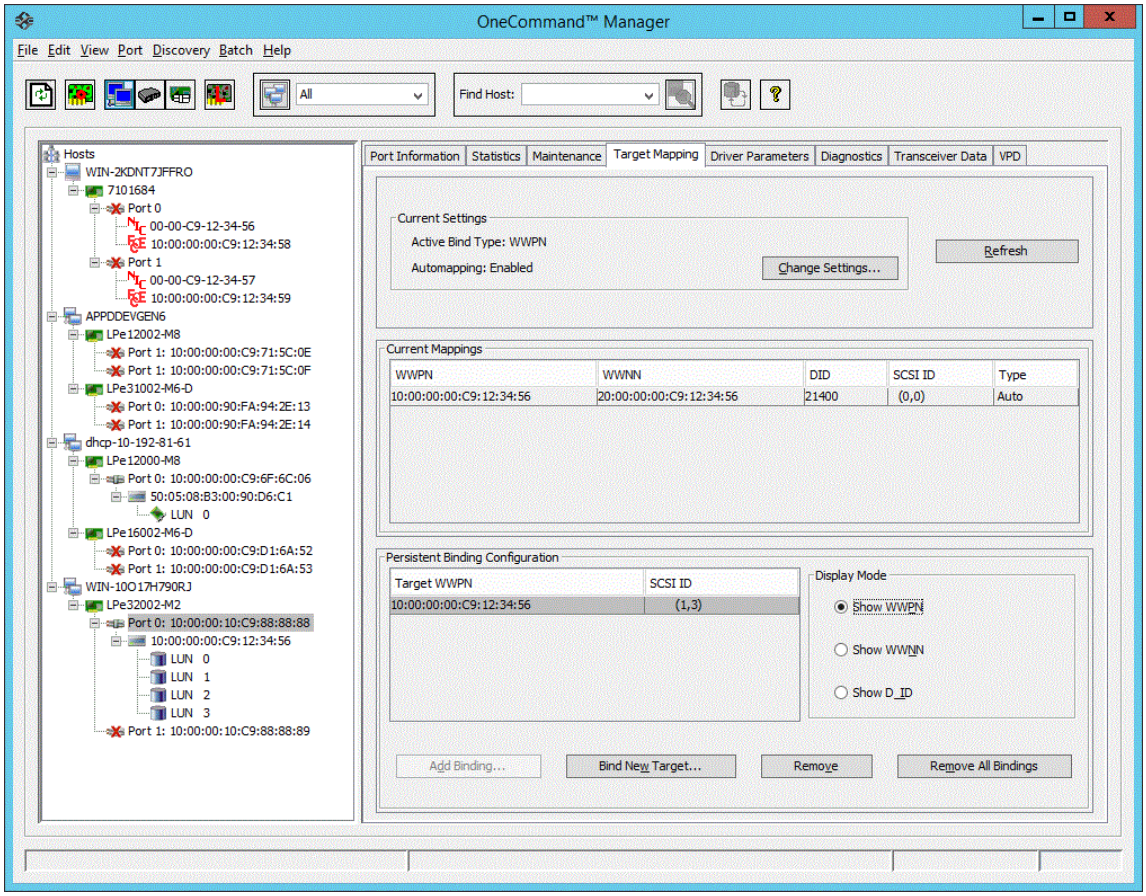

The following Target Mapping fields are displayed:

- Current Settings area:

  NOTE For Linux and VMware ESXi, this area is N/A.

  — **Active Bind Type** – WWPN, WWNN, or a destination identifier (D_ID).

  — **Automapping** – The current state of SCSI device automapping: enabled (default) or disabled.
- Current Mappings area:

  — This table lists current mapping information for the selected FC or FCoE function.
- Persistent Binding Configuration area:

  NOTE For Linux and VMware ESXi, this area is N/A.

  This table lists persistent binding information for the selected FC/FCoE function (not available on VMware ESXi servers being managed through the CIM interface).

  NOTE For Linux and VMware ESXi, this area is N/A.
- Display Mode area:

  — Select the method by which you want to display information in the Persistent Binding Configuration table.

  For information on changing settings, see Section 8.1.13.1, Changing Automapping Settings.

  For information on adding a binding, see Section 8.1.13.2, Adding a Persistent Binding.

For information on binding a new target, see Section 8.1.13.3, Binding a Target that Does Not Appear in the Persistent Binding Table.

To remove a single binding, select the binding and click **Remove**.

To remove all bindings, click **Remove All Bindings**.

8.1.13 Using Automapping and Persistent Binding (Windows Only)

NOTE This option is not available in read-only mode.

Set up persistent binding on remote and local adapters. Global automapping assigns a binding type, target ID, SCSI Bus, and SCSI ID to the device. The binding type, SCSI Bus, and SCSI ID can change when the system is rebooted. With persistent binding applied to one of these targets, the WWPN, SCSI Bus, and SCSI ID remain the same when the system is rebooted.

The driver refers to the binding information at during system boot. When you create a persistent binding, the OneCommand Manager application tries to make that binding dynamic. However, the binding must meet all of the following criteria to be dynamic:

- The SCSI ID (target/bus combination) specified in the binding request must not be mapped to another target. For example, the SCSI ID must not already appear in the 'Current Mappings' table under 'SCSI ID'. If the SCSI ID is already in use, the binding cannot be made dynamic, and a reboot is required.
- The target (WWPN, WWNN, or DID) specified in the binding request must not be mapped to a SCSI ID. If the desired target is already mapped, then a reboot is required.
- The bind type (WWPN, WWNN, or DID) specified in the binding request must match the currently active bind type shown in the Current Settings area of the **Target Mapping** tab. If they do not match, then the binding cannot be made active.

8.1.13.1 Changing Automapping Settings

To change automapping settings, perform these steps:

1. Select the **Host** or **Fabric** view.
2. In the discovery-tree, select the FC or FCoE function you want to set up with persistent binding.
3. Select the **Target Mapping** tab (Figure 42). All targets are displayed.

Figure 42 Target Mapping Tab

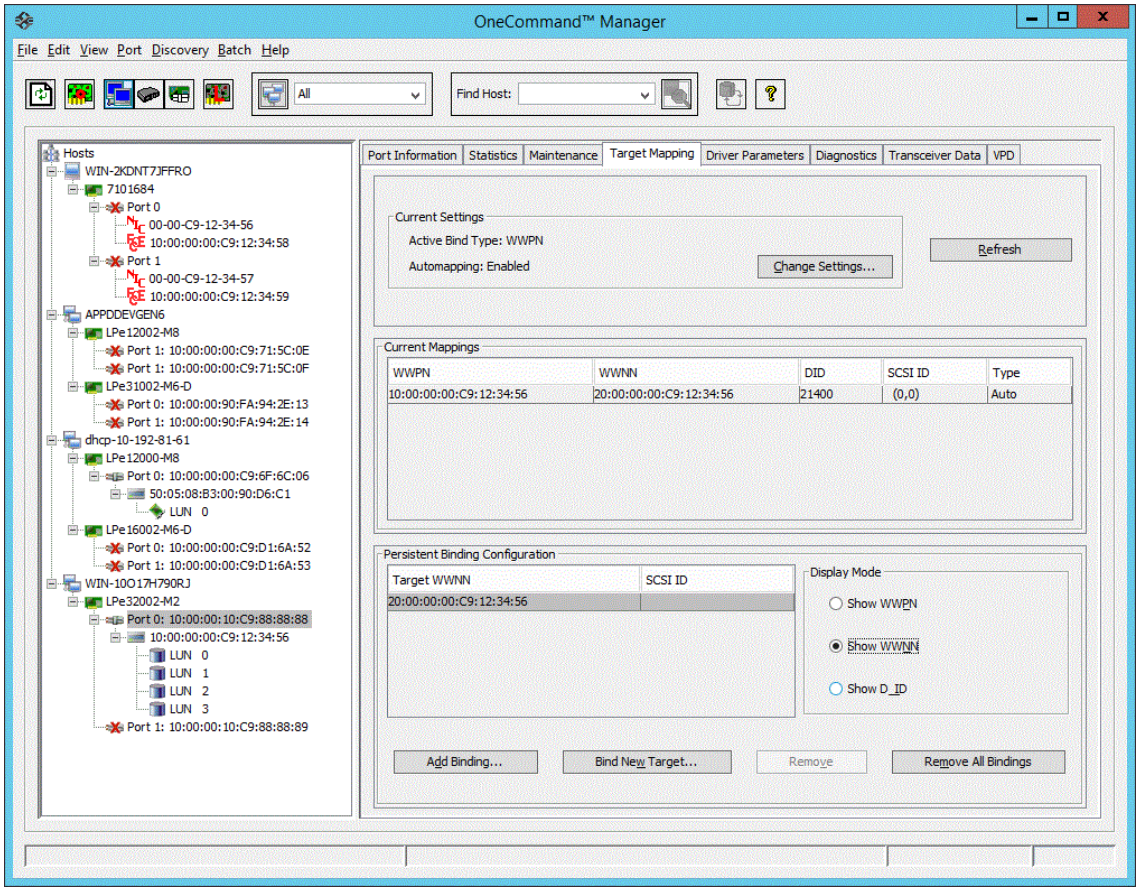

4. Target mappings are displayed by WWPN, WWNN, or D_ID. **PB** indicates mapping from persistent binding, and **Auto** indicates an automapped target. In the Display Mode section, choose the display mode you want to use.
5. If you want to make changes, click **Change Settings**. The **Mapped Target Settings** dialog appears. You can enable or disable auto-mapping and change the active bind type. Click **OK**.
6. Reboot the system for changes to take effect.

8.1.13.2 Adding a Persistent Binding

To add a persistent binding, perform these steps:

1. Select the **Host** or **Fabric** view.
2. In the discovery-tree (Figure 4), select the FC or FCoE function you want to set up with persistent binding.
3. Select the **Target Mapping** tab (Figure 42). All targets are displayed. In the Persistent Binding Configuration table, click the target that you want to bind.
4. Click **Add Binding**. The **Add Persistent Binding** dialog (Figure 43) is displayed.

Figure 43 Add Persistent Binding Dialog

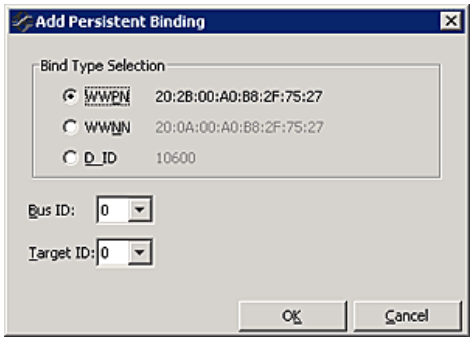

5. Select the bind type that you want to use (**WWPN**, **WWNN**, or **D_ID**).
6. Select the **Bus ID** and **target ID** that you want to bind, and click **OK.**

NOTE Automapped targets have entries only in the second column of the Current Mappings table. Persistently bound targets have entries in the second and third columns. In this case, the third column contains the SCSI Bus and target numbers you specified in the **Add Persistent Binding** dialog. This binding takes effect only after the local machine is rebooted.

8.1.13.3 Binding a Target that Does Not Appear in the Persistent Binding Table

NOTE It is possible to specify a SCSI bus and target that have already been used on behalf of a different FC target. Attempting to bind a target already in the Persistent Binding table on the **Target Mapping** tab results in an error message:

```
Target already in target list.
```

Click **Add Binding**.

To bind a target that does not appear in the Persistent Binding table on the **Target Mapping** tab, perform these steps:

1. Select the **Host** or **Fabric** view.
2. In the discovery-tree (Figure 4), select the FC or FCoE function you want to set up with persistent binding.
3. Select the **Target Mapping** tab (Figure 41). All targets are displayed.
4. Click **Bind New Target**. The **Bind New Target** dialog is displayed (Figure 44).

Figure 44 Bind New Target Dialog

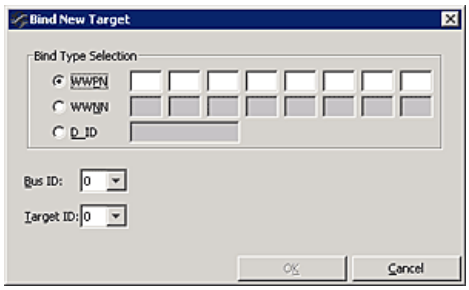

5. Select the type of binding you want to use, and type the WWPN, WWNN, or D_ID that you want to bind to the target.
6. Select the Bus ID and Target ID that you want to bind, and click **OK.**

NOTE A target does not appear on the target list if automapping is disabled and the target is not already persistently bound.

8.1.14 Masking and Unmasking LUNs (Windows)

LUN masking refers to whether or not a LUN is visible to the operating system. A masked LUN is not available and is not visible to the operating system. You can use the **LUN Masking** tab (Figure 45) to mask or unmask LUNs at the host level.

NOTE The **LUN Masking** tab (Figure 45) is not shown in the Virtual Port view because LUN masking is not available for virtual ports.

Figure 45 LUN Masking Tab

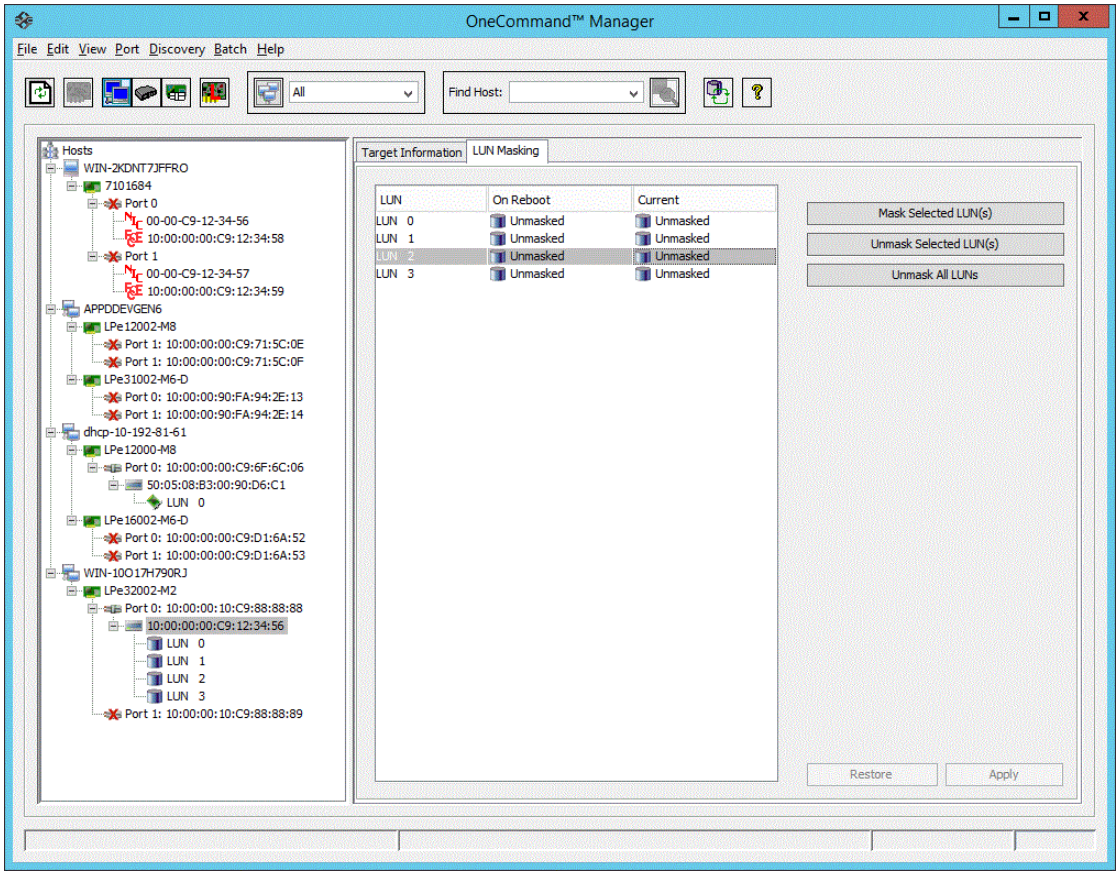

8.1.14.1 LUN Masking Conventions and Guidelines

LUN icons in the discovery-tree (Figure 4) reflect the live mask state currently in use by the driver. Green LUN icons indicate unmasked LUNs. Gray LUN icons indicate masked LUNs. Red text indicates that a LUN mask has been changed, but not applied (saved).

The following LUN Masking information is displayed:

- LUN – The FC LUN number.
- On Reboot – The On Reboot column shows the mask configuration currently saved to the configuration file on disk (Solaris) or to the registry (Windows). Usually, for a specific LUN, the states reported in the On Reboot and Current column are identical. However, there are times where these do not match. For example, the hbacmd utility can be used to change only the Current mask state for a LUN and not touch the On Reboot mask state contained in the configuration file.
- Current – The Current column displays the live mask state currently in use by the driver. When you first see the **LUN Masking** tab, the mask states displayed in the Current column are identical to the mask states for the corresponding LUNs in the discovery-tree (Figure 4).

To change the mask status of a LUN, perform these steps:

1. Select the **Host** view.
2. From the discovery-tree (Figure 4), select the target whose LUN masking state you want to change. A set of LUNs appears below the selected target.
3. Select the **LUN Masking** tab (Figure 45). This tab contains a list of the same set of LUNs that appear below the FC/FCoE target in the discovery-tree (Figure 4).
4. In the LUN list of the **LUN Masking** tab, select one or more LUNs. The **Mask Selected LUNs**, **Unmask Selected LUNs**, **Unmask All LUNs**, **Restore**, and **Apply** buttons become active as appropriate. For example, if the LUN is currently unmasked, only the **Mask Selected LUN(s)** button is active.
5. Change the mask status: click **Mask Selected LUN(s)**, **Unmask Selected LUN(s)** or **Unmask All LUNs** as appropriate. Mask status changes appear in red text.

NOTE To return all mask settings to their status before you started this procedure, click **Restore** before you click **Apply**. If you click **Apply**, changes cannot be cancelled by clicking **Restore**. To unmask all LUNs, click **Unmask All LUNs**. This button is always active. Make sure to also click **Apply** to commit the changes.

6. Click **Apply** to commit the changes. An informational message is displayed that confirms the mask status has changed and the red text changes to black.

8.1.15 Managing ExpressLane LUNs

The OneCommand Manager application allows you to set special priority queuing for selected LUNs by making them ExpressLane LUNs (Figure 46). ExpressLane LUN performance is superior to that of regular LUNs. You can enable ExpressLane LUNs attached to both physical and virtual ports.

You can also assign a frame priority to an ExpressLane LUN if the adapter and the switch support it.

ExpressLane LUN assignments persist across reboots.

NOTES

- Masked LUNs cannot be ExpressLane enabled because they are not presented to the host. Conversely, ExpressLane LUNs cannot be masked.
- For Linux operating systems, if ExpressLane LUNs are created, the VPort must be re-created after a system boot because VPorts do not persist across system reboots. If the VPort is re-created with the same WWPN to which the ExpressLane LUN was previously assigned, and the same LUN is then detected, it becomes an ExpressLane LUN again.

Figure 46 LUN Information Tab (ExpressLane LUN with Frame Priority Selection Supported)

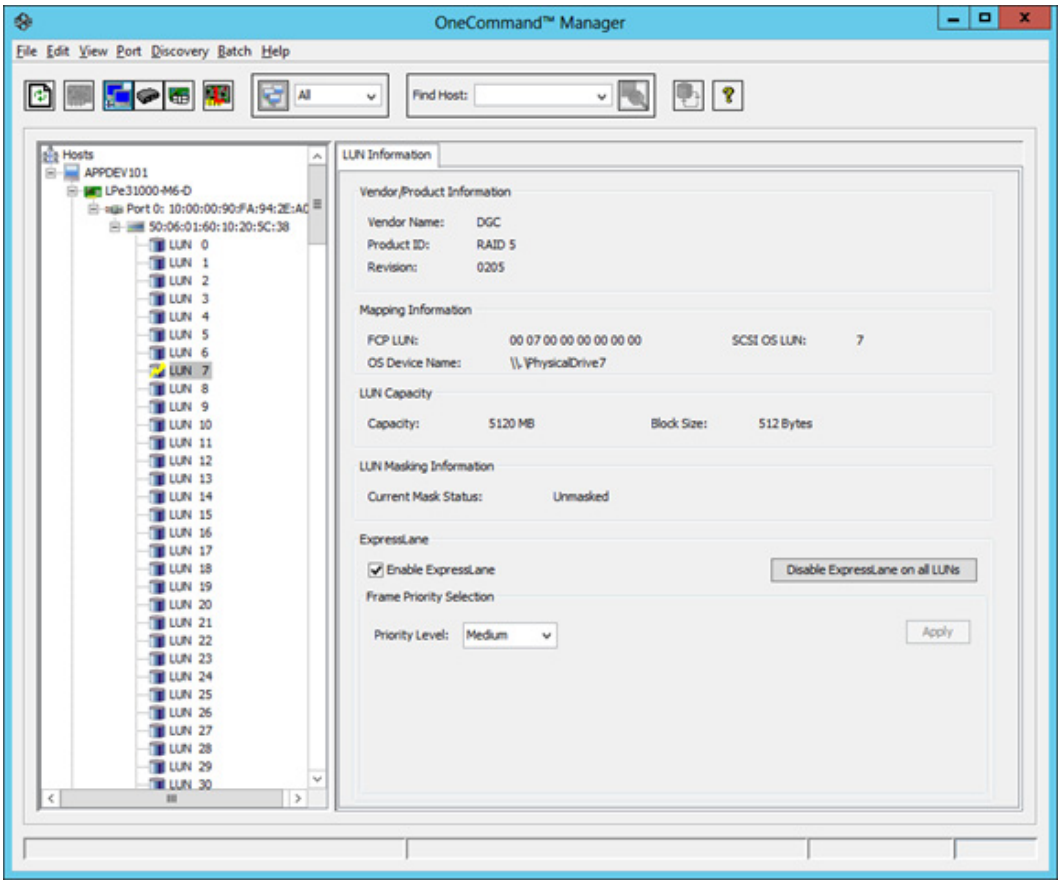

To enable an ExpressLane LUN, perform these steps:

NOTE ExpressLane must be enabled on the Driver Parameters tab to create an ExpressLane LUN. See Section 8.1.19, Configuring the FC and FCoE Driver Parameters, for more information.

1. Select the **Host**, **Fabric,** or **Virtual Port** view.
2. From the discovery-tree, select a LUN under the adapter on which you want to enable ExpressLane. The **LUN Information** tab appears (Figure 46).
3. Select the **Enable ExpressLane** check box.
4. Click **Apply**. The **LUN** icon in the discovery-tree changes to the **ExpressLane LUN** icon.

To disable an ExpressLane LUN, perform these steps:

1. Select the **Host**, **Fabric,** or **Virtual Port** view.
2. From the discovery-tree, select a LUN under the adapter on which you want to disable ExpressLane. The **LUN Information** tab appears (Figure 46).
3. Clear the **Enable ExpressLane** check box to disable the selected LUN.
4. Click **Apply**.

To disable all ExpressLane LUNs, perform these steps:

1. Select the **Host**, **Fabric,** or **Virtual Port** view.
2. From the discovery-tree, select a LUN under the adapter on which you want to disable ExpressLane. The **LUN Information** tab appears (Figure 46).

3. Click **Disable ExpressLane for all LUNs on this target**.
4. A dialog appears warning you that you are about to delete all ExpressLane LUNS on this target. Click **OK**.

   All ExpressLane LUN icons in the discovery-tree (for the selected adapter port) will change to the regular **LUN** icon and any assigned frame priority is set to 0.

8.1.15.1 Selecting a Frame Priority

If the adapter and switch support it, you can assign a frame priority to the ExpressLane LUN. Switches can provide up to three priority levels; **Low**, **Medium**, and **High,** but they might provide fewer options.

To select a frame priority, perform these steps:

NOTE The `EnableXLane` driver parameter must be enabled on the Driver Parameters tab to set frame priorities. See Section 8.1.19, Configuring the FC and FCoE Driver Parameters, for more information.

1. Select the **Host**, **Fabric,** or **Virtual Port** view.
2. From the discovery-tree, select a LUN under the adapter on which you want to enable ExpressLane. The **LUN Information** tab appears (Figure 46).
3. Select the **Enable ExpressLane** check box if it is not already selected. The **LUN** icon in the discovery-tree changes to the **ExpressLane LUN** icon.
4. Select a frame priority from the **Priority Level** drop-down list.

NOTE If the switch connected to the FC or FCoE initiator does not support LUN-specific frame priority levels using the Get Fabric Object (GFO), you must manually enter the frame priority values in the range of 0–127 for all ExpressLane enabled LUNs as depicted in Figure 47.

Figure 47 LUN Information Tab (Frame Priority Not Supported by the Switch)

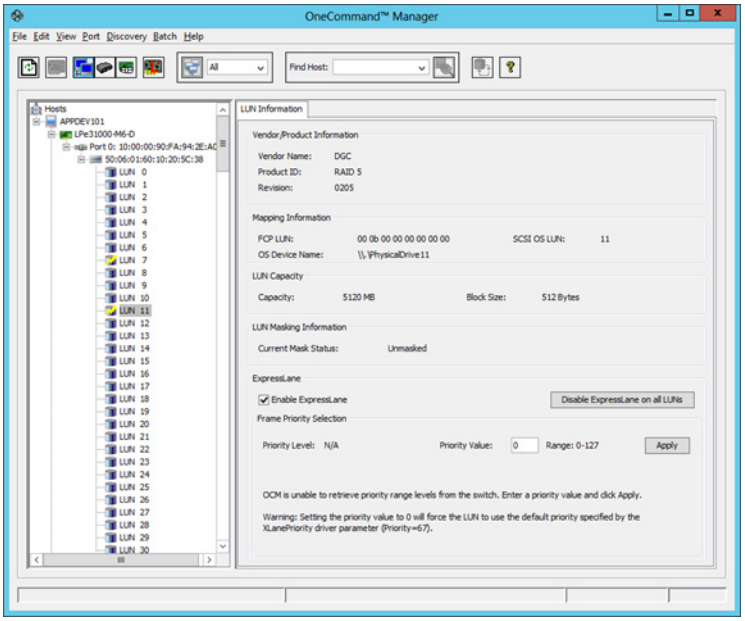

5. Click **Apply**.

   If problems occurred when assigning the frame priority, the **LUN Information** tab displays a message with a suggested solution.

8.1.16 Changing the FC and FCoE WWPN and WWNN

The **Maintenance** tab (Figure 48) enables you to change the WWPN and the WWNN of a selected FC or FCoE function. For example, you can use an installed adapter as a standby in case another installed adapter fails. By changing the standby adapter's WWPN or WWNN, it can assume the identity and configuration (for example, driver parameters, persistent binding settings, and so on) of the failed adapter.

NOTE You cannot change FC WWPN and WWNN when FA-PWWN is enabled on the adapter port.

Three options exist for referencing WWNs:

- Factory Default WWN – As shipped from the factory. This value cannot be changed.
- Non-volatile WWN – Values that are saved in non-volatile adapter's flash memory that survives a reboot or a power outage.
- Volatile WWN – A temporary value that is saved in volatile memory on the flash. If volatile WWNs are set, they are used instead of the non-volatile WWNs.

NOTE Volatile WWN changes require a warm system reboot to take effect. Volatile WWN changes are lost on systems that power-cycle the adapters during the reboot.

To avoid address conflicts, do not assign a WWPN with the same WWPN as another FC or FCoE function on your SAN. The OneCommand Manager application checks the WWPN you specify against all the other detected WWPNs and, if a duplicate is found, an error is displayed and the WWPN is not changed.

CAUTION Changing volatile WWNs takes the selected adapter offline. Make sure that this adapter is not controlling a boot device and all I/O activity on this adapter is stopped before proceeding, or unexpected behavior or data loss can result.

In an environment where preboot management exists, a WWPN or WWNN modified by the OneCommand Manager application can be overridden by preboot management, such as Lenovo System X BOFM and industry-standard CLP.

For example, in an environment with CLP or BOFM: the OneCommand Manager application modifies the WWNN or WWPN. The OneCommand Manager application requires a reboot to complete the change. After reboot, the CLP string is sent during system boot and rewrites the WWNN or WWPN, or EFIBoot finds the BOFM protocol and uses the default WWNN or WWPN by the command from the BOFM.

To change an FC or FCoE function's WWPN or WWNN, perform these steps:

1. Perform one of the following tasks:
   - From the **View** menu, click **Group Adapters by Host Name**.
   - From the toolbar, click  **Group Adapters by Host Name**.
   - From the **Host Grouping** menu, select **Group Adapter by Fabric Names**.
2. In the discovery-tree, select the FC or FCoE function that you want to change.
3. Select the **Maintenance** tab (Figure 48).

Figure 48 Maintenance Tab

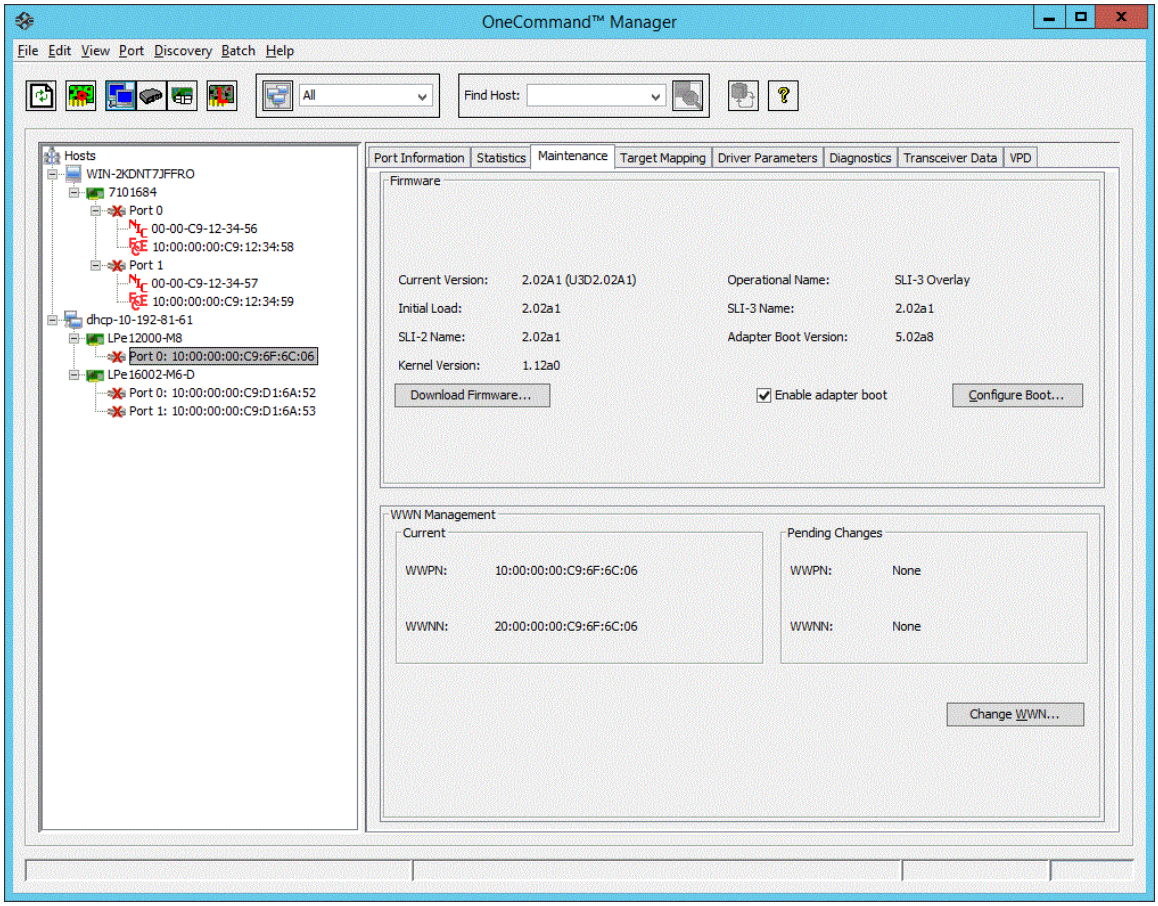

4. Click **Change WWN**. The warning in Figure 49 appears.

Figure 49 Warning About Changing WWN

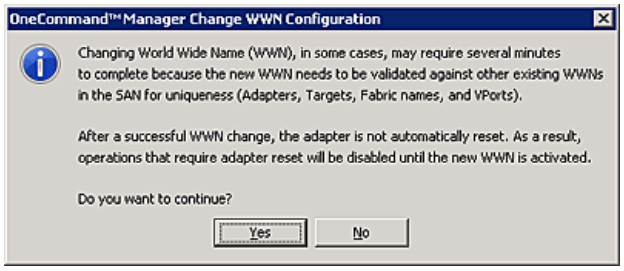

5. Click **Yes**. The **OneCommand Manager Change WWN Configuration** dialog appears (Figure 50).

Figure 50 OneCommand Manager Change WWN Configuration Dialog

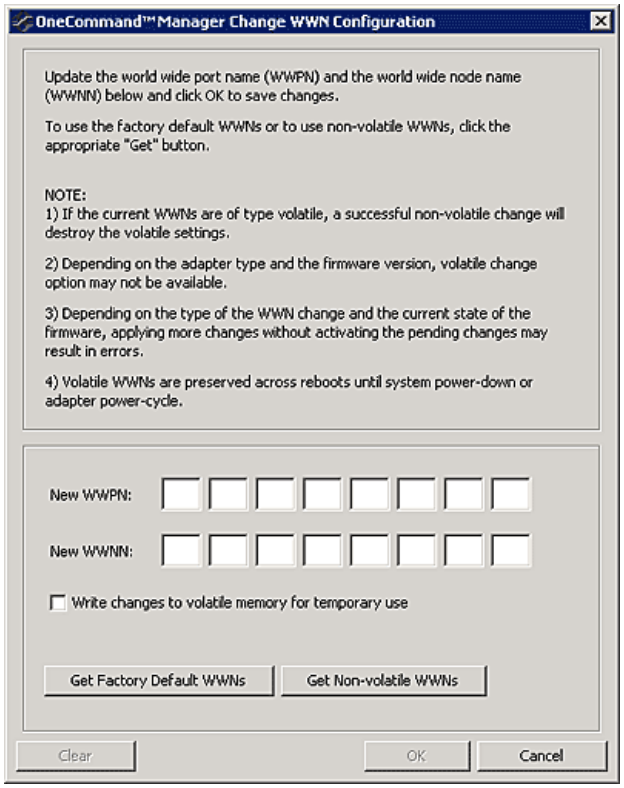

6. Perform one of the following tasks:
   - Enter a new WWPN and WWNN.
   - Click **Get Factory Default WWNs** to load the settings that were assigned to the FC or FCoE function when the adapter was manufactured to the New WWPN and WWNN settings. These values can then be modified if desired and saved as Volatile or Non-volatile WWNs.
   - Click **Get Non-volatile WWNs** to load the current non-volatile WWN settings to the New WWPN and WWNN settings. These values can then be modified if desired and saved to volatile or non-volatile memory. You can edit the data returned from the button.
7. Select **Write changes to volatile memory for temporary use** check box to save the New WWPN and New WWNN settings as Volatile WWNs. If cleared, the New WWPN and New WWNN settings are saved as non-volatile WWNs.

   NOTE If the adapter or firmware does not support Volatile WWNs, the **Write changes to volatile memory for temporary use** check box is disabled.
8. Click **OK**. After checking for a duplicate WWPN, the new WWPN and new WWNN values are saved for Volatile or Non-volatile use. The new WWPN and WWNN appear in the Pending Changes section in the WWN Management area of the **Maintenance** tab until the system is rebooted.
9. Reboot the system for the changes to take effect. After rebooting, the changes are applied and appear in the Current section of the **Maintenance** dialog.

8.1.17 Changing the FC and FCoE Port Names

NOTE This option is not available in read-only mode.

The OneCommand Manager application allows you to change the FC or FCoE adapter port names in the discovery-tree.

For example, you may want to identify a particular FC or FCoE function with the role it supports, such as a tape drive, scanner, or some other device. Use any characters you want for names, and names can be up to 255 characters in length. You can also revert to the adapter's default name.

NOTE Although you can change the FC or FCoE function's displayed name from the default WWPN, the change occurs in the discovery-tree (Figure 4) only. The function's WWPN is still active; it is simply replaced for display purposes with the name you enter. For example, the **Port WWN** field of the **Port Information** tab is not changed. Also, any changes you make to the names in your discovery-tree are seen only by you; users running the OneCommand Manager application on another host do not see your name changes.

To change the name of an FC or FCoE function, perform these steps:

1. From the discovery-tree (Figure 4), select the FC or FCoE port that you want to change by performing one of the following tasks:
   - Select **Edit Name** from the **Port** menu.
   - From the discovery-tree, right-click the port that you want to change and select **Change Name**.
2. Edit the name in the discovery-tree.

To use the FC or FCoE function's default name, perform these steps:

1. From the discovery-tree (Figure 4), select the FC or FCoE function that you want to change.
   - Perform one of the following tasks:
   - Select **Use Default Name** from the **Port** menu.
   - From the discovery-tree, right-click the port that you want to change and select **Restore Default Name**.

8.1.18 Resetting the FC and FCoE Functions

You can reset remote and local FC and FCoE functions.

CAUTION Do not reset functions while copying or writing files. This action could result in data loss.

NOTE For FCoE functions, a reset is necessary only to activate the updated driver parameters (that require a reset) or the FIP settings. It does not actually perform an adapter-function-level reset.

To reset the FC or FCoE function, perform these steps:

1. In the discovery-tree (Figure 4), select the FC or FCoE function that you want to reset.
2. Perform one of the following tasks:
   - From the **Port** menu, click **Reset Port**.
   - From the toolbar, click  **Reset**.

   The Reset Adapter warning appears (Figure 51).

Figure 51 Reset Adapter Warning

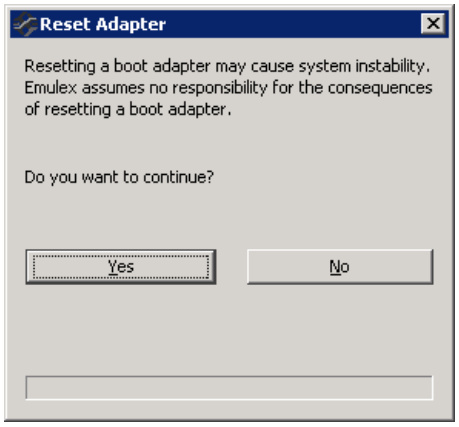

3. Click **Yes** to perform the reset.

   The reset can require several seconds to complete. While resetting, the status bar shows **Reset in progress**. When the reset is finished, the status bar shows **Reset Completed**.

8.1.19 Configuring the FC and FCoE Driver Parameters

NOTE This option is not available in read-only mode.

The OneCommand Manager application displays available driver parameters along with their defaults and maximum and minimum settings. A description of the selected parameter is also provided. This section contains information you must be aware of when working with driver parameters. For a more detailed description of specific driver parameters, refer to the appropriate Emulex driver user guide.

NOTE In Solaris and Linux, you can also specify parameters when loading the driver manually. Refer to the appropriate driver user guide for instructions.

8.1.19.1 Activation Requirements

A parameter has one of the following activation requirements:

- Dynamic – The change takes effect while the system is running.
- Reset – Requires an FC or FCoE reset from the utility before the change takes effect.
- Reboot – Requires reboot of the entire machine before the change takes effect. In this case, you are prompted to perform a reboot when you exit the utility.

8.1.19.2 Host Driver Parameters Tab

The **Host Driver Parameters** tab (Figure 52) enables you to view and edit the adapter driver parameter settings contained in a specific host. The host driver parameters are global values, and apply to all adapters in that host unless they are overridden by parameters assigned to a specific adapter using the adapter **Driver Parameters** tab. For each parameter, the tab shows the current value, the range of acceptable values, the default value, and whether the parameter is dynamic. A dynamic parameter allows the change to take effect without resetting the adapter or rebooting the system.

For information on changing parameters for a single adapter, see Section 8.1.19.3, Setting the Driver Parameters. For information on changing parameters for the host, see Section 8.1.19.9, Setting the Driver Parameters for All FC and FCoE Functions in a Host.

NOTE If there are no discovered FC or FCoE functions, the entire Host **Driver Parameters** tab (Figure 52) is dimmed. This event occurs because there are no drivers to which the host driver parameters apply.

Figure 52 Host Driver Parameters Tab

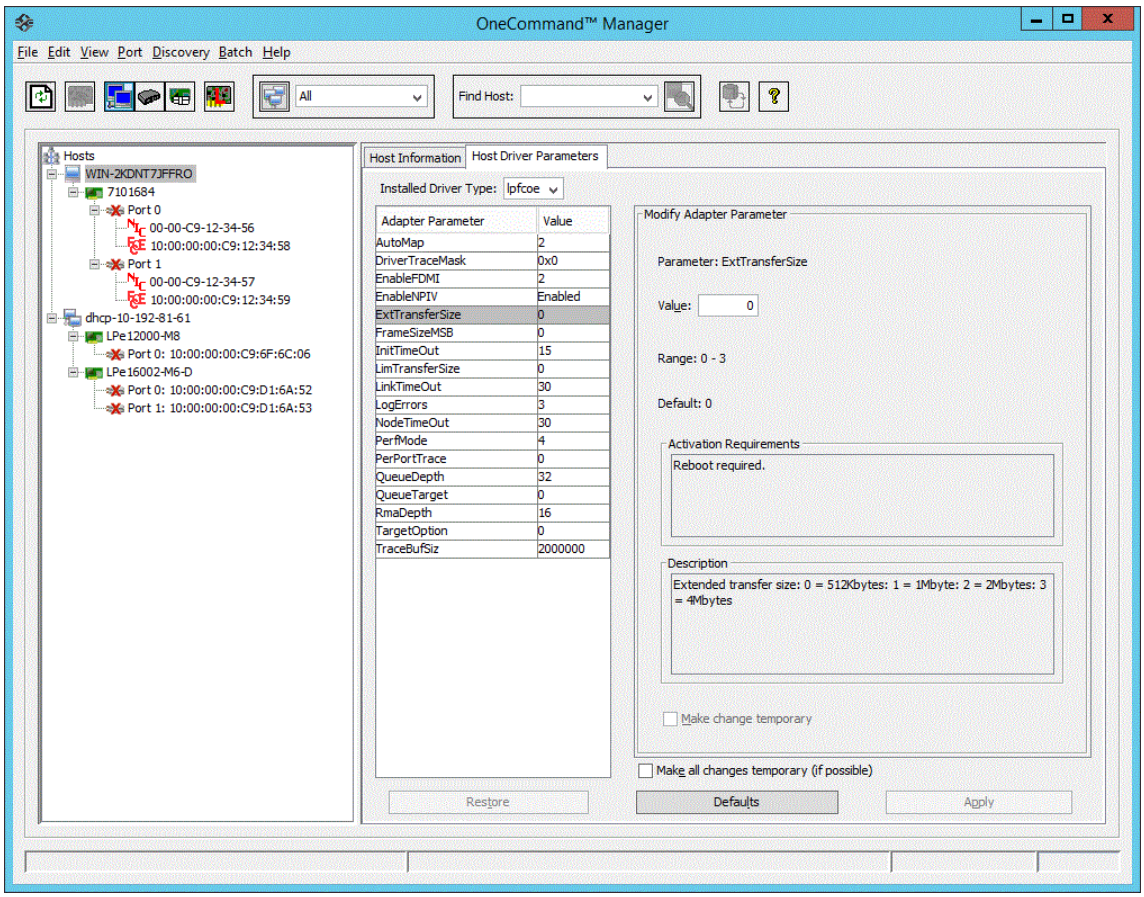

The following **Host Driver Parameters** tab fields are displayed:

- **Installed Driver Type** – The current drivers installed on this host. If more than one driver type is installed, the **Installed Driver Types** drop-down list shows a list of all driver types that are installed on the adapters in the host and enables you to select the particular driver type to configure.
- **Adapter Parameter table** – A list of adapter driver parameters for the selected driver type and their current values.
- Modify Adapter Parameter area:
  - Adapter-specific information is displayed in this area, which includes value, range, default, activation requirements, and description.

8.1.19.3 Setting the Driver Parameters

The **Driver Parameters** tab for FC or FCoE functions and hosts enables you to modify driver parameters for a specific FC or FCoE function or all FC and FCoE functions in a host.

For example, if you select a host in the discovery-tree (Figure 4), you can globally change the parameters for all FC and FCoE functions in that host. If you select an FC or FCoE function in the discovery-tree, you can change parameters for only that FC or FCoE function.

For each parameter, the **Driver Parameters** tabs show the current value, the range of acceptable values, the default value, and the activation requirement. You can also restore parameters to their default settings.

You can apply driver parameters for one FC or FCoE function to other FC and FCoE functions in the system using the **Driver Parameters** tab, thereby simplifying multiple adapter configuration. See Section 8.1.20, Creating a Batch Mode Driver Parameters File, for more information.

NOTE The Linux 2.6 kernel only supports setting some of the driver parameters for individual FC or FCoE functions. Some driver parameters must be applied to all FC and FCoE functions contained in the host. Refer to the *Emulex Driver for Linux for LightPulse Adapters User Guide* for more information.

8.1.19.4 Setting the Driver Parameters for a Single FC or FCoE Function

To change the driver parameters for a single FC or FCoE function, perform these steps:

1. Select the **Host** or **Fabric** view.
2. In the discovery-tree (Figure 4), select the FC or FCoE function that you want to change.
3. Select the **Driver Parameters** tab (Figure 53). The parameter values for the selected FC or FCoE function are displayed.

NOTE The LinkSpeed (Windows) or link-speed (Linux/Solaris) driver parameters are not shown if the adapter supports forced link speed. The link speed is configured using the **Firmware Parameters** tab. See Section 7.1.3, Viewing FC Firmware Parameters, for more information.

Figure 53 Driver Parameters Tab – Adapter Selected

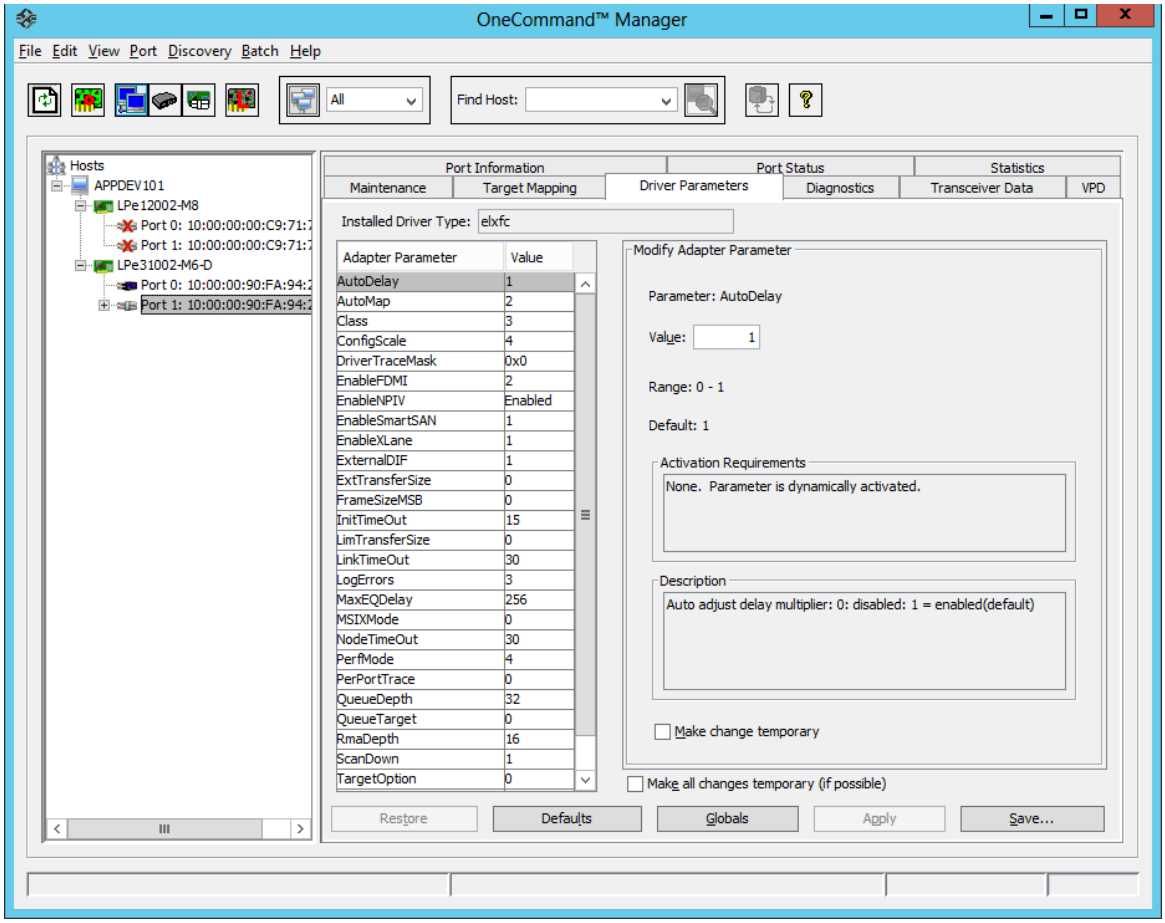

4. Click the parameter that you want to change. A description of the parameter appears on the right side of the tab.
5. Enter a new value in the **Value** field in the same hexadecimal or decimal format as the current value or select a value from the drop-down list. If you enter a value and the current value is in hexadecimal format, it is prefaced by 0x (for example, 0x2d). You can enter a new hexadecimal value without the 0x. For example, if you enter ff10, this value is interpreted and displayed as 0xff10.
6. If you want the change to be temporary (causing the parameter to revert to its last permanent setting when the system is rebooted), select the **Make change temporary** check box. This option is available only for dynamic parameters.
7. If you are making changes to multiple parameters, and you want all the changes to be temporary, select the **Make all changes temporary** check box. This setting overrides the setting of the **Make change temporary** check box. Only dynamic parameters can be made temporary.
8. Click **Apply**.

8.1.19.5 Restoring All Parameters to Their Earlier Values

If you changed parameters, but did not click **Apply** in the **Driver Parameters** tab (Figure 53) and you want to restore the parameters to their last saved values, click **Restore**.

8.1.19.6 Resetting All Default Values

To reset all parameter values to their default (factory) values, click **Defaults** in the **Driver Parameters** tab (Figure 53).

8.1.19.7 Setting an Adapter Parameter Value to the Host Adapter Parameter Value

To set an adapter parameter value to the corresponding host parameter value, perform these steps:

1. Select the **Host** or **Fabric** view.
2. In the discovery-tree (Figure 4), select the adapter port.
3. Select the **Driver Parameters** tab (Figure 53).
4. Click **Globals**. All parameter values are now the same as the global, or host, values.
5. To apply the global values, click **Apply**.

8.1.19.8 Saving the Adapter Driver Parameters to a File

To save a desired adapter parameter configuration for using with the Batch Driver Parameter Update feature, click **Save** in the **Driver Parameters** tab (Figure 53). To apply your configuration changes, click **Apply**.

Each definition is saved in a comma-delimited file with the following format:

<parameter-name>=<parameter-value>

The file is saved in the Emulex repository directory.

- In Windows: `\Program Files\Emulex\Util\Emulex Repository`
- In Linux: `/usr/sbin/ocmanager/RMRepository`
- In VMware ESXi: `/tmp/RMRepository`
- In Solaris: `/opt/ELXocm/RMRepository`

The OneCommand Manager application can then uses the Batch Driver Parameter Update function to apply these saved settings to all compatible adapters on the SAN.

NOTE Host driver parameters and persistent binding settings cannot be saved.

8.1.19.9 Setting the Driver Parameters for All FC and FCoE Functions in a Host

To change the driver parameters for all FC and FCoE functions installed in a host, perform these steps:

1. Perform one of the following tasks:
   - From the **View** menu, click **Group Adapters by Host Name**.
   - From the toolbar, click  **Group Adapters by Host Name**.
2. In the discovery-tree, click the host whose adapter driver parameters you want to change.
3. Select the **Host Driver Parameters** tab (Figure 54). If adapters with different driver types are installed, the **Installed Driver Types** menu shows a list of all driver types and driver versions that are installed. Select the driver whose parameters you want to change. This menu does not appear if all the adapters are using the same driver.
4. Click the parameter that you want to change. A description of the parameter appears on the right side of the tab.

Figure 54 Host Driver Parameters Tab – Host Selected

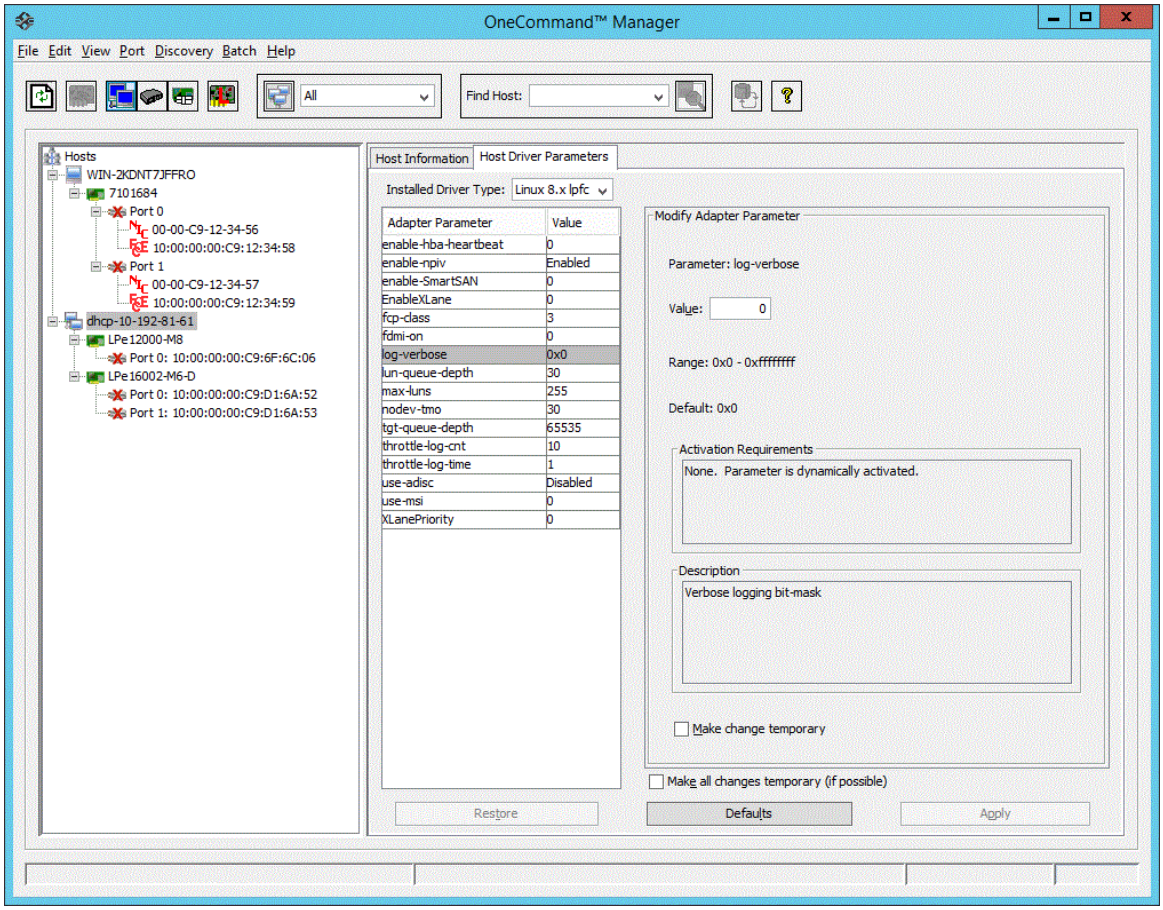

5. Enter a new value in the **Value** field in decimal or hexadecimal format, depending on how the current value is presented. If the value is in hexadecimal format, it is prefaced by 0x (for example, 0x2d).
6. To make a change temporary (the parameter will revert to its last permanent setting when the system is rebooted), select the **Make changes temporary** check box. This option is available only for dynamic parameters.
7. To make changes to multiple parameters, select the **Make all changes temporary (if possible)** check box. Only dynamic parameters can be made temporary.
8. Click **Apply**.

8.1.19.10 Changing Non-dynamic Parameter Values (Linux)

To change non-dynamic parameter values for Linux, perform these steps:

1. Navigate to the `/usr/sbin/ocmanager` directory, and run the scripts to stop the OneCommand Manager application processes. Type the following command:

   ```
   ./stop_ocmanager
   ```
2. Stop all I/O to FC and FCoE attached devices.
3. Unload the FC or FCoE driver. Type the following command:

   ```
   modprobe -r lpfc
   ```
4. Reload the driver. Type the following command:

   ```
   modprobe lpfc
   ```
5. Start the `elxhbamgr` service (remote service). Type the following command:

```
./start_ocmanager
```

The OneCommand Manager application discovery service starts automatically when you start the application.

NOTE For changes to persist after a reboot, you must create a new ramdisk image. Refer to the *Emulex Driver for Linux for LightPulse Adapters User Guide* for more information.

8.1.20 Creating a Batch Mode Driver Parameters File

NOTE This option is not available in read-only mode.

You can apply driver parameters for one FC or FCoE function to other FC or FCoE functions in the system using the **Driver Parameters** tab. When you save the driver parameters for an adapter, you create a `.dpv` file, which contains parameters for that adapter. After you create the `.dpv` file, the OneCommand Manager application enables you to assign the `.dpv` file parameters to multiple adapters in the system.

To create the `.dpv` file, perform these steps:

1. Select the **Host** or **Fabric** view.
2. Select the FC or FCoE function whose parameters you want to apply to other FC or FCoE functions from the discovery-tree (Figure 4).
3. Select the **Driver Parameters** tab (Figure 53).
4. Set the driver parameters.
5. After you define the parameters for the selected adapter, click **Apply**.
6. Click **Save**. The **Save Driver Parameters** dialog appears (Figure 55). You can save the file to a different directory or change its name.

Figure 55 Save Driver Parameters Dialog

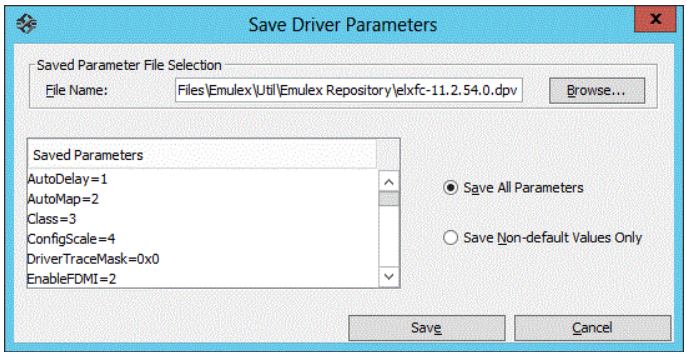

7. Use the two radio buttons to choose the type of parameters to save. You can save all parameters or only those parameters whose current values differ from their corresponding default values.

   A list of the saved parameters and their current values appear in the **Saved Parameters** list.
8. Click **Save**.

8.1.20.1 Assigning Batch Mode Parameters

To assign batch mode parameters to adapters, perform these steps:

1. From the **Batch** menu, select **Update Driver Parameters**. (You do not need to select any discovery-tree [Figure 4] elements at this time.)
2. When the **Batch Driver Parameter Update** dialog appears, click **Browse**.

Figure 56 Batch Driver Parameters Update Dialog

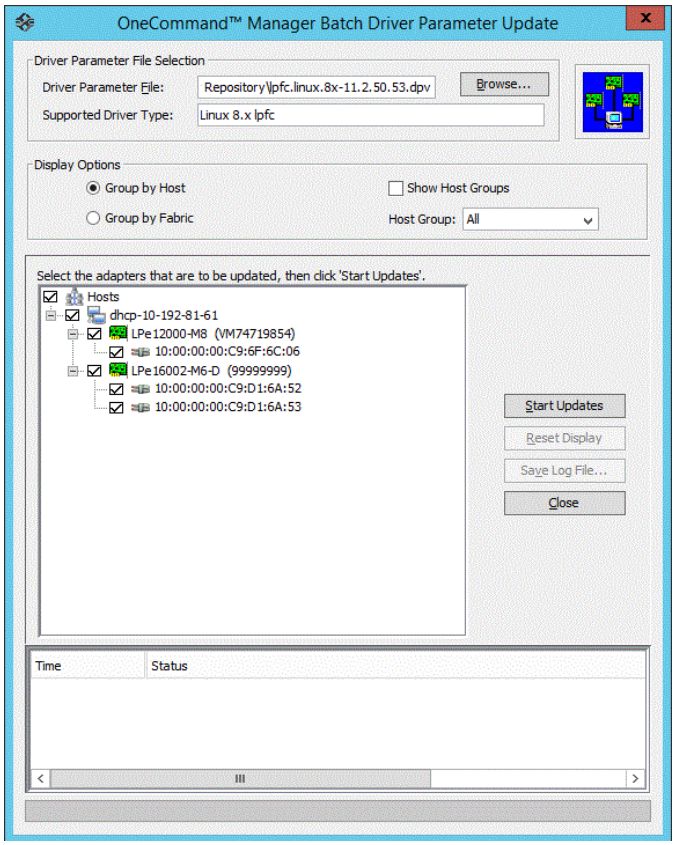

3. The **Driver Parameter File Selection** dialog appears (Figure 56). Select the file you want to use and click **OK**. A dialog appears notifying you that the OneCommand Manager application is searching for compatible adapters.

   After compatible FC and FCoE functions are found, the **Driver Parameter File** field of the **Batch Driver Parameter Update** dialog displays the selected file's path. The **Supported Models** text field displays a list of all adapter models that are compatible with the selected file. The set of compatible adapters appears in the dialog's discovery-tree.

   Using the Display Options settings, you can choose how adapters are displayed in the discovery-tree. Selecting the **Group by Host** radio button displays adapters in a host-centric view.Selecting the **Group by Fabric** radio button shows hosts in a fabric-centric view with their fabric addresses. The WWPN and host name for each available FC or FCoE function is displayed under its respective fabric.

   You can also display host groups by selecting the **Show Host Groups** check box. To display a particular host group, choose that group from the **Host Group** selection box.

   Check boxes next to the host, adapter, and FC and FCoE functions entries are used to select or clear an entry. Checking an adapter selects or removes all FC or FCoE functions on that adapter; checking a host removes or selects all eligible adapters for that host.

4. Make your selections, and click **Start Updates**. The OneCommand Manager application **Batch Driver Parameter Update** dialog (Figure 56) shows the current status of the update. When the update completes, a final summary shows the number of FC and FCoE functions that were successfully processed, and the number of FC and FCoE functions for which one or more parameter updates failed.

5. You can click **Save Log File** to save a report of the update.

8.1.21 Configuring FIP

The **FIP** tab enables you to configure FCoE Initialization Protocol (FIP) for FCoE functions on LPe16202/OCe15100 adapters in NIC+FCoE mode.

To configure FIP, perform these steps:

1. From the discovery-tree, select the FCoE function whose FIP properties you want to configure.
2. Select the **FIP** tab (Figure 57).
3. Set the parameters you want and click **Apply Changes**.

Figure 57 FIP Tab for FCoE Adapters

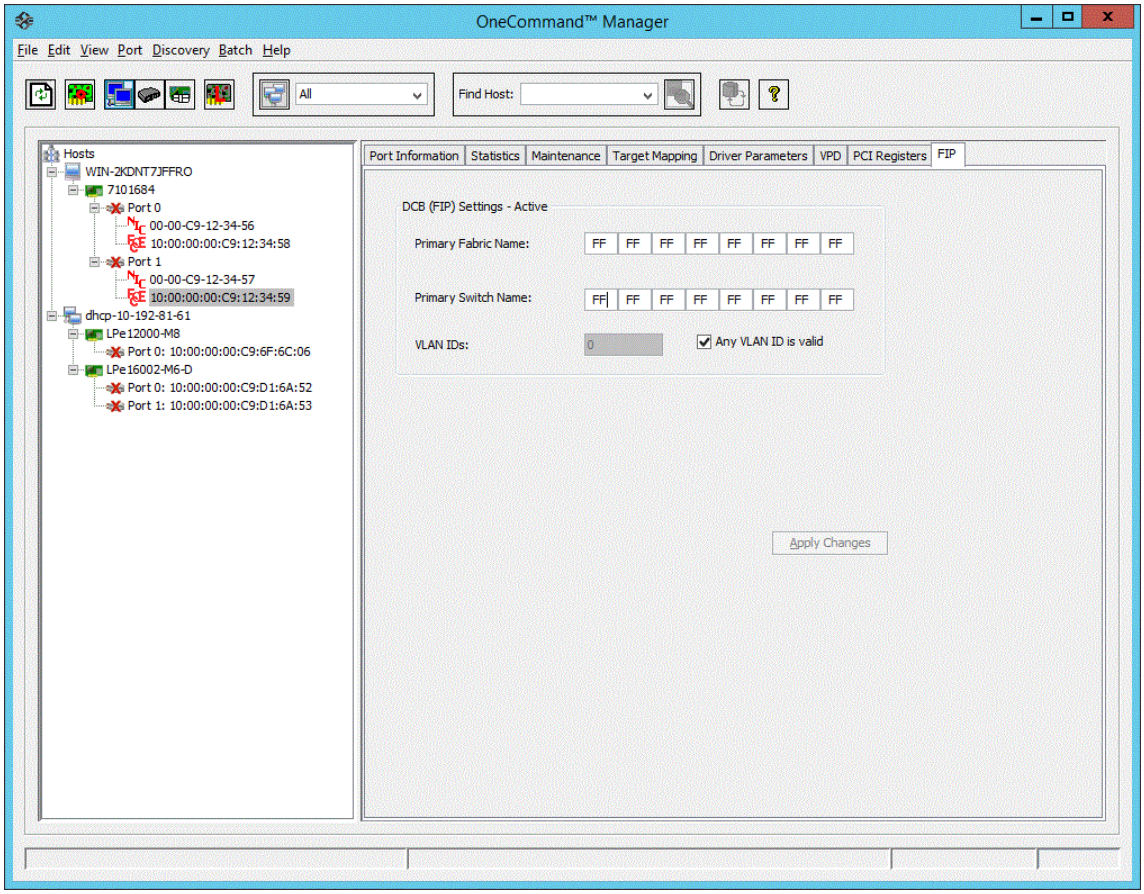

The following **FIP** tab field descriptions are displayed:

- **Primary Fabric Name** – Indicates the FC Fabric's WWN to which to connect. If the Primary Fabric Name is wild, that is, all 0xFFs, then connection to any fabric name is allowed.
- **Primary Switch Name** – Indicates the FC Switch's WWN to which to connect. If the Primary Switch Name is wild, that is, all 0xFFs, then connection to any switch name is allowed.
- **VLAN ID** – Determines the VLAN where the adapter FCoE services are available. It can have a value from 0 to 4095 and supports wild card values if **Any** is selected.
- **Any VLAN ID is valid** check box – If selected, the **VLAN ID** field of the FCoE forwarder can be any valid value.

8.1.22 Configuring Boot from an FC or FCoE SAN

You can use the OneCommand Manager application to configure a system to boot from an attached FC or FCoE LUN. Boot from SAN allows servers on a storage network to boot their operating systems directly from a SAN storage device, typically identified by its WWPN and a LUN located on the device. By extending the server system BIOS, boot from SAN functionality is provided by the BootBIOS contained on a Emulex adapter in the server. When properly configured, the adapter then permanently directs the server to boot from a LUN on the SAN as if it were a local disk.

NOTE Boot from SAN is not supported through the CIM interface.

8.1.22.1 Boot Types

Using the **Maintenance** tab, you can enable, disable, or configure boot from SAN for x86 BootBIOS, EFIBoot, and OpenBoot (also known as FCode).

- x86 BootBIOS works with the existing BIOS on x64 and x86 systems.
- OpenBoot (FCode) works with the existing system BIOS on Solaris SPARC systems using the SFS driver and on Linux PowerPC systems.
- EFIBoot works with x64-based systems and provides 64-bit system boot capability through the use of the EFI Shell.

Emulex provides Universal Boot and Pair Boot code images that contain multiple types of boot code. These images provide multiplatform support for boot from SAN. Universal Boot and Pair Boot transparently determine your system platform type and automatically run the proper boot code image in the adapter. These code images reside in adapter flash memory, allowing easier adapter portability and configuration between servers.

The adapters store the boot configuration data for each of these boot types.

NOTES

- x86 and OpenBoot share the same configuration memory space. You cannot configure an adapter for both x86 and OpenBoot at the same time. If you try, a message is displayed, stating that the existing boot type configuration will be overwritten by the new configuration.
- Boot from SAN configuration does not affect current system operation. The changes only take effect upon reboot if you have configured it correctly.

8.1.22.2 Boot Device Parameters

The boot LUN for all three boot types is in the range of 0 to 255. EFIBoot and OpenBoot (FCode) also support an 8-byte LUN, which you can use instead of the single-byte LUN. You must select which LUN type to configure.

- For OpenBoot, you must also provide a Target ID parameter.
- You must boot the host to configure boot from SAN with the OneCommand Manager application.
- You must work from a running host that supports the OneCommand Manager application. Often, this host has booted from a direct-attached drive. With the OneCommand Manager application, you can configure a direct boot host to boot from a SAN. You can modify an existing boot from SAN configuration or configure boot from SAN on an adapter for installation in another host so it can boot from SAN.
- You must know the boot code type that the adapter has; the OneCommand Manager application cannot detect this information. Without knowing this information, you could configure a boot type but not be able to boot from it because the adapter lacks the correct boot code.
- You must know the boot code type that the system supports; the OneCommand Manager application cannot detect this information. You can configure any boot type, but if the system does not support that type, it cannot boot from SAN.

- If you manage adapters on a remote host that is running a version of the OneCommand Manager application that does not support boot from SAN, the **Configure Boot** button does not appear.

  NOTE You can configure boot from SAN before boot by using the Emulex Boot BIOS setup command line interface that runs during system startup. Refer to the *Boot for LPe16202/OCe15100 Adapters User Guide* or the *Boot for the Fibre Channel Protocol User Guide* for details.

- One of the following FC or FCoE adapter drivers must be installed:
  - Windows: Storport Miniport driver
  - Linux: Emulex driver
  - Solaris: emlxs FCA Driver
  - VMware: Emulex driver

To configure boot from SAN, perform these steps:

1. Select the **Host** or **Fabric** view.
2. In the discovery-tree (Figure 4), click the FC or FCoE adapter function on which you want to enable boot from SAN.
3. Select the **Maintenance** tab (Figure 37 or Figure 38), select the **enable adapter boot** check box (if available), and click **Configure Boot**. The **Boot from SAN Configuration** dialog appears (Figure 58).

   NOTE The **Configure Boot** button is disabled if the **Enable Adapter Boot** check box is not selected. If boot code is not present on the adapter, the **Enable Adapter Boot** check box and **Configure Boot** button are not displayed on the **Maintenance** tab.

Figure 58 Boot from SAN Configuration Dialog

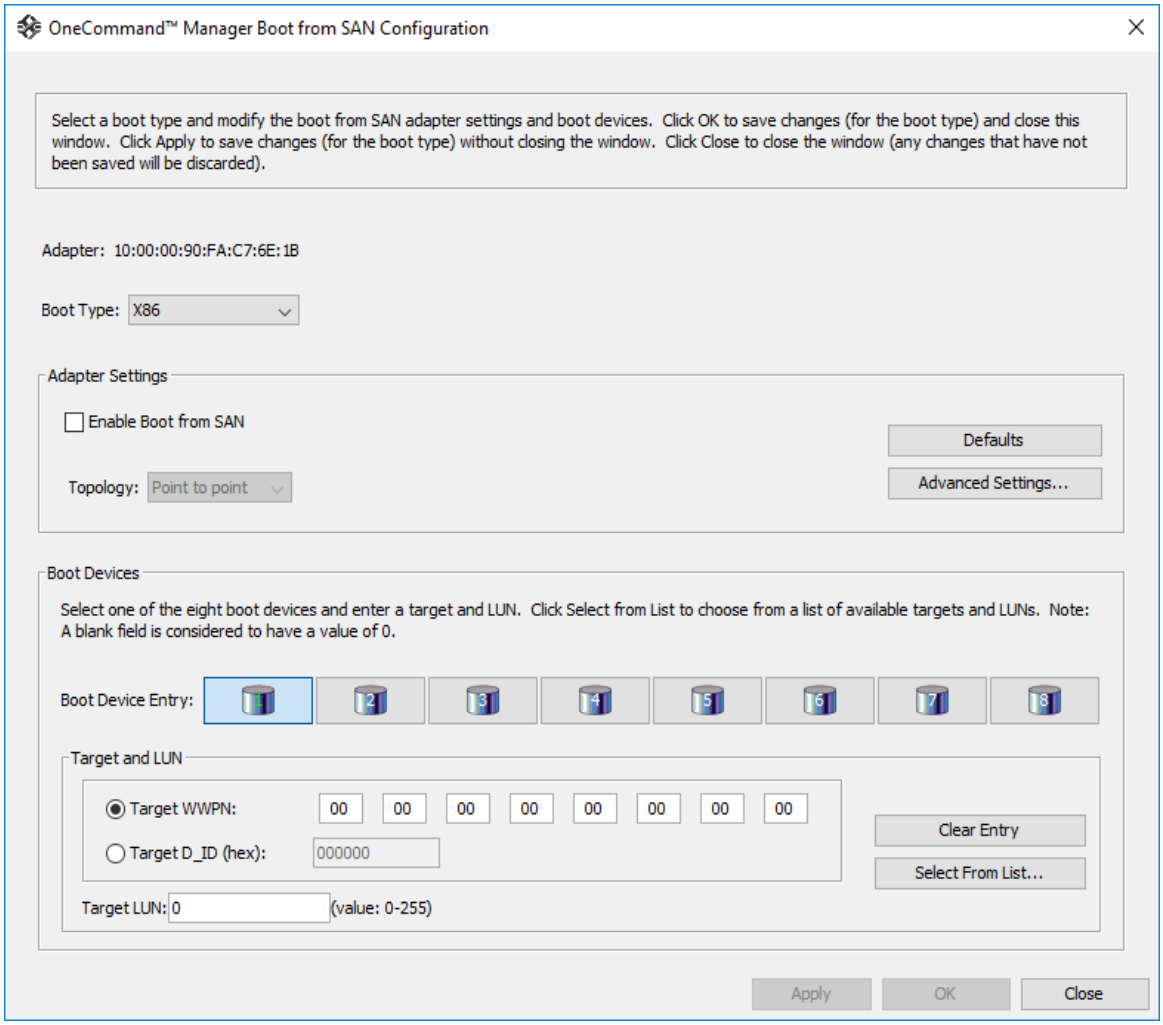

The **Boot from SAN Configuration** dialog varies for each boot type. Figure 58 depicts the boot from SAN configuration for the x86 type boot.

4. Verify that the **Adapter** field contains the WWPN of the FC or FCoE function and boot BIOS version to make sure you configure the correct adapter FC or FCoE function and that it has the BIOS boot code version you want.
5. From the **Boot Type** menu, select **x86**, **EFIBoot**, or **OpenBoot**.

NOTE **x86** and **OpenBoot** share the same configuration memory space. You cannot configure an adapter for both **x86** and **OpenBoot** at the same time. When you select one of these boot types and the configuration region is configured for the other boot type, a message appears warning that making changes overwrites the other boot-type configuration.

If you modified the settings for the current boot type and then change to a new boot type, a message appears telling you to save the current settings before changing to the new boot type.

6. Select **Enable Boot from SAN** check box, and for FC functions, set the Topology.

   Topology options are:

- **Auto, Loop First** (default)
- **Auto, Point to Point First**
- **Loop**
- **Point to Point**

7. To configure autoscan, spinup delay, and other advanced settings, see Section 8.1.23, Configuring Advanced Settings (Boot from SAN).
8. For x86 and EFIBoot, select one or more boot devices. For OpenBoot, select only one boot device.
9. Perform one of the following tasks on the **Boot from SAN Configuration** dialog (Figure 58):
   - Select **Target World Wide Port Names**, type the numbers, and click **OK**.
   - Select **Target D_ID**, type the numbers, and click **OK**.
   - Select **Target LUN**, type the number, and click **OK**.

     For EFIBoot and OpenBoot, type in an 8-byte LUN (hexadecimal) and a target ID for the LUN. Also, you must enter the LUN value in big endian order (most-significant byte, or big end first) and enter all 16 characters, including leading zeros.
   - Click **Select from List**, select the target from a list of discovered LUNs (if available), and click **OK** on the **Select Boot Device** window (Figure 59). You can manually enter the target and LUN from the **Boot from SAN Configuration** dialog; however, it is easier to select an existing LUN from this window (Figure 59.) The OneCommand Manager application attempts to update the boot parameters. If successful, a window appears with a confirmation message. Click **OK** on this confirmation window.

Figure 59 Select Boot Device Window (for x86 or EFIBoot)

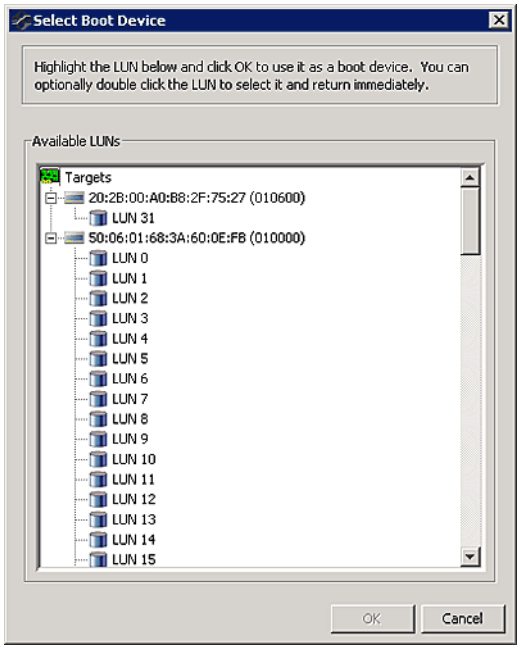

10. On the **Boot from SAN Configuration** dialog (Figure 58), click **Apply** to save your changes but leave the dialog open, or click **OK** to apply the changes and close the dialog.

    NOTE Click **Close** to close the **Boot from SAN Configuration** dialog without saving your changes. A message appears to discard your changes.
11. Reboot the system for your changes to take effect.

8.1.23 Configuring Advanced Settings (Boot from SAN)

The OneCommand Manager application provides advanced settings for each boot type. From the **Boot from SAN Configuration** dialog (Figure 58), click **Advanced Settings**. A boot type-specific dialog allows you to enable options such as spinup delay and autoscan. If you do not use advanced settings, the default values are used.

If you make changes you must click **OK** to save the changes and close the dialog. You can click **Cancel** and close the dialog without saving the changes.

NOTE If you do not enter the advanced settings and the configuration for the boot type is new, default values are used. The default settings are given with descriptions of the Advanced Adapter Settings dialogs in the following sections.

8.1.23.1 x86 Boot Advanced Adapter Settings Dialog

Use the **x86 Boot Advanced Adapter Settings** dialog (Figure 60) to configure **advanced settings** for the selected x86 adapter. All check boxes are cleared (off) by default. All changes require a reboot to activate.

Figure 60 x86 Boot Advanced Adapter Settings Dialog

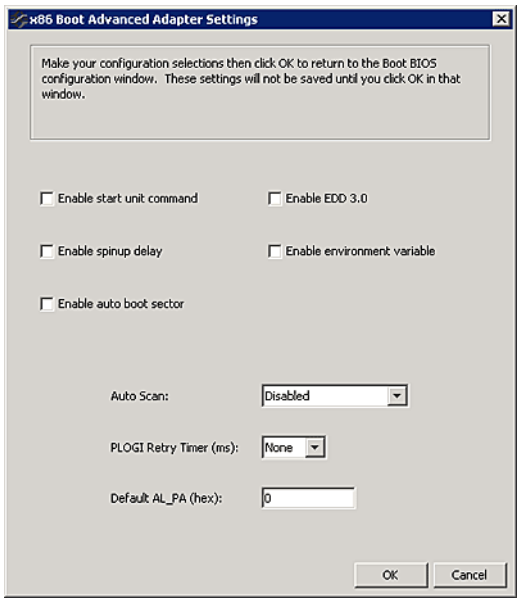

8.1.23.1.1 x86 Boot Advanced Adapter Settings Definitions

- **Enable start unit command** – Issues the SCSI start unit command. You must know the specific LUN to issue.
- **Enable EDD 3.0** – Enables the EDD option showing the path to the boot device. (Available on Intel Itanium servers only.)
- **Enable spinup delay** – If at least one boot device has been defined, and the spinup delay is enabled, the BIOS searches for the first available boot device.
  - If a boot device is present, the BIOS boots from it immediately.
  - If a boot device is not ready, the BIOS waits for the spinup delay and, for up to three additional minutes, continues the boot scanning algorithm to find another multi-boot device.

NOTE The default topology is auto topology with loop first. Change this topology setting, if necessary, before configuring boot devices.

  - If no boot devices have been defined and auto scan is enabled, then the BIOS waits for five minutes before scanning for devices.
  - In a private loop, the BIOS attempts to boot from the lowest target AL_PA it finds.
  - In an attached fabric, the BIOS attempts to boot from the first target found in the NameServer data.
- **Enable environment variable** – Sets the boot controller order if the system supports the environment variable.
- **Enable auto boot sector** – Automatically defines the boot sector of the target disk for the migration boot process, which applies only to HP MSA1000 arrays. If there is no partition on the target, the default boot sector format is 63 sectors.
- **Auto Scan** – With auto scan enabled, the first device issues a Name Server Inquiry. The boot device is the first DID, LUN 0, or not LUN 0 device returned, depending on the option you select. Only this device is the boot device, and it is the only device exported to the Multi-boot menu. Auto Scan is available only if none of the eight boot entries is configured to boot through DID or WWPN. Use the **Configure Boot Devices** menu to configure eight boot entries for fabric point-to-point, public loop, or private loop configurations. Set to one of the following values:
  - Disabled (default)
  - Any First Device
  - First LUN 0 Device
  - First non-LUN 0 Device
- **PLOGI Retry Timer (ms)** – Sets the interval for the PLOGI (port log in) retry timer. This option is especially useful for Tachyon-based RAID arrays. Under very rare occasions, a Tachyon-based RAID array resets itself, and the port goes offline temporarily in the loop. When the port comes back online, the PLOGI retry interval scans the loop to discover this device. The default setting is None (0 ms). Sets the PLOGI Retry Timer to one of the following values:
  - None (default)
  - 50 ms
  - 100 ms
  - 200 ms
- **Default AL_PA number (hex)** – This number has a range of 00 to EF (the default is 0). This value changes the AL_PA of the selected adapter. This option is not available for LPe16202/OCe15100 adapters in NIC+FCoE mode.

8.1.23.2 EFIBoot Advanced Adapter Settings Dialog

Use the **EFIBoot Advanced Adapter Settings** dialog (Figure 61) to configure the advanced settings for the selected EFIBoot adapter.

Figure 61 EFIBoot Advanced Adapter Settings Dialog

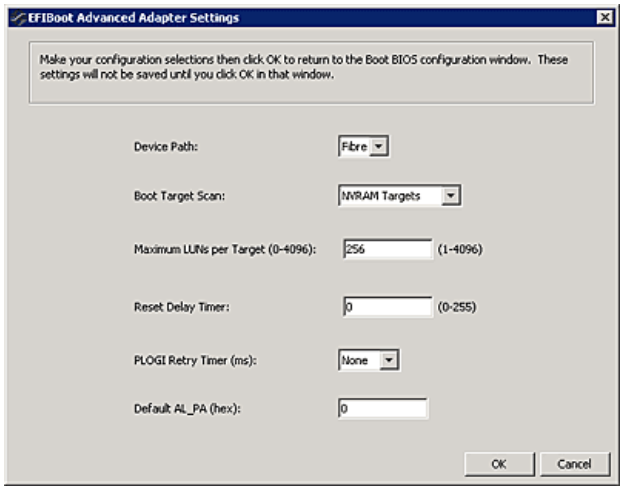

8.1.23.2.1 EFIBoot Advanced Adapter Settings Field Definitions

- **Device Path** – Makes the Fibre driver appear as a SCSI driver.
  - **Fibre (default)**
  - **SCSI**
- **Boot Target Scan** – This option is available only if none of the eight boot entries are configured to boot by using DID or WWPN.
  - **NVRAM Targets (default)** – Discovers only LUNs that are saved to the adapter nonvolatile random access memory (NVRAM).
  - **Discovered Targets** – Discovers all devices that are attached to the FC port. Discovery can take a long time on large SANs.
  - **None**.
  - **EFIBootFCScanLevel: NVRAM Targets and EFIBootFCScanLevel: Discovered Targets** – Allows third-party software to toggle between Boot Path from NVRAM and Boot Path from Discovered Targets by manipulating an EFI system NVRAM variable.
- **Maximum LUNs per Target** – Sets the maximum number of LUNs that are polled during device discovery. The range is 1 to 4096. The default is 256.
- **Reset Delay Timer in seconds** – Sets a value for delay device discovery. The range is 0 to 255. The default is 0.
- **PLOGI Retry Timer** – Sets the interval for the PLOGI (port log in) retry timer. This option is especially useful for Tachyon-based RAID arrays. Under rare occasions, a Tachyon-based RAID array resets itself and the port goes offline temporarily in the loop. When the port comes online again, the PLOGI retry interval scans the loop to discover this device.
  - **None** (default)
  - **50 ms**
  - **100 ms**
  - **200 ms**
- **Default AL_PA number** – The range is 0x 00 to EF. The default is 0x00. This option changes the AL_PA (Arbitrated Loop Physical Address) of the selected adapter (not available for LPe16202/OCe15100 adapters in NIC+FCoE mode).

8.1.23.3 OpenBoot Advanced Adapter Settings Dialog

Use the **OpenBoot Advanced Adapter Settings** dialog (Figure 62) to configure the advanced adapter settings for the selected OpenBoot adapter.

Figure 62 OpenBoot Advanced Settings Dialog

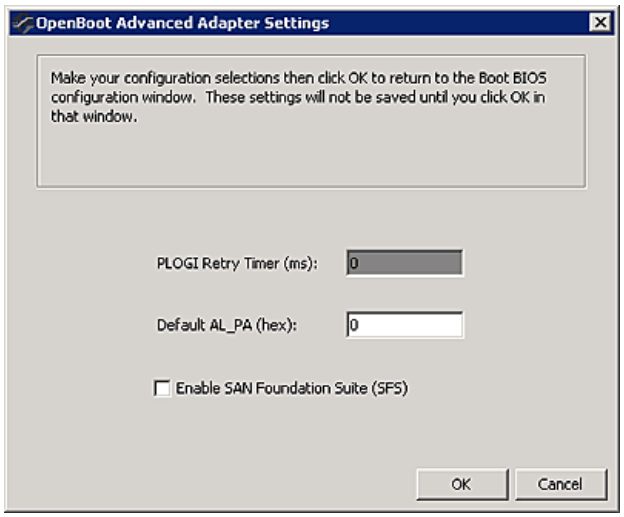

8.1.23.3.1 OpenBoot Advanced Adapter Field Definitions

- **PLOGI Retry Timer (ms)** – Sets the PLOGI Retry timer value. The range is 0 to 0xFF.
- **Default AL_PA (hex)** – Sets the default AL_PA. The range is 0 to 0xEF. The default is 0 (not available for LPe16202/OCe15100 adapters in NIC+FCoE mode).
- **Enable the SAN Foundation Suite (SFS)** – Select this check box to enable the SAN Foundation Suite (SFS) driver (the emlxs driver for Solaris).

8.1.24 Using FC-SP DHCHAP Authentication (Windows and Solaris)

Use the **DHCHAP** tab to view and configure FC-SP DHCHAP. You can authenticate an adapter to a switch.

8.1.24.1 DHCHAP Considerations

- DHCHAP is available only on LPe12000-series adapters.
- DHCHAP is available only for physical FC functions, not for virtual ports.
- The authentication driver parameters are only available on local hosts. The OneCommand Manager application GUI does not display this driver parameter for any remote hosts.

After DHCHAP has been activated and configured, manually initiate authentication per adapter by clicking **Initiate Authentication** or by inducing a fabric login (FLOGI) time in accordance with the FC-SP standard to the switch. A FLOGI can also be caused by bringing the link between the switch and adapter down and then up (not available in read-only mode).

Authentication must be enabled at the driver level. Authentication is disabled by default. To enable DHCHAP using the **Driver Parameters** tab, enable the enable-auth parameter in Windows or the enable-auth parameter in Solaris.

8.1.24.2 DHCHAP Tab

The **DHCHAP** tab (Figure 63) enables you to configure authentication.

Figure 63 DHCHAP Tab

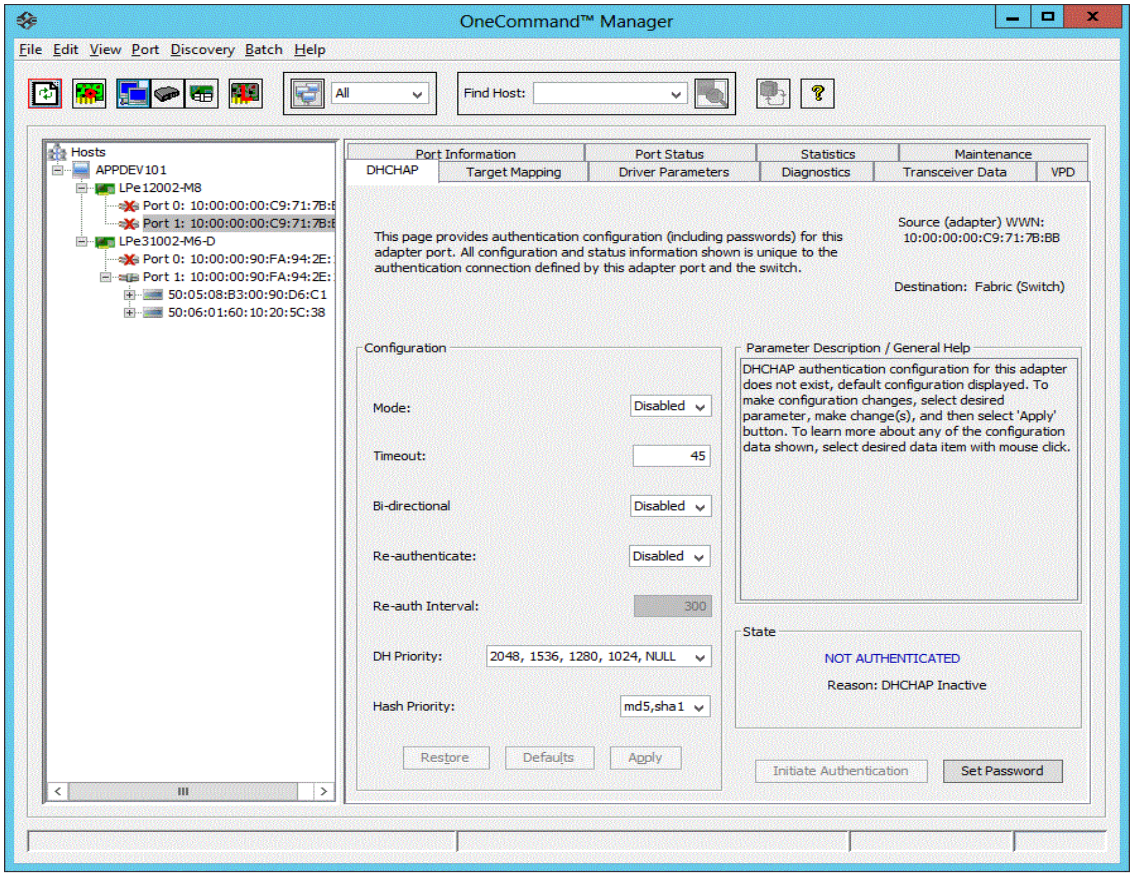

The following **DHCHAP** tab fields are displayed:

- **Source** – The WWPN of the FC function.
- **Destination** – The fabric (switch).
- Configuration area:
  - **Mode** – The mode of operation. Three modes are available:
    - **Enabled** – The FC function initiates authentication after issuing an FLOGI to the switch. If the connecting device does not support DHCHAP authentication, the software still continues with the rest of the initialization sequence.
    - **Passive** – The FC function does not initiate authentication, but participates in the authentication process if the connecting device initiates an authentication request.
    - **Disabled** – The FC function does not initiate authentication or participate in the authentication process when initiated by a connecting device. This mode is the default mode.
  - **Timeout** – During the DHCHAP protocol exchange, if the switch does not receive the expected DHCHAP message within a specified time interval, authentication failure is assumed (no authentication is performed). The time value ranges from 20 to 999 seconds.
  - **Bi-directional** – If enabled, the FC driver supports authentication initiated by either the switch or the FC function. If disabled, the driver supports only FC function initiated authentication.
  - **Re-authenticate** – If enabled, the FC driver can periodically initiate authentication.
  - **Re-auth Interval** – The value in minutes that the FC driver uses to periodically initiate authentication. Valid interval ranges are 10 to 3600 minutes. The default is 300 minutes.

- DH Priority – The priority of the five supported DH Groups (Null group, and groups 1, 2, 3, and 4) that the FC driver presents during the DHCHAP authentication negotiation with the switch.
- Hash Priority – The priority of the two supported hash algorithms (MD5 and SHA1) that the FC driver presents during the DHCHAP authentication negotiation with the switch (default is MD5 first, then SHA1 ,2 , 3...).

- State area:
  - State – Possible states are Not Authenticated, Authentication In Progress, Authentication Success, and Authentication Failed.

8.1.24.3 Changing Authentication Configuration

To view or change authentication configuration, perform these steps:

1. In the discovery-tree (Figure 4), select the FC function whose configuration you want to view or change.
2. Select the **DHCHAP** tab (Figure 63).

   NOTE If the fields on this tab are dimmed, authentication has not been enabled at the driver level.
3. Change the configuration values as you want.

   NOTE You can only configure DHCHAP on the local host.
4. Click **Apply**. You are prompted for the current password (local password) to validate the configuration change request. The verification request only appears if a local password has been defined for this adapter.
5. Enter the password, and click **OK**.

   To return settings to the status before you started this procedure, click **Restore** before you click **Apply**.

   NOTE If you click **Apply**, changes cannot be cancelled.

   To return all settings to the default configuration, click **Defaults**.

   CAUTION Be careful, because this action also resets any passwords to NULL for this configuration.

   To initiate an immediate authentication, click **Initiate Authentication**. This request is sent to the driver, even if you have not made any changes to the setup.

   NOTE To successfully authenticate with the switch using DHCHAP, you only have to set the configuration mode to enabled and set the local password. The local password must be set to the identical value as the switch for the DHCHAP authentication to succeed.

8.1.24.4 Changing Your Password

To change your password, perform these steps:

1. From the discovery-tree (Figure 4), select the FC function whose password you want to change.
2. Select the **DHCHAP** tab (Figure 63) and click **Set Password**. The **Password** dialog is displayed.
3. Choose **Set Local Password** or **Set Remote Password**.
   - The FC driver uses the local password when the adapter initiates authentication to the switch (typical use).
   - The FC driver uses the remote password if the switch authenticates with the adapter. This situation is only possible when bi-directional is selected on the **DHCHAP** tab (Figure 63).
4. If you want to see the password characters entered in the dialog, select the **Show Characters check box.**
5. Provide the current value for the password to validate the **set new password** request (unnecessary if this is the first time the password is set for a given adapter).
6. Enter the new password.

7. Select alphanumeric or hexadecimal format.
8. Click **OK**.

CAUTION Do not forget the password after one has been assigned. After a password is assigned to an adapter, subsequent DHCHAP configuration settings for that adapter including, the default configuration or new passwords, require you to enter the existing password to validate your request (that is, no further changes can be made without the password).

NOTE Additional help is available by clicking **Help** on the **Set Password** dialog.

8.1.24.5 Viewing the Error and Event Log

For Solaris systems, a simple shell script checks the `/var/adm/messages` file for recent Emulex driver DHCHAP events and outputs them to a default location.

To view the Error and Event Log, click **Event Log History** on the **DHCHAP** tab (Figure 63).

8.2 NIC Functions

This section describes managing NIC functions for LPe16202/OCe15100 adapters in NIC+FCoE mode.

NOTE If you change NIC driver properties (such as packet size) using the Windows Driver Properties applet while the OneCommand Manager application is running, the OneCommand Manager application may no longer be able to manage LightPulse adapters. As a result, most data fields show N/A, adapter settings are not configurable, and firmware download is not available. If this situation occurs, exit and restart the OneCommand Manager application.

8.2.1 Viewing NIC Function Information

When you select a NIC function from the discovery-tree (Figure 4), the NIC **Port Information** tab (Figure 64) contains general attributes associated with the selected NIC function.

The NIC **Port Information** tab also allows you to enable or disable network boot on the selected port.

To view general NIC function information, perform these steps:

1. Select the **Host** view.
2. Select NIC function in the discovery-tree (Figure 4).
3. Select the **Port Information** tab (Figure 64).

Figure 64 NIC Port Information Tab

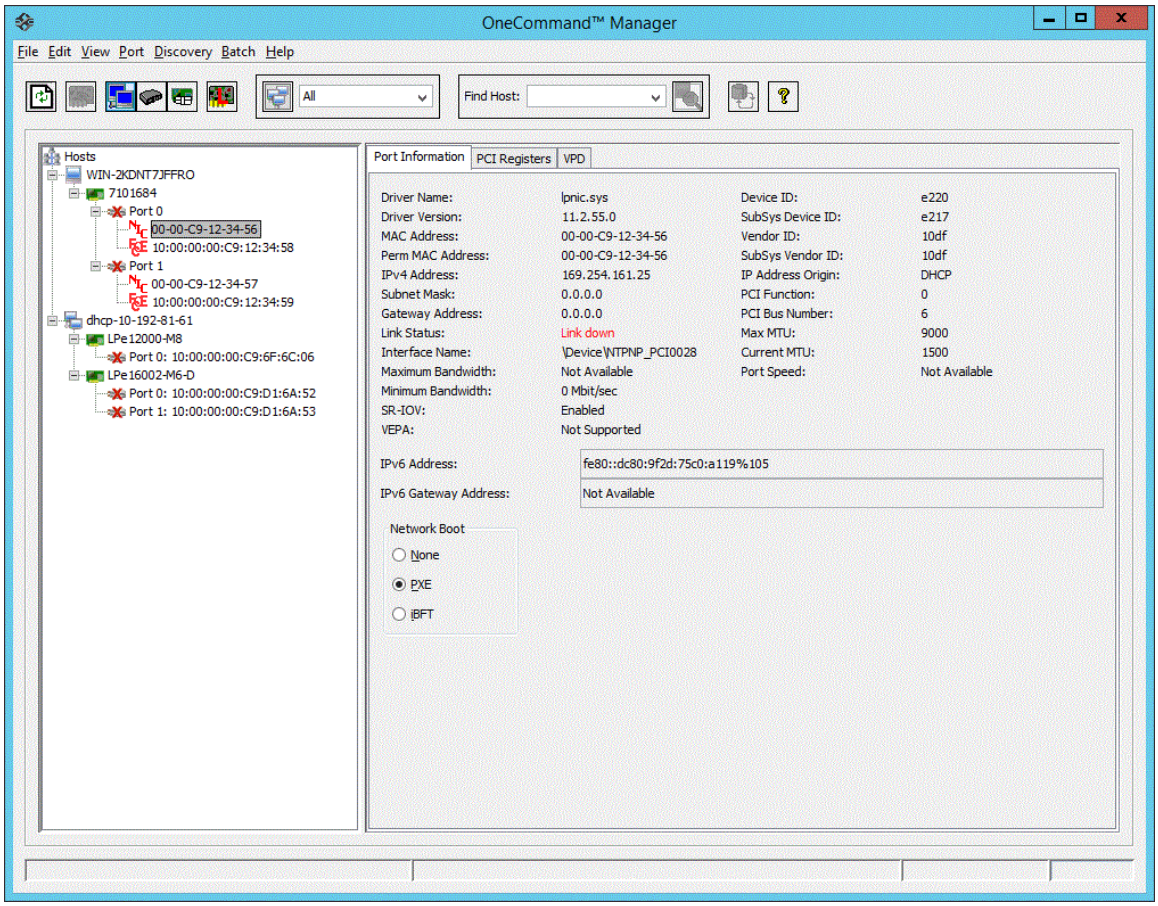

The following NIC **Port Information** tab fields are displayed:

- **Driver Name** – The NIC driver file name.
- **Driver Version** – The NIC driver version.
- **MAC Address** – The NIC MAC address currently assigned to the NIC function.
- **Perm MAC Address** – The original factory-assigned NIC MAC address.
- **IPv4 Address** – The IPv4 address for the NIC function.
- **Subnet Mask** – The subnet mask for the NIC function.
- **Gateway Address** – The NIC initiator gateway address.
- **Link Status** – The status of the link on the selected NIC function.

- **Interface Name** – The interface assigned to this NIC function by the host operating system.
- **Maximum Bandwidth** – The maximum bandwidth (that is, speed) at which the NIC function is guaranteed to run.
- **Minimum Bandwidth** – The minimum bandwidth (that is, speed) at which the NIC function is guaranteed to run.
- **SR-IOV** – If the adapter supports SR-IOV, it is automatically enabled; multiple VFs can be created on a NIC PF for an adapter port. These VFs are assigned by virtual operating systems, such as Microsoft Hyper-V and Linux KVM to virtual machines (VMs). Each VM can be assigned one or more VFs by the guest operating system running on the VM. To the guest operating system, the VF is an independent NIC function with its own MAC address, and it is available for network I/O.
- **VEPA** – Not supported on this adapter.
- **Device ID** – The PCI device ID assigned to the NIC function.
- **Subsys Device ID** – The PCI subsystem ID assigned to the NIC function.
- **Vendor ID** – The PCI vendor ID assigned to the NIC function.
- **Subsys Vendor ID** – The PCI subsystem vendor ID assigned to the NIC function.
- **IP Address Origin** – The origin of the IP address (DHCP or Static).
- **PCI Function** – The PCI function number assigned to the NIC function.
- **PCI Bus Number** – The PCI bus number assigned to the NIC function.
- **Max MTU** – The maximum transmission unit for Ethernet traffic.
- **Current MTU** – The current transmission unit for Ethernet traffic.
- **Port Speed** – The current port speed of the selected port.
- **Bandwidth Limit** – The QoS bandwidth restriction on the NIC function.
- **IPv6 Address** – The IPv6 address for the NIC function.
- **IPv6 Gateway Address** – The IPv6 gateway address for the NIC function.
- Network Boot area:
  - **None** – Select this radio button to disable network boot on the selected port.
  - **PXE** – Select this radio button to enable PXE boot on the selected NIC function.
  - **iBFT** – Not supported on this adapter.

NOTE PXE boot is supported only on the first function of each physical port on the adapter.

8.2.2 Viewing NIC VPD Information

The **VPD** tab (Figure 65) displays vital product data (if available) for the selected NIC function, such as the product name, part number, serial number, and so on.

To view VPD information, perform these steps:

1. Select the **Host** view.
2. In the discovery-tree, select the NIC function whose VPD information you want to view.
3. Select the **VPD** tab (Figure 65).

Figure 65 NIC VPD Tab

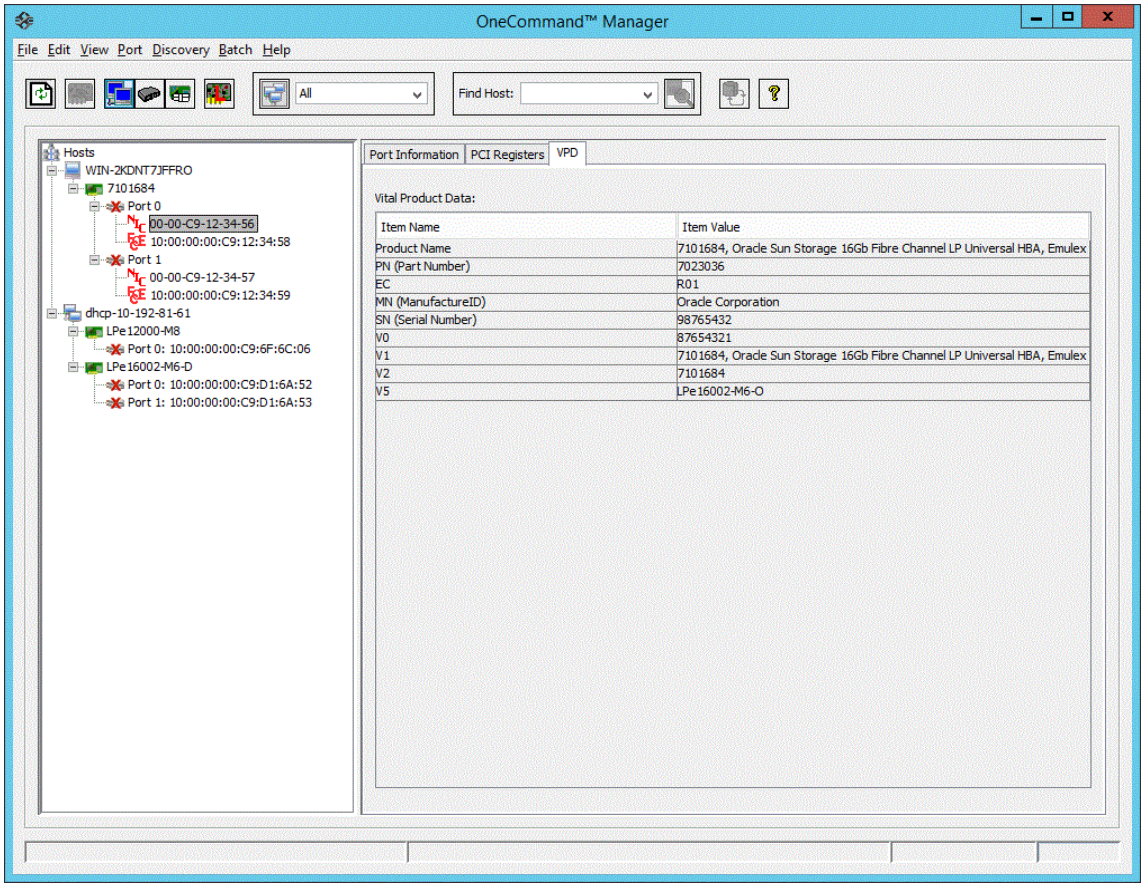

The following **VPD** tab fields are displayed:

- **Product Name** – Product information about the selected adapter port.
- **PN (Part Number)** – The adapter's part number.
- **SN (Serial Number)** – The adapter's serial number.
- **VO** – Vendor-unique data. **V** indicates a vendor-specific field. An adapter may have none, one, or more of these fields defined. Valid values for this field are VO (the letter O, not the number zero) and V*x* (where *x* is a number).

NOTE Some adapters may show additional VPD information such as EC (EC level) and MN (manufacturer ID).

8.2.3 Guest Operating System Discovery and Management from the Base Host Operating System

When the OneCommand Manager application is installed on a guest operating system, the guest operating system and VF are discovered by the OneCommand Manager application running on the host operating system. The guest operating system host appears as a remote host in the discovery-tree (Figure 66).

Figure 66 OneCommand Manager Application Running on the Base Host after Discovering the Guest Host

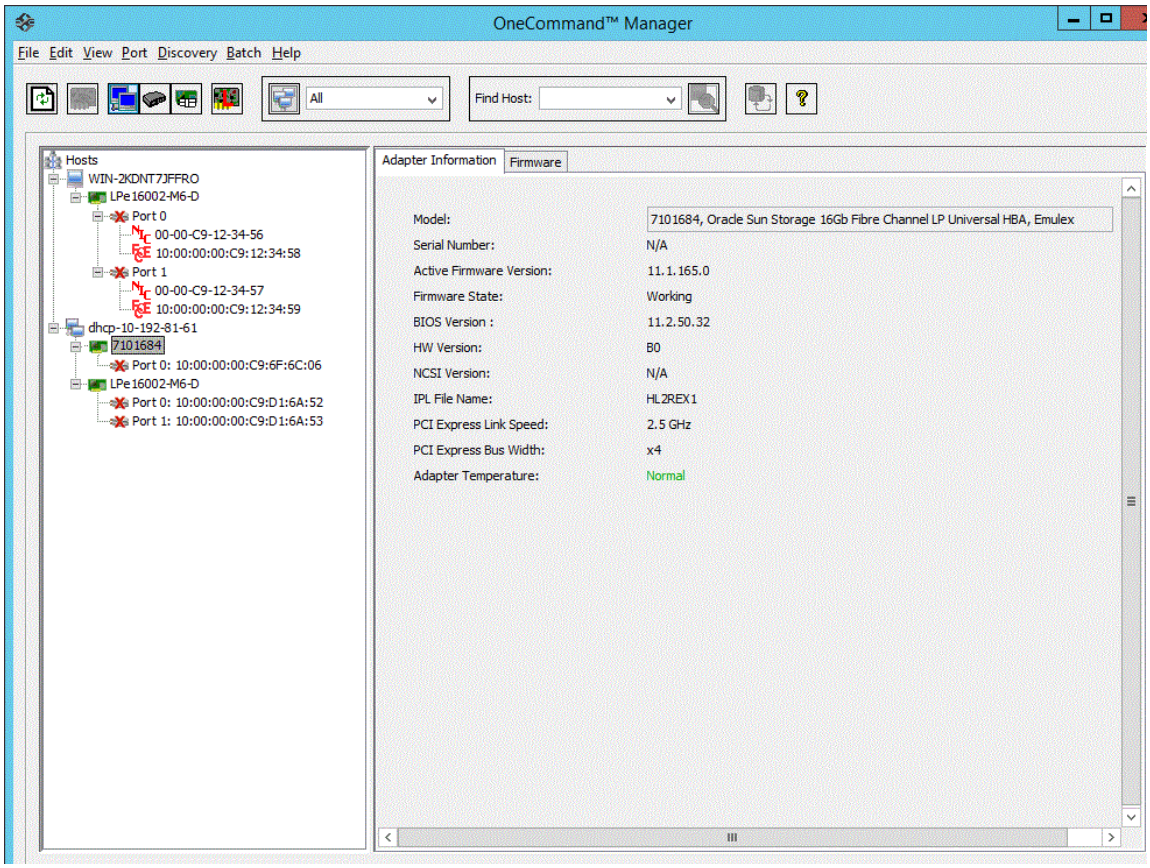

The NIC function for the discovered guest operating system matches one of the VFs in the base host operating system as seen in Figure 67.

Select the VF under the base operating system host in the discovery-tree to display the **Port Information** tab (Figure 67). Some of the information displayed in the tab is obtained from the remote management agent running on the guest operating system.

Figure 67 VF Selected Showing the Port Information Tab for the Discovered NIC in the Guest Operating System

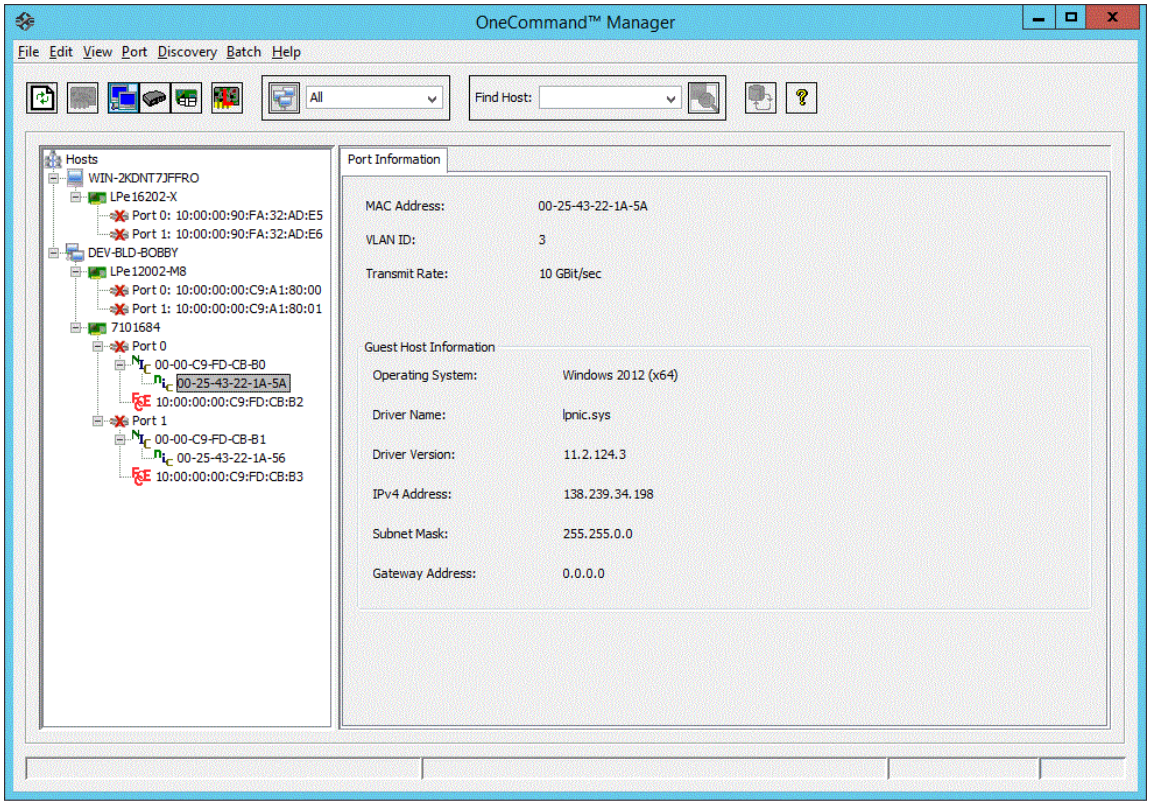

The following **Port Information** tab fields are displayed:

- **MAC Address** – The NIC MAC address currently assigned to the port.
- **VLAN-ID** – The VLAN identifier used by the NIC port.
- **Transmit Rate** – The rate at which data is transmitted over the port in MB/s.
- Guest Host Information area:

  NOTE The following fields are only available when the OneCommand Manager application is installed on the guest operating system running on this VF and the guest host has been discovered (over TCP/IP) by the OneCommand Manager application running on the base host operating system. See Section 5.1, Discovery Using the TCP/IP Access Protocol.

  — **Operating System** – The operating system and version installed on the selected host.
  — **Driver Name** – The NIC driver file name.
  — **Driver Version** – The NIC driver version.
  — **IPv4 Address** – The IPv4 address for the NIC port.
  — **Subnet Mask** – The subnet mask for the NIC port.
  — **Gateway Address** – The NIC initiator gateway address.

8.2.3.1 Running the OneCommand Manager Application on a Guest Operating System

When the OneCommand Manager application is installed and runs on a VM's guest operating system, it runs in a local-only/read-only mode. Therefore, only the guest host containing the adapter and NIC port is displayed in the discovery-tree. Because active NIC property management is not available, OneCommand Manager, running on the guest operating system, cannot perform operations that could adversely affect the host operating system or other guest operating systems. These operations include, firmware download, diagnostics, or DCB changes from the guest operating system. By preventing remote access to the host operating systems, the guest operating system cannot affect the operation of the adapter.

8.2.4 Viewing NIC PCI Registers

The NIC **PCI Registers** tab (Figure 68) displays base PCI registers. See Section 12.1.1.3, Viewing the PCI Registers for FC PCI register information.

Figure 68 NIC PCI Registers Tab

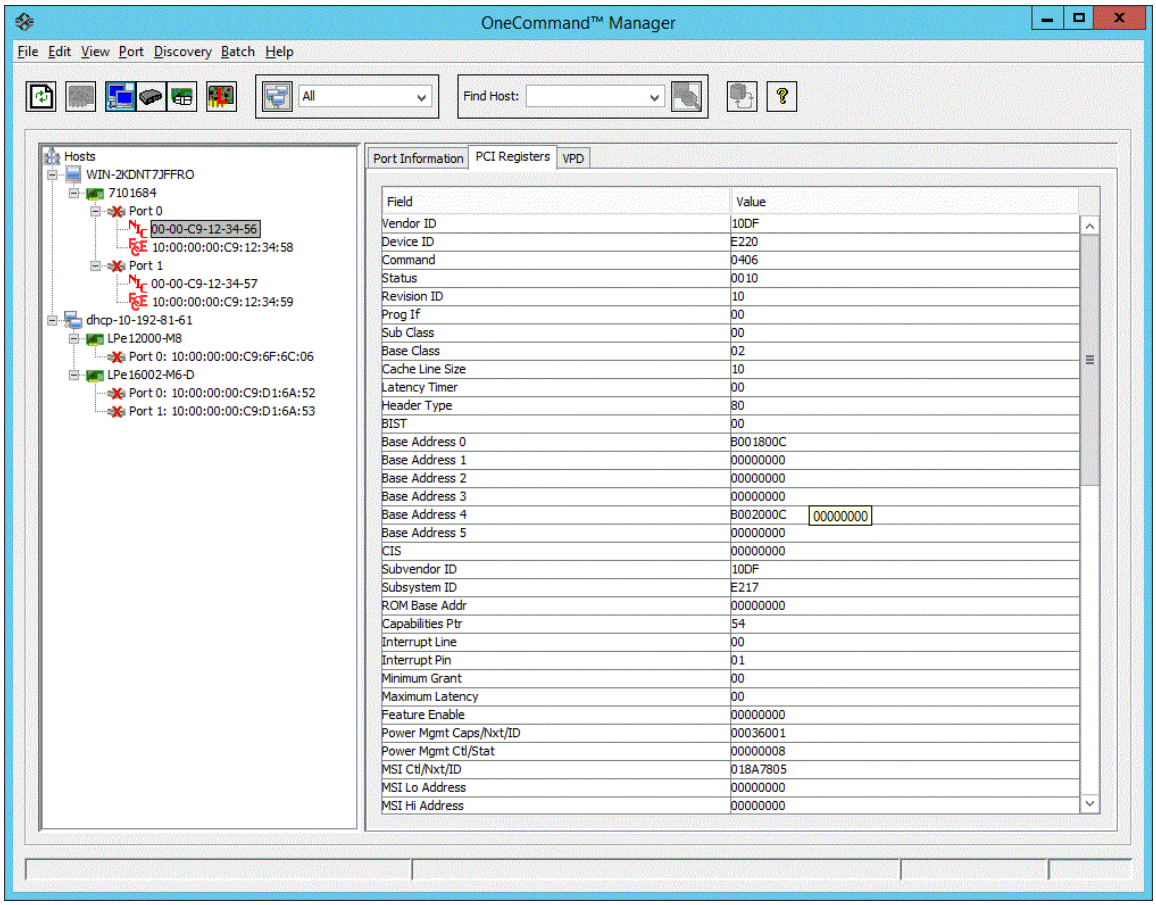

To view NIC PCI registers, perform these steps:

1. From the discovery-tree, select the NIC function whose PCI information you want to view.
2. Select the NIC **PCI Registers** tab (Figure 68).

Chapter 9: Updating Adapter Firmware

The OneCommand Manager application enables you to update firmware for a single adapter or simultaneously for multiple adapters.

CAUTION Updating firmware or boot code on an LPe12000-series adapter that is being used to boot from SAN may cause unpredictable behavior. After the update has completed, an adapter reset is issued, which may cause a loss of connectivity to the SAN and possible loss of data.

To update firmware on an LPe12000-series adapter, make sure that the adapter is not currently being used to boot from SAN. Do one of the following:

- Move the adapter to be updated to a non-boot from SAN host, and perform the update from that location.
- If the host with the target adapter is also hosting other boot from SAN adapters, perform a boot from SAN using one of the other boot from SAN adapters. The target adapter can now be updated.

9.1 Updating Firmware for a Single Adapter

NOTE This option is not available in read-only mode.

Using the **Maintenance** or **Firmware** tab, you can update firmware on local and remote adapters. Before you can perform this procedure, do the following:

- Download the firmware file from http://www.broadcom.com to a local drive.
- Make sure that the Emulex driver is installed.
- Make sure that the OneCommand Manager application is installed.
- If the adapter is already connected to a boot device, check that the system is in a state in which this type of maintenance can be performed:
  - I/O activity on the bus has been stopped.
  - Cluster software, or any other software that relies on the adapter to be available, is stopped or paused.

NOTES

- For vendor-branded adapters, refer to the vendor's website or contact the vendor's technical support for the firmware files.
- You cannot update firmware with the OneCommand Manager application on an Oracle-branded adapter.
- For LPe12000-series adapters, you update the firmware and boot code on each FC port/function. The firmware and boot code are two separate binaries. You must flash both the firmware and boot binaries to update LPe12000-series adapters.
- For LPe16000-series, LPe31000-series, and LPe32000-series adapters, you update the firmware for the entire adapter.

To update firmware for a single adapter or an adapter port, perform these steps:

1. Select the **Host** or **Fabric** view.
2. In the discovery-tree (Figure 4), select the adapter or FC port whose firmware you want to update.
3. Select the **Maintenance** tab for LPe12000-series adapters (Figure 40) or the **Firmware** tab for all other adapters (Figure 24), and click **Download Firmware**. The **Firmware Download** dialog appears (Figure 69).

Figure 69 Firmware Download Dialog

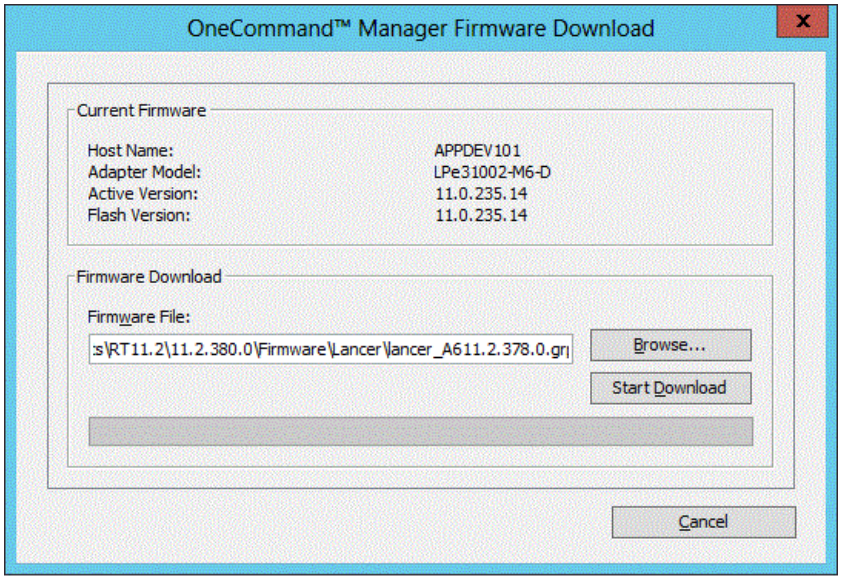

4. Using the **Firmware Download** dialog (Figure 69), navigate to the image file you want to download. The firmware image may be specified either by entering the image file's full path name in the **Firmware File** field or by clicking **Browse**.

   If you click **Browse**, the **Firmware File Selection** dialog appears. Select the file you want to use and click **OK**. The **Firmware Download** dialog appears.
5. Click **Start Download**. A warning dialog appears.
6. Click **Yes** to continue.

   A status bar shows the progress of the download. The adapter in the discovery-tree (Figure 4) is displayed in black text when the update is complete.

   NOTE The adapter in the discovery-tree may change to blue during the download, but this is normal.
7. Click **Close**. The **Firmware** tab displays the updated firmware information for the selected adapter.

   If you are updating the firmware on a multiport FC adapter, repeat steps 1 through 7 to update the firmware on the second port or use the Section 9.2, Updating Firmware for Multiple Adapters, procedure.

9.2 Updating Firmware for Multiple Adapters

Use batch mode to install firmware on multiple adapters in a single step. Batch firmware loading is restricted to a single firmware file and to all accessible, compatible adapters. Batch mode is not available in read-only mode.

NOTE Stop other OneCommand Manager application functions while batch loading is in progress.

Before you can perform a batch update, the firmware file must be downloaded from http://www.broadcom.com to a directory on your local drive.

NOTE VMware ESXi hosts managed through the CIM interface lists all adapters regardless of whether the selected firmware can update the adapter. You must manually deselect the nonmatching adapters.

To update firmware for multiple adapters, perform these steps:

1. From the **Batch** menu, select **Download Firmware**.

   You do not need to select a particular discovery-tree element for this operation.
2. When the **Batch Firmware Download** dialog appears (Figure 70), click **Browse**.

Figure 70 Batch Firmware Download Dialog, Selecting Adapters to Update

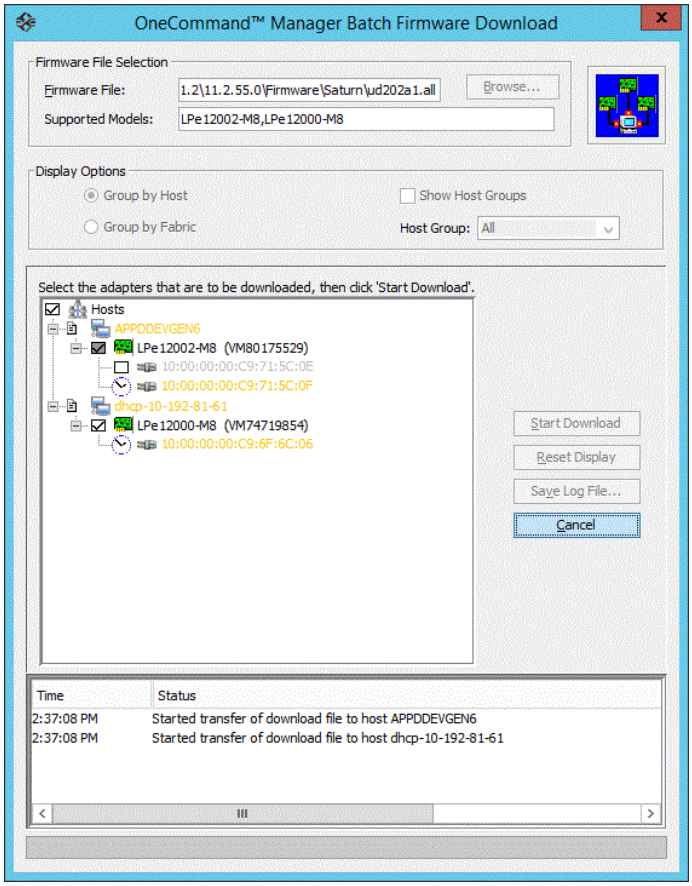

3. The **Firmware File Selection** dialog appears. Select the file you want to use and click **OK**. A dialog appears notifying you that the OneCommand Manager application is searching for compatible adapters.

   After compatible adapters are found, the Firmware File text area of the main **Batch Download** dialog displays the selected image file's path. The **Supported Models** text field displays a list of all adapter models that are compatible with the selected image file. The set of compatible adapters appears in the dialog's discovery-tree.

   Using the Display Options settings, you can choose how adapters are displayed in the discovery-tree. Selecting the **Group by Host** radio button displays adapters in a host-centric view. Selecting the **Group by Fabric** radio button shows hosts in a fabric-centric view with their fabric addresses. The WWPN and host name for each downloadable port is displayed under its respective fabric.

   You can also display host groups by selecting the **Show Host Groups** check box. To display a particular host group, choose that group from the **Host Group** menu.

   Check boxes next to the host and adapter entries are used to select or clear an entry. Selecting an adapter selects or removes that adapter; selecting a host removes or selects all eligible adapters for that host.

   For adapters where each individual port or ASIC can have new firmware downloaded, you can select the ports or ASICs on the adapter to which you want to download firmware.

4. Make your selections, and click **Start Download**. When downloading begins, the tree-view displays the progress. As firmware for a selected adapter is being downloaded, it appears orange in the tree-view. After successful downloading is complete, the entry changes to green. If the download fails, the entry changes to red (Figure 71).

Figure 71 Batch Firmware Download Dialog, Download Complete

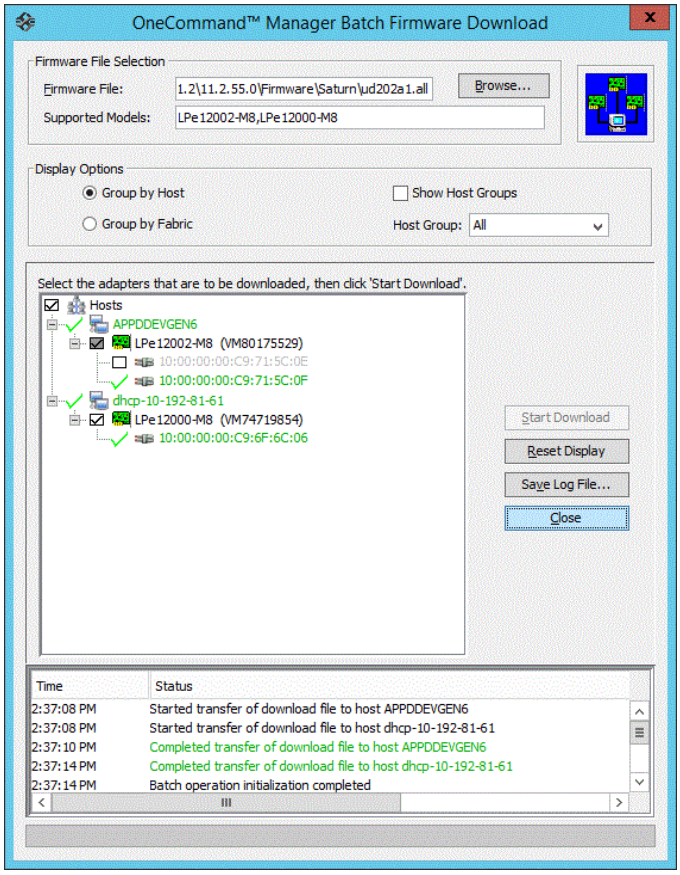

When downloading is finished, click **Save Log File** to save a copy of the activity log.

Chapter 10: DCB Configuration

This chapter describes how to configure DCB parameters on LPe16202/OCe15100 NIC+FCoE adapter ports.

To view the DCB parameters for NIC+FCoE adapter ports, perform these steps:

1. From the discovery-tree, select the NIC+FCoE adapter port whose DCB properties you want to view.
2. Select the **DCB** tab (Figure 72).

Figure 72 DCB Tab (NIC+FCoE Adapter Port Selected)

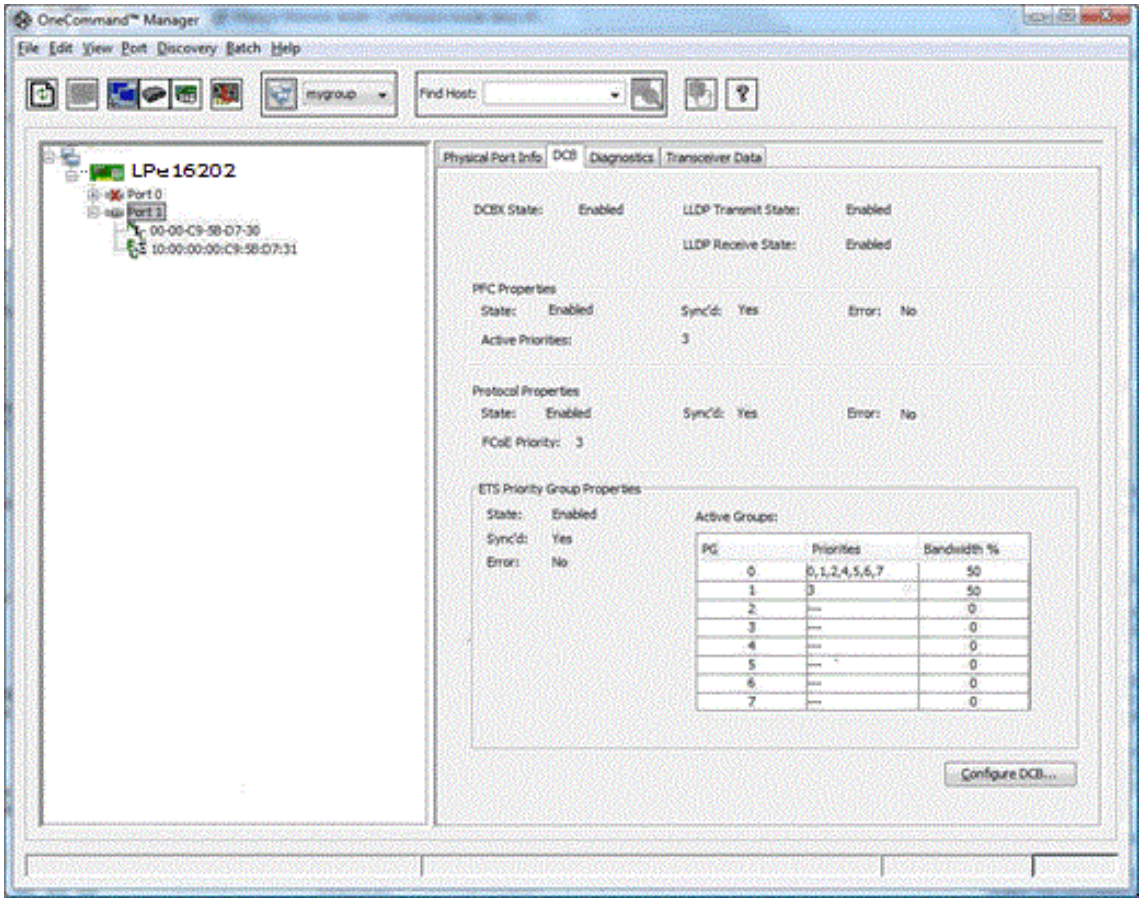

The following DCB tab fields are displayed:

- **DCBX State** – The current DCBX state (enabled or disabled).
- **LLDP Transmit State** – DCBX uses LLDP to exchange parameters between two link peers. For the DCBX protocol to operate correctly, both LLDP Rx and Tx must be enabled. If either Rx or Tx is disabled, DCBX is disabled.
- **LLDP Receive State** – DCBX uses LLDP to exchange parameters between two link peers. For the DCBX protocol to operate correctly, both LLDP Rx and Tx must be enabled. If either Rx or Tx is disabled, DCBX is disabled.
- PFC Properties area:
  - **State** – Enabled means that flow control in both directions (Tx and Rx) is enabled. Disabled means that priority-flow control is currently disabled. The priority value, if shown, is not applicable. This situation may be caused by:
    - The switch port priority-flow control being set to On instead of Auto.
    - The switch port uses port flow control instead of priority flow control.

    - PFC is disabled at the adapter or switch.
  - **Active Priorities** – Lists the priorities with PFC set to enabled.
  - **Sync'd** – If yes, the PFC priorities have been set by the peer. This parameter cannot be set.
  - **Error** – The error state. This capability indicates whether an error has occurred during the configuration exchange with the peer. Error is also set to yes when the Compatible method for the capability fails.
- Protocol Properties area:
  - **State** – Enabled means that protocol properties are enabled. Disabled means protocol properties are currently disabled.
  - **Priority** – Lists the protocol priority.
  - **Sync'd** – If yes, the protocol priorities have been set by the peer. This parameter cannot be set.
  - **Error** – The error state. This capability indicates whether an error has occurred during the configuration exchange with the peer. Error is also set to yes when the Compatible method for the capability fails.
- FCoE Properties area (FCoE Ports Only):
  - **State** – The FCoE state. It can be enabled or disabled.
  - **Active Priority** – The current active priority assigned for FCoE.
  - **Sync'd** – If yes, the FCoE priority has been set by the peer. This parameter cannot be set.
  - **Error** – The FCoE error state. This capability indicates whether an error has occurred during the configuration exchange with the peer. Error is also set to yes when the Compatible method for the capability fails.
- ETS Priority Group area:
  - **State** – The priority group state. It can be enabled or disabled.
  - **Sync'd** – If yes, the priority groups have been set by the peer. This parameter cannot be set.
  - **Error** – The error state. This capability indicates whether an error has occurred during the configuration exchange with the peer. Error is also set to yes when the Compatible method for the capability fails.
- Active Groups:
  - **PG** – The priority group number. It can be 0 to 7.
  - **Priorities** – The priorities that are assigned to each priority group. It is represented in comma-separated format.
  - **Bandwidth %** – The percentage of available link bandwidth allocated to a particular priority group.

To configure DCB for NIC+FCoE adapter ports, perform these steps:

1. From the **DCB** tab, click **Configure DCB**. The **Configure DCB** dialog appears (Figure 73).

NOTE An error message is displayed if you try to configure more priority groups than the adapter supports. The **Max Configurable PGs** field in the ETS priority groups area shows the number of priority groups supported by the adapter.

Figure 73 Configure DCB Dialog (NIC+FCoE)

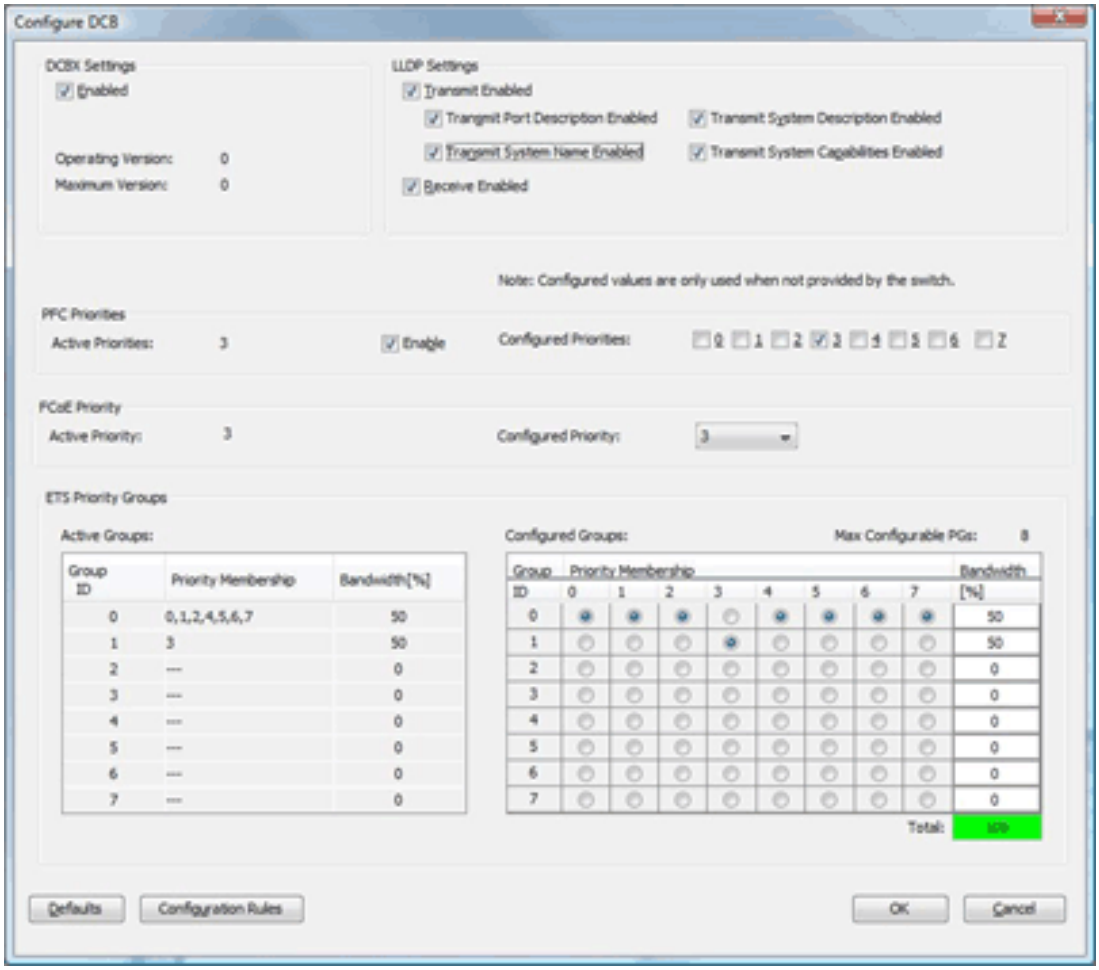

2. Configure the settings.

The following **Configure DCB** dialog fields are displayed:

- DCBX Settings area:
  - **Enabled** – DCBX can be enabled or disabled. With DCBX enabled, the configured values are used only if the switch does not provide them. With DCBX disabled, the configured values are used.
  - **Operating Version** – The operating version of the DCBX protocol. The system adjusts as needed to operate at the highest version supported by both link partners. This setting cannot be changed.
  - **Maximum Version** – The highest DCBX protocol version supported by the system. Version numbers start at zero. The DCBX protocol must be backward compatible with all previous versions. This setting cannot be changed.
- LLDP Settings area:
  - **Transmit Enabled** – LLDP Transmit can be enabled or disabled.
  - **Transmit Port Description Enabled** – Provides a description of the port in an alpha numeric format. If the LAN device supports RFC 2863, the value equals the ifDescr object.
  - **Transmit System Name Enabled** – Provides the system's assigned name in an alpha numeric format. If the LAN device supports RFC 3418, the value equals the sysName object.

  - **Transmit System Description Enabled** – Provides a description of the network entity in an alpha numeric format. This description includes system's name and versions of hardware, operating system, and networking software supported by the device. If the LAN device supports RFC 3418, the value equals the sysDescr object.
  - **Transmit System Capabilities Enabled** – Indicates the primary function or functions of the device and whether or not these functions are enabled on the device. The capabilities are indicated by two octets. Bits 0 through 7 indicate Other, Repeater, Bridge, WLAN AP, Router, Telephone, DOCSIS cable device, and Station, respectively. Bits 8 through 15 are reserved.
  - **Receive Enabled** – LLDP Receive can be enabled or disabled.
- PFC Priorities area:
  - **Active Priorities** – The priorities that are marked active for PFC.
  - **Enable** – If selected, PFC is enabled.
  - **Configured Priorities** – The priorities that are configured, but might not yet be active.
- FCoE Priority area (FCoE Ports Only):
  - **Active Priority** – The active FCoE priority.
  - **Configured Priority** – The configured FCoE priority.
- ETS Priority Groups area:
  - Active Groups:
    - **Group ID** – The priority group ID.
    - **Priority Membership** – The different priorities that are assigned to the various priority groups. This information is the currently active configuration.
    - **Bandwidth** – The bandwidths that are assigned to different priority groups. This information is the currently active configuration.
  - Configured Groups:
    - **Group ID** – The priority group ID.
    - **Priority Membership** – The configured priority membership grouping.
    - **Bandwidth %** – The configured value of bandwidth for the different priority groups.
    - **Max Configurable PGs** – The maximum number of priority groups that can be configured.

3. Click **OK** to apply and save your changes.

To see the configuration rules that apply to configuring priorities, click **Configuration Rules**. You must observe these rules when configuring priority groups for adapter ports. Configuration rules are displayed based on the type of adapter ports you are configuring.

- FCoE configuration rules:
  - One and only one priority is configured for the FCoE priority.
  - You can select a maximum of two PFC priorities, and one of them must match the FCoE priority.
  - The PFC Priority Group to which the FCoE priority is assigned must contain no other priorities.
  - The additional PFC priority must be assigned to a priority group that has no other priorities.
  - Bandwidths of all the priority groups must total 100%.
- NIC configuration rules:
  - One and only one PFC priority can be configured.
  - The PFC priority must be assigned to a priority group that has no other priorities.
  - Bandwidths of all the priority groups must total 100%.

To reset the DCB configuration to its default values, click **Defaults**.

Chapter 11: Exporting SAN Information

11.1 Creating a SAN Report

The OneCommand Manager application enables you to create reports about discovered SAN elements. Reports are generated in `.xml` and `.csv` format and include all the SAN information that is displayed through the various OneCommand Manager application tabs.

NOTE Creating a SAN report can take several minutes for a large SAN.

To create a SAN report, perform these steps:

1. From the **File** menu, select **Export SAN Info**.
2. Browse to a folder and enter a file name with the `.xml` or `.csv` extension.
3. Click **Save** to start the export process.

   During the export process, progress is displayed in the lower-right side of the progress bar. On Windows, you cannot change views, reset, or download firmware during the export process.

Chapter 12: Diagnostics

This section describes the diagnostic tests that can be run on LightPulse adapters.

12.1 FC Diagnostics

This section describes the diagnostics (Figure 74) available for FC adapters. For diagnostics for LPe16202/OCe15100 adapters in NIC+FCoE mode, see Section 12.2, Diagnostics for LPe16202/OCe15100 Adapters in NIC+FCoE Mode.

Use the **Diagnostics** tab to perform the following tasks:

- View flash load list, PCI registers, and wakeup parameter information.
- Run the following tests on Emulex adapters installed in the system:
  - PCI Loopback
  - Internal Loopback
  - External Loopback
  - Power-on self-test (POST)
  - Echo (End-to-End)
  - Quick Test

  These tests are not available in read-only mode.
- Run a diagnostic dump and retrieve dump files from remote hosts (this option is not available in read-only mode).
- Control adapter beaconing (this option is not available in read-only mode).

CAUTION Running a PCI Loopback, Internal Loopback, External Loopback, or POST test on an LPe12000-series adapter (for example, LPe12000) that is being used to boot from SAN is not advisable. After the tests have completed, the system performs an adapter reset, which may cause a loss of connectivity to the SAN and possible loss of data. To perform these tests on an LPe12000-series adapter, you must make sure that the adapter is not currently being used to boot from SAN. Perform one of the following action:

- Move the adapter to be tested to a non-boot from SAN host, and perform the tests from that location.
- If the host with the adapter that needs to be tested is also hosting other boot from SAN adapters, perform a boot from SAN using one of the other boot from SAN adapters. The target adapter can now be tested, since it is no longer being used for boot from SAN.

12.1.1 Viewing Flash Contents, PCI Registers, and Wakeup Information

The **Diagnostics** tab (Figure 74) shows PCI register dump information and flash memory contents. The information is read-only and is outlined in the following section.

Figure 74 PCI Registers and Flash Contents of the Diagnostics Tab

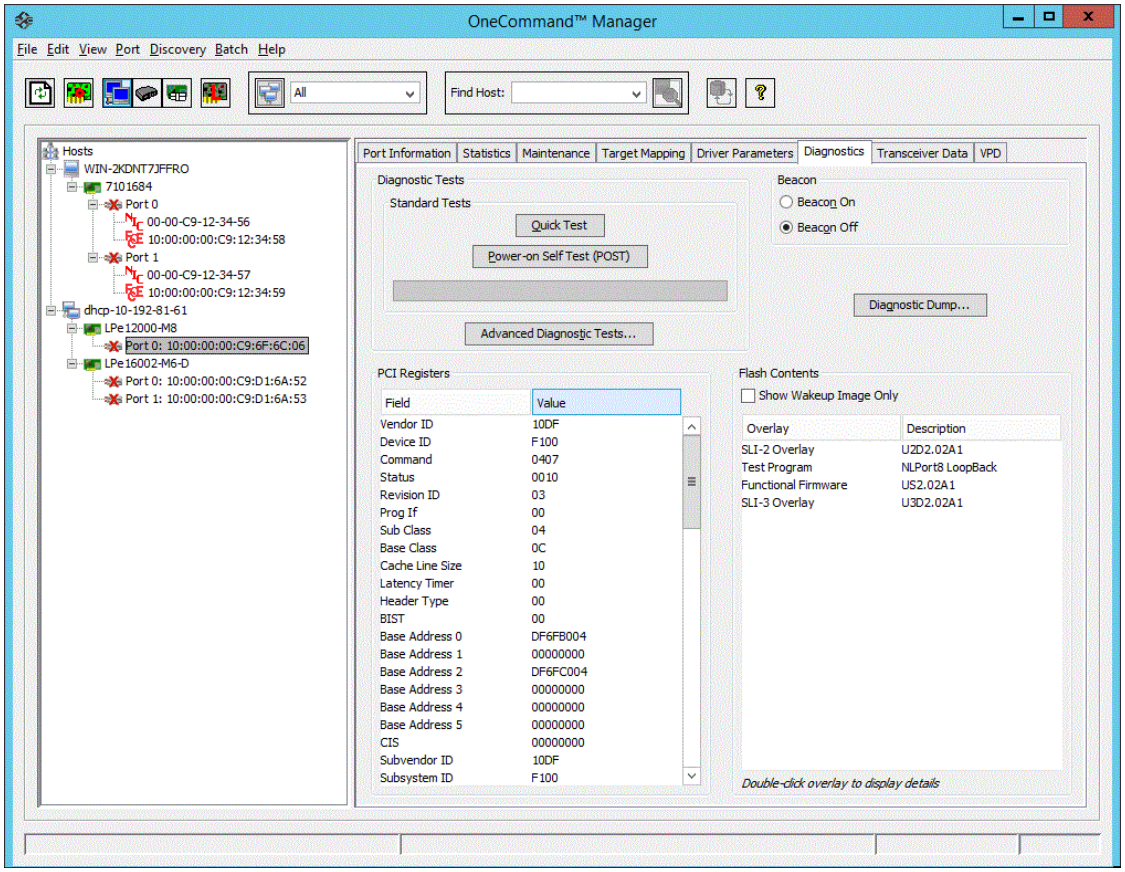

12.1.1.1 Viewing Flash Contents

If you select the **Show Wakeup Image Only** check box, the flash overlays that are not loaded when the system is booted no longer display. This check box defaults to not selected.

12.1.1.2 Viewing Overlay Details

If you double-click a flash overlay, another window appears with details about that overlay (Figure 75).

Figure 75 Overlay Detail Window

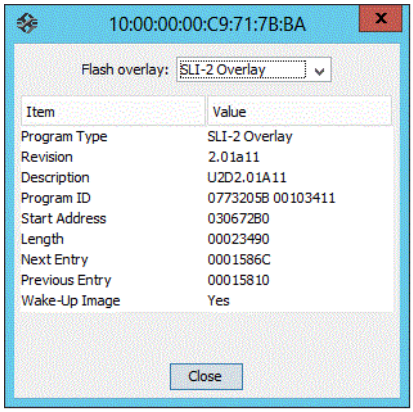

To see the details of a different flash overlay image, you can either close the details window and double-click on another overlay name, or choose a different overlay name from the **Flash overlay** menu.

12.1.1.3 Viewing the PCI Registers

The PCI Registers appear directly on the **Diagnostics** tab (Figure 74).

12.1.2 Running a Quick Test

The **Diagnostics** tab enables you to run a quick diagnostics test on a selected FC port. The Quick Test consists of 50 PCI Loopback test cycles and 50 Internal Loopback test cycles. This test is not available in read-only mode or on LightPulse adapters in ESXi hosts.

To use quick test, perform these steps:

1. From the discovery-tree (Figure 4), select the FC port on which you want to run the Quick Test.
2. Select the **Diagnostics** tab (Figure 74) and click **Quick Test**. A warning message appears (Figure 76).

Figure 76 Run Quick Test Warning

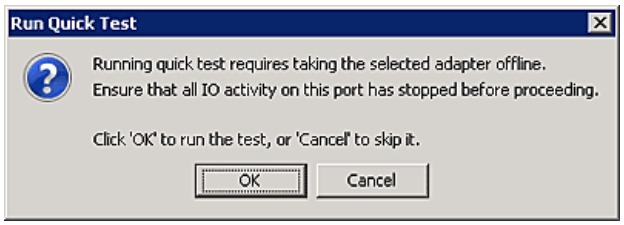

3. Click **OK** to run the test. The **Quick Diagnostic Test** window appears displaying the PCI Loopback and Internal Loopback test results.

12.1.3 Running a POST

NOTE The POST is supported only on LPe12000-series adapters.

The POST is a firmware test normally performed on an adapter after a reset or restart. The POST does not require any configuration to run. This test is not available in read-only mode.

To run the POST, perform these steps:

1. From the discovery-tree (Figure 4), select the FC port on which you want to run the POST.
2. Select the **Diagnostics** tab (Figure 74) and click **Power-on Self Test (POST)**. A warning dialog appears.
3. Click **OK**. A **POST** window appears displaying POST information.

NOTE After the test starts, it cannot be cancelled. It must run to completion.

12.1.4 Using Beaconing

The beaconing capability enables you to force a specific adapter's LEDs to blink in a particular sequence. The blinking pattern acts as a beacon, making it easier to locate a specific adapter among racks of other adapters. This option is not available in read-only mode.

If you enable beaconing, the two LEDs blink rapidly in unison for 24 seconds, after which the LEDs report the adapter health status for 8 seconds. When the 8 seconds are up, the adapter returns to Beaconing mode. This cycle repeats indefinitely until you disable beaconing or you reset the adapter.

NOTE The beaconing buttons are disabled if the selected adapter does not support beaconing.

To enable or disable beaconing, perform these steps:

1. From the discovery-tree (Figure 4), select the FC port whose LEDs you want to set.
2. Select the **Diagnostics** tab (Figure 74) and click **Beacon On** or **Beacon Off**.

12.1.5 Running D_Port Tests

D_Port is a diagnostic mode supported by Brocade switches for 16GFC and faster. D_Port tests enable you to detect physical cabling issues that result in increased error rates and intermittent behavior. If activated, D_Port runs a series of tests including local electrical loopback, loopback to the remote optics, loopback from the remote port to the local optics, and a full device loopback test with data integrity checks. It also provides an estimate of cable length to validate that a proper buffering scheme is in place. The various loopback tests allow some level of fault isolation so you can distinguish faults from marginal cable, optics modules, and connector or optics seating.

12.1.5.1 D_Port Test Considerations

- Basic connectivity diagnostics are already supported by Emulex HBAs. The OneCommand Manager application has diagnostic modes that support validation of connection to the switch. The functionality that Brocade offers provides the ability to diagnose marginal cable conditions (for example, dust in the optics) that result in higher error rates.
- Do not enable D_Port on the switch port.
- D_Port tests run with the physical connection in an offline diagnostic state, so normal I/O cannot be sent through the physical port while the test is in progress. While the port is in D_Port mode, the link appears down on that port, similar to an unplugged cable.
- When using D_Port in a boot from SAN configuration, the configuration must have redundant paths to the boot LUN and only one of the redundant adapter ports should be set to D_Port.
- For more information about D_Port, refer to the Brocade website at www.brocade.com.
- D_Port is also referred to as ClearLink.

The **D_Port Tests** button on the **Diagnostics** tab enables you to run D_Ports tests on LPe16000, LPe31000, and LPe32000-series adapters (Figure 77).

Figure 77 Diagnostics Tab for LPe16000-Series Adapters (D_Port Tests... Button Depicted)

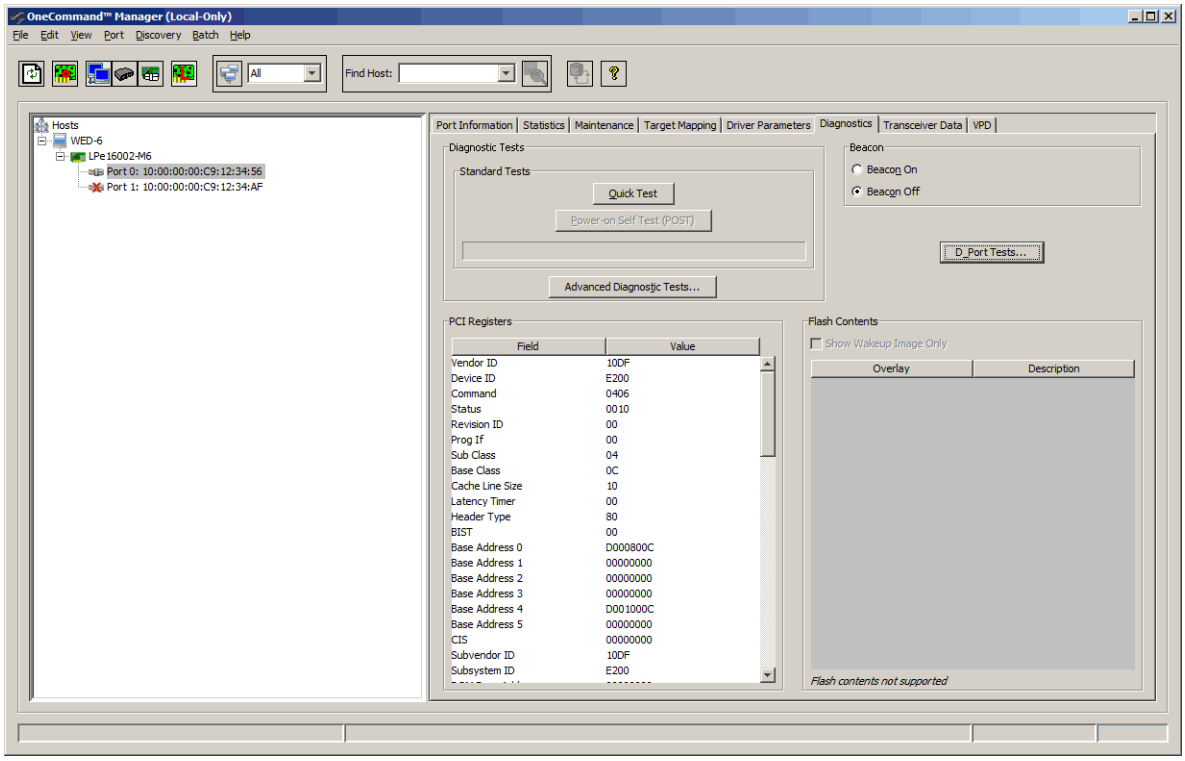

To run a D_Port test, perform these steps:

1. From the discovery-tree, select the FC port on which you want to run the D_Port test.
2. Select the **Diagnostics** tab (Figure 74) and click **D_Port Tests**. The **D_Port Tests** window appears (Figure 78).
3. Click **Start Tests**. The start time is displayed.

Figure 78 D_Port Tests Dialog

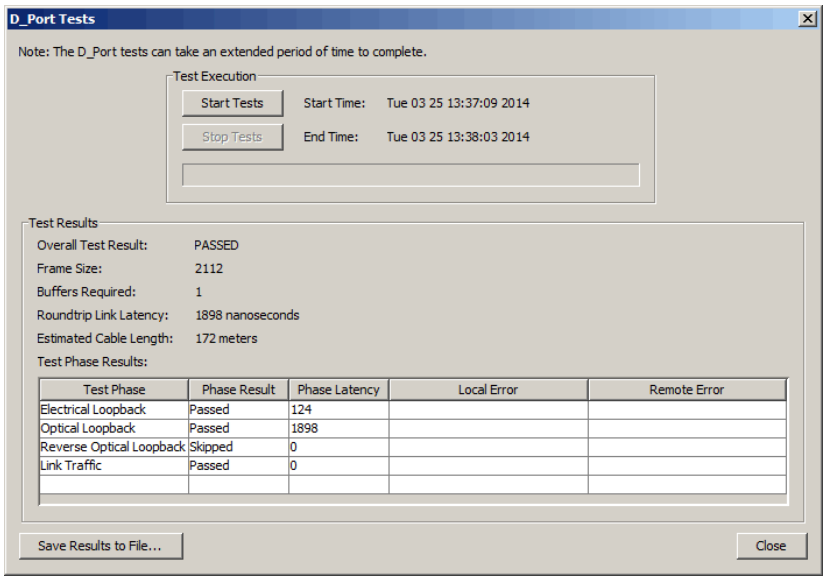

The following **D_Port Tests** dialog fields are displayed:

- Test Results area:
  - **Overall Test Result** – Displays PASSED or FAILED depending upon the outcome of all the test phases.
  - **Frame Size** – The size of the frames used in each test phase.
  - **Frame Count** – The number of frames generated during each test phase.
  - **Roundtrip Link Latency** – Estimated roundtrip link latency calculated by switch during the execution of all tests.
  - **Estimated Cable Length** – Estimated cable length calculated by switch during the execution of all tests.
- Test Phase Results area:
  - **Test Phase** – The name of the test run.
  - **Phase Result** – The result of the test run. Possible results are Pass, Fail, or Skipped.
  - **Phase Latency** – The round trip legacy (in ns.) calculated during the execution of the test.
  - **Local Error** – The errors, if any, detected on the local side of the test.
  - **Remote Error** – The errors, if any, detected on the remote side of the test.

To stop running D_Port tests, click **Stop Tests**. The stop time is displayed.

To save the test results to a file, click **Save Results to File**. You are prompted to enter the file name of the save file.

NOTE If the SFP or adapter firmware do not support running D_Port diagnostics, clicking **Start Tests** causes an error message to be displayed, and the tests are not run.

12.1.6 Creating Diagnostic Dumps

NOTE This option is not available in read-only mode.

The diagnostic dump capability enables you to create a dump file for a selected FC port. Dump files contain various information, such as firmware version, driver version, and so on, that is particularly useful when troubleshooting an adapter. You can also retrieve dump files from remote hosts.

For LPe12000-series HBAs, see Section 12.1.6.1, Creating Diagnostic Dumps for LPe12000-Series Adapters. For LPe16000-series, LPe31000-series, and LPe32000-series adapters, see Section 12.1.6.2, Creating Diagnostic Dumps for All Other FC Adapters.

CAUTION Disruption of service can occur if a diagnostic dump is run during I/O activity.

12.1.6.1 Creating Diagnostic Dumps for LPe12000-Series Adapters

To start a diagnostic dump, perform these steps:

1. From the discovery-tree (Figure 4), select an FC port whose diagnostic information you want to dump.
2. Select the **Diagnostics** tab (Figure 74) and click **Diagnostic Dump**. The **Diagnostic Dump** dialog appears (Figure 79). You can specify how many files you want to retain using the Dumps Retained counter. Click **Delete Existing Dump Files** to remove existing dump files for the selected FC port from your system.

Figure 79 Diagnostic Dump Dialog

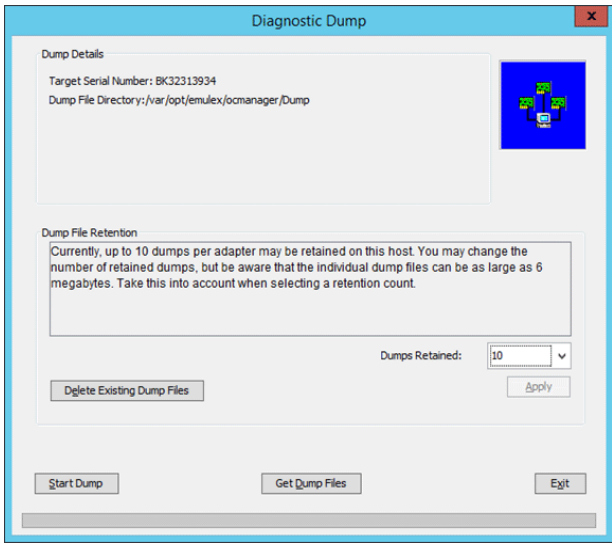

3. Click **Start Dump**. A warning message about taking the adapter offline appears.

   NOTE For VMware systems, you must set a dump directory before initiating a dump. The dump directory must be a storage partition (a datastore) under the directory `/vmfs/volumes`.

4. Click **OK**. Dump files are created. The file location depends upon your operating system:
   - Windows – In the `Dump` directory under the OneCommand Manager Installation Directory `Util\Dump\`.
   - Solaris – `/opt/ELXocm/Dump`.
   - Linux – `/var/log/emulex/ocmanager/Dump`.
   - VMware – a dump directory you created under `/vmfs/volumes`.

   Two files are created:
   - *<Hostname_WWPN_Date-Time>*`.dmp`
   - *<Hostname_WWPN_Date-Time>*`.txt`

5. To list the dump files in the local system or to obtain remote host dump files and copy them to your local system, click **Get Dump Files**. The **Diagnostic Dump File Transfer** dialog appears (Figure 80).

   NOTE The **Start Copy** button is disabled when a local adapter port is selected.

Figure 80 Diagnostic Dump File Transfer Dialog

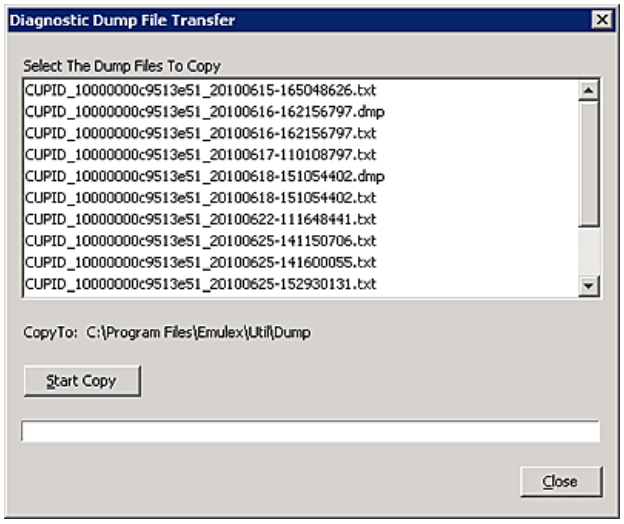

6. Select the files you want to copy (multiple selections are available), and click **Start Copy**. The remote dump files are copied to your local dump folder. The local dump folder locations are described in step 4.

12.1.6.2 Creating Diagnostic Dumps for All Other FC Adapters

NOTE This option is not available in read-only mode.

The diagnostic dump capability enables you to create an Enhanced FAT Dump (EFD) dump file for a selected adapter. Dump files contain various information, such as firmware version, driver version, and so on, that is particularly useful when troubleshooting an adapter. You can also retrieve dump files from remote hosts.

To start a diagnostic dump, perform these steps:

1. From the discovery-tree (Figure 4), select an adapter whose diagnostic information you want to dump.
2. Select the **Firmware** tab (Figure 24), and click **Diagnostic Dump**. The **Diagnostic Dump** dialog appears (Figure 93).

   For hosts being managed through the CIM interface, the **Set Dump Directory** button enables you to set the dump directory for ESXi host dumps (VMware only).
3. Specify how many files you want to retain using the Files Retained counter. Click **Delete Existing Dump Files** to remove existing dump files for the selected adapter from your system.

Figure 81 Diagnostic Dump Dialog

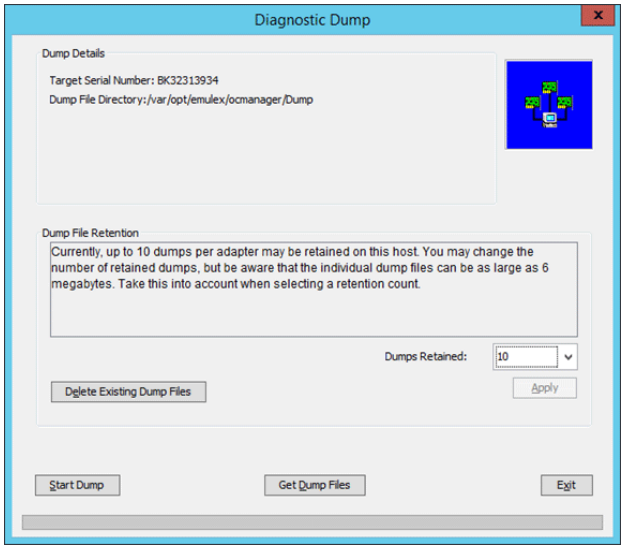

4. Click **Start Dump**. Dump files are created. The file location depends upon your operating system:

   NOTE For VMware systems, you must set a dump directory before initiating a dump. The dump directory must be a storage partition (a datastore) under the directory `/vmfs/volumes`.

   — Windows – `%ProgramFiles%Util\Dump\`.
   — Solaris – `/opt/ELXocm/Dump`.
   — Linux – `/var/log/emulex/ocmanager/Dump`.
   — VMware – a dump directory you create under `/vmfs/volumes`.

   Two files are created:

   - *<Hostname_WWPN_Date-Time>*.efd
   - *<Hostname_WWPN_Date-Time>*.txt

5. To obtain remote host dump files and copy them to your local system, click **Get Dump Files**. The **Diagnostic Dump File Transfer** dialog appears (Figure 94).

   NOTE The **Get Dump Files** button is disabled if a local adapter port is selected.

Figure 82 Diagnostic Dump File Transfer Dialog

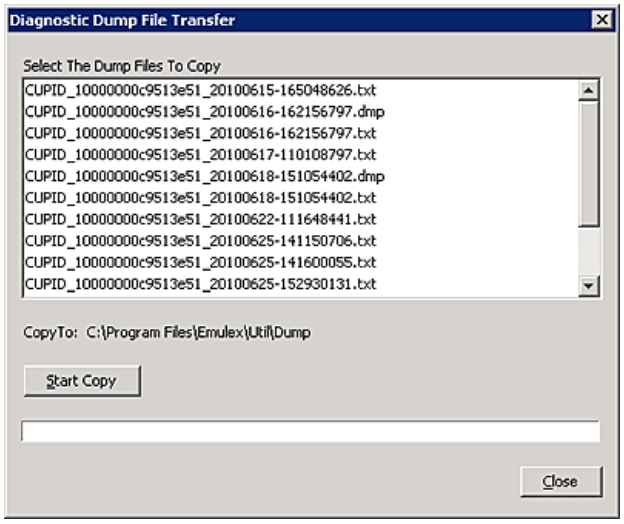

6. Select the files you want to copy (multiple selections are available), and click **Start Copy**. The remote dump files are copied to your local dump folder. The local dump folder locations are described in step 4.

12.1.7 Running Advanced Diagnostic Tests

The advanced diagnostics capability gives you greater control than the quick test over the type of diagnostics tests that run. Through advanced diagnostics, you can specify which tests to run, the number of cycles to run, and actions to take in the event of a test failure. Advanced diagnostics capability is not available in read-only mode.

To run advanced diagnostics tests, click **Advanced Diagnostic Tests** on the **Diagnostics** tab (Figure 74). The **Diagnostic Test Setup** dialog appears (Figure 83).

Figure 83 Diagnostic Test Setup Dialog

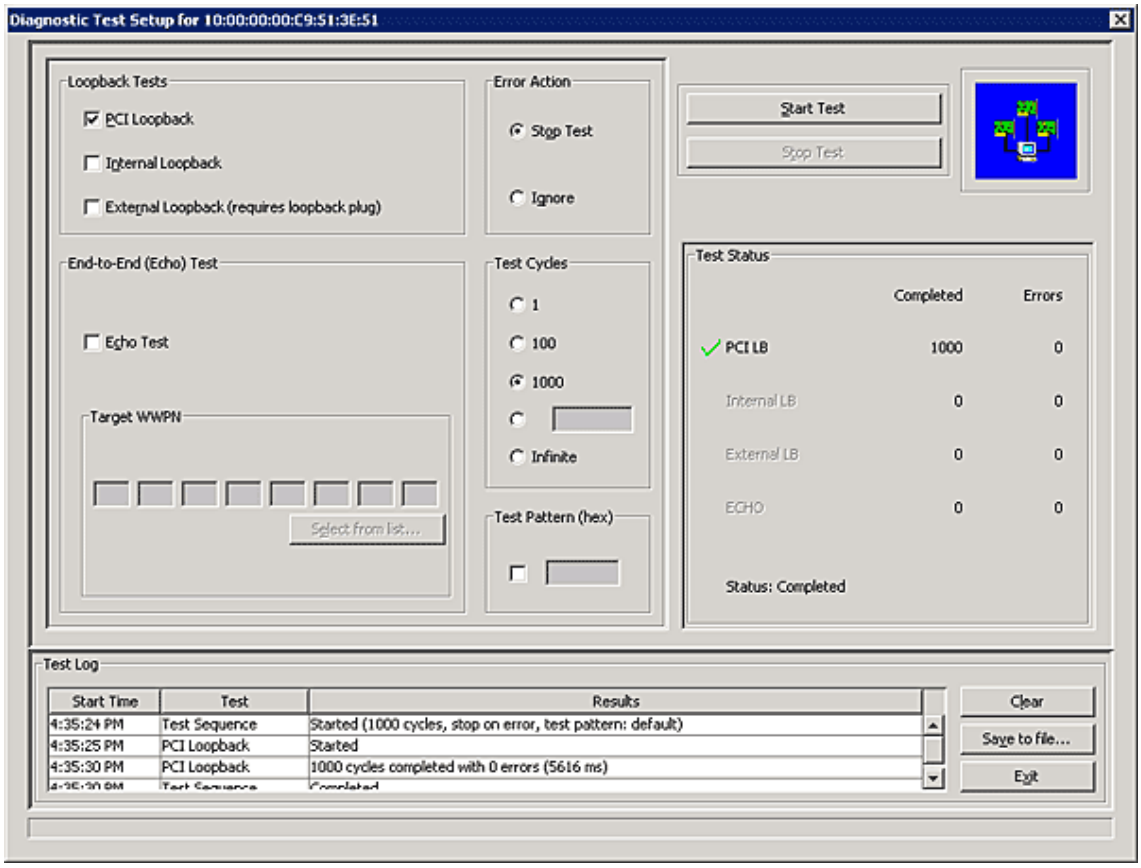

The following **Diagnostic Test Setup** dialog fields are displayed:

- Loopback Tests area:
  - **PCI Loopback**
  - **Internal Loopback**
  - **External Loopback**

  NOTE For details about these tests, see Section 12.1.7.1, Running Loopback Tests.
- End-to-End (Echo) Test
  - **Echo Test**
  - **Target WWPN**

  NOTES
  - For details about this test, see Section 12.1.7.2, Running End-to-End (Echo) Tests.
  - You cannot run the External Loopback test and the Echo test concurrently. If you select the **External Loopback** check box, the Echo test section is disabled, and if you select the **Echo Test** check box, the External Loopback section is disabled.
- Error Action area:

  Error Action enables you to define the actions to be performed in the event of a test failure. Two error action options exist:
  - **Stop Test** – Do not log the error and halt the test. No further tests are run.

- **Ignore** – Log the error and proceed with the next test cycle.

- Test Cycles area:

  Test Cycles enables you to specify test cycles three ways:

  - Select an established cycle count by clicking on the corresponding radio button.
  - Enter a custom cycle count in the blank field in the Test Cycles area.
  - Set the test to run until you manually click **Stop Test**, by selecting the **Infinite** radio button.

- Test Pattern area:

  Enter a custom test pattern to be used in tests that transfer data. The test pattern can be up to eight hexadecimal bytes.

- Test Status area:

  The Test Status area displays how many completed cycles of each test ran, in addition to the number of errors for each test.

- Test Log area:

NOTE For details about test logs, see Section 12.1.7.3, Saving the Log File.

12.1.7.1 Running Loopback Tests

You can run the following loopback tests:

- PCI Loopback Test – A firmware controlled diagnostic test in which a random data pattern is routed through the PCI Bus without being sent to an adapter link port. The returned data is subsequently validated for integrity.
- Internal Loopback Test – A diagnostic test in which a random data pattern is sent down to an adapter link port, then is immediately returned without actually going out on the port. The returned data is subsequently validated for integrity.
- External Loopback Test – A diagnostic test in which a random data pattern is sent down to an adapter link port. The data goes out the port and immediately returns by way of a loopback connector. The returned data is subsequently validated for integrity.

NOTE You cannot run the External Loopback test and Echo test concurrently. If you select **External Loopback**, the Echo Test section is disabled, and if you select **Echo Test**, the External Loopback section is disabled.

To run loopback tests, perform these steps:

1. From the **Diagnostics** tab (Figure 74), click **Advanced Diagnostics Tests** (Figure 83). From the **Loopback Test** section of the dialog, choose the type of Loopback test you want to run and define the Loopback test parameters.

NOTE You must insert a loopback plug in the selected adapter before running an External Loopback test.

2. Click **Start**. The following warning appears (Figure 94).

Figure 84 Run Diagnostic Tests Warning

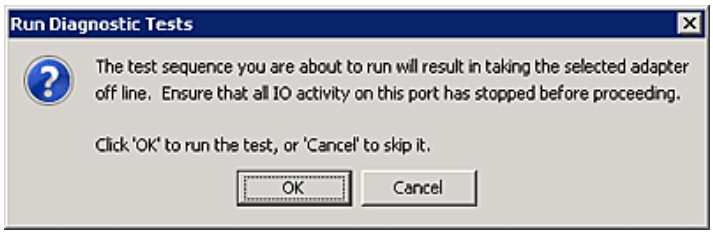

3. Click **OK**. If you choose to run an External Loopback test, the following window appears (Figure 85).

Figure 85 Advanced Diagnostic Tests Warning Window for External Loopback

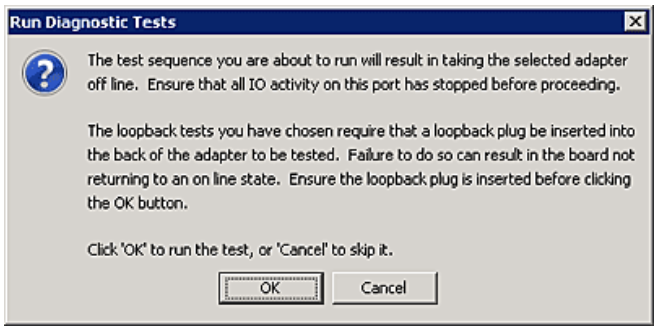

4. Click **OK**. The progress bar indicates that the test is running.

   Periodic test feedback, consisting of the current loopback test/cycle plus the completion status of each type of test, is displayed in the **Test Log** section of the dialog. Either click **Clear** to erase the contents of the log display, or click **Save to File** to save the log file.

After starting the tests, you can click **Stop Tests** to stop the tests before they complete. Depending upon the tests being run, it may take some time before they stop.

12.1.7.2 Running End-to-End (Echo) Tests

The end-to-end test enables you send an `echo` command and response sequence between an adapter port and a target port.

NOTES

- Not all remote devices respond to an `echo` command.
- You cannot run the Echo test and the External Loopback test concurrently. If you select **Echo Test**, the External Loopback test is disabled.

To run end-to-end echo tests, perform these steps:

5. From the **Diagnostics** tab (Figure 74), click **Advanced Diagnostic Tests** (Figure 83).
6. Check **Echo Test**, and enter the WWPN for the target.

   Click **Select From List** if you do not know the actual WWPN of the test target. The **Select Echo Test Target** dialog appears (Figure 86).

   Select the port to test from the tree-view and click **Select**. All relevant information for the selected port is automatically added to the Target Identifier section of the **Diagnostics** dialog.

Figure 86 Select Echo Test Target Window

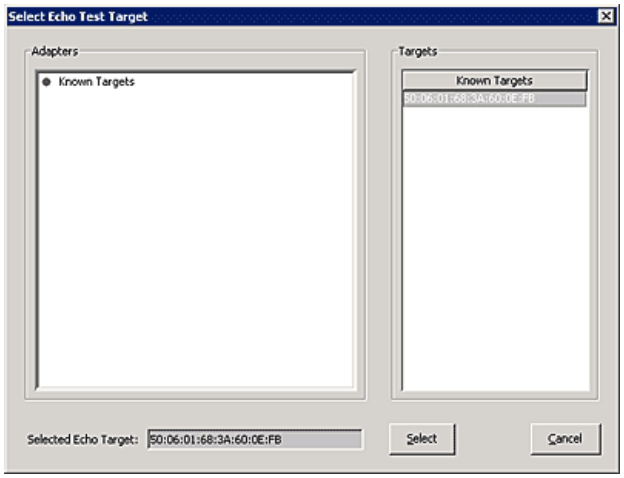

7. Define the other parameters you want to use and click **Start Test**. The following warning window appears (Figure 87).

Figure 87 Advanced Diagnostic Tests Warning Window

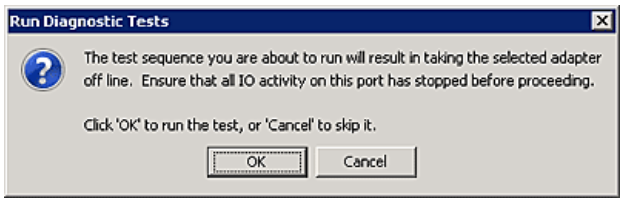

8. Click **OK**. A result screen appears, and the test results appear in the Test Log. Either click **Clear** to erase the contents of the log display, or click **Save to File** to save the log file.

12.1.7.3 Saving the Log File

You can save the test log to a log file for later viewing or printing. When new data is written to a saved file, the data is appended to the end of the file. Each entry has a two-line header that contains the identifier of the adapter being tested and the date and time of the test. Over time, the data accumulates to form a chronological history of the diagnostics performed on the adapter.

The default location is:

- In Windows: The OneCommand Manager application install directory on your local drive.
- In Solaris: `/opt/ELXocm/Dump`
- In Linux: `/var/opt/emulex/ocmanager/Dump`
- In VMware Server: A default directory does not exist for VMware

After writing an entry into the log, you are prompted to clear the display. The default name of the saved file is `DiagTest.log`. An example of a saved log file is shown in Figure 88.

Figure 88 Example of a DiagTest.log Window

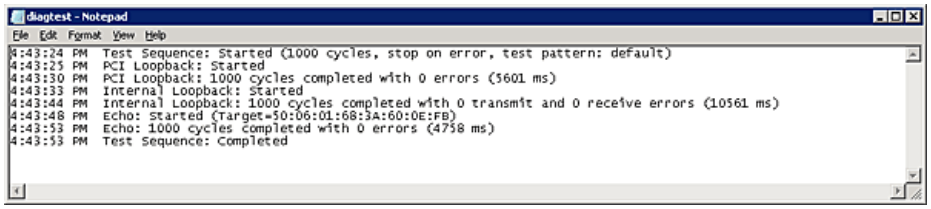

To save the log file, perform these steps:

1. After running a test from the **Diagnostic Test Setup** dialog (Figure 83), click **Save to File**. The **Select Diagnostic Log file Name** dialog appears. The default name of a saved file is `DiagTest.log`.
2. Browse to the desired directory, change the log file name if you want and click **Save** (Figure 83).

12.2 Diagnostics for LPe16202/OCe15100 Adapters in NIC+FCoE Mode

This section describes the diagnostics for LPe16202/OCe15100 adapters in NIC+FCoE mode. For FC adapter diagnostics, see Section 12.1, FC Diagnostics.

NOTE Diagnostics are not available in read-only mode. See Section 4.3, Changing Management and Read-Only Mode, for more information.

To run diagnostics, select the port on which you want to run diagnostics and select the **Diagnostics** tab. The **Diagnostics** tab appears (Figure 89).

Figure 89 Diagnostics Tab

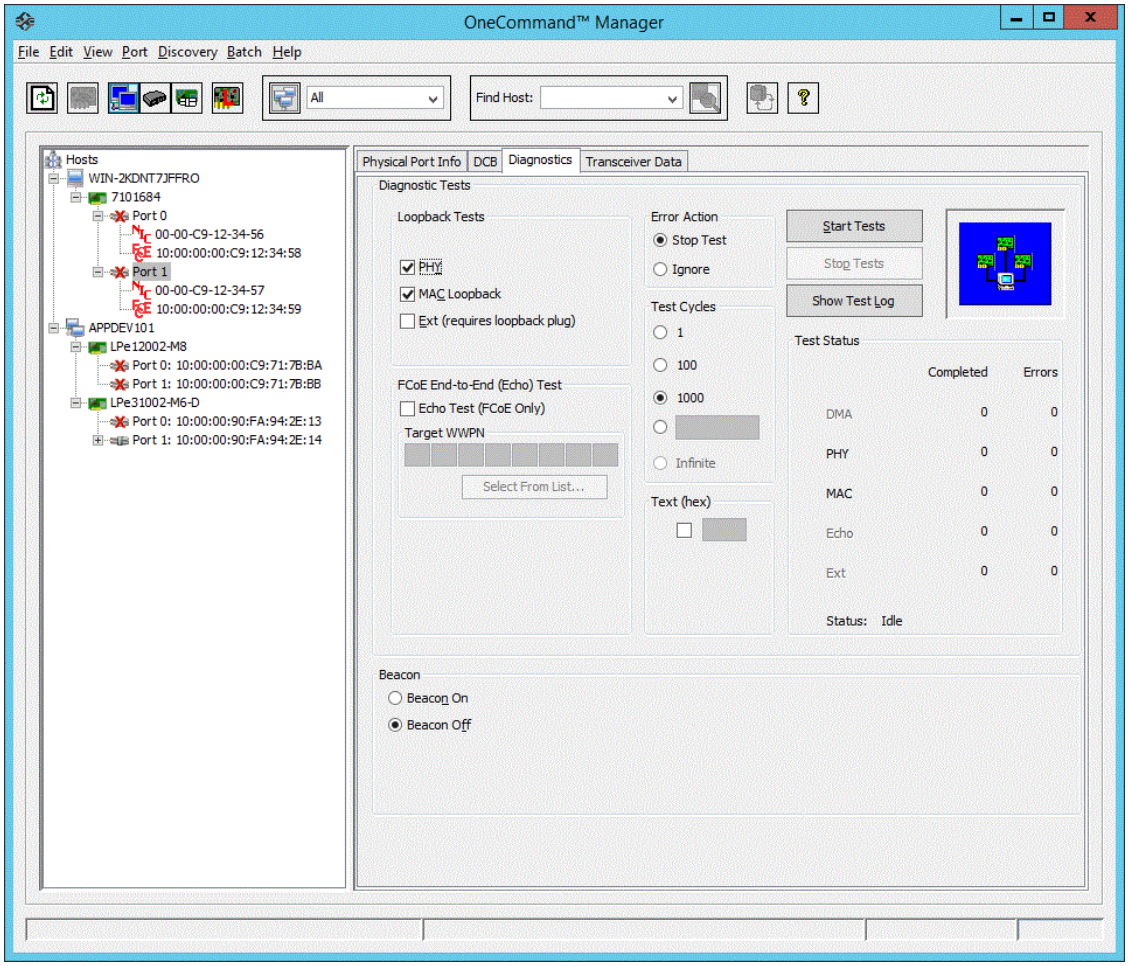

The following **Diagnostic Test Setup** dialog fields are displayed:

- Loopback Tests area:
  - **PHY**
  - **MAC Loopback**
  - **Ext**

NOTE For details about these tests, see Section 12.2.1, Running Loopback Tests.

- FCoE End-to-End (Echo) Test
  - **Echo Test (FCoE Only)**
  - **Target WWPN**

NOTES

- For details about this test, see Section 12.2.2, Running FCoE End-to-End (Echo) Tests.
- You cannot run the External Loopback test and the Echo test concurrently. If you select **External Loopback**, the Echo test section is disabled, and if you select **Echo Test (FCoE Only)**, the External Loopback section is disabled.

- Error Action area:

  Error Action enables you to define the actions to be performed in the event of a test failure. Two error action options exist:

  — **Stop Test** – Do not log the error and halt the test. No further tests are run.

  — **Ignore** – Log the error and proceed with the next test cycle.
- Test Cycles area:

  Test Cycles enables you to specify test cycles three ways:

  — Select an established cycle count by clicking on the corresponding radio button.

  — Enter a custom cycle count in the blank field in the Test Cycles area.

  — Set the test to run until you manually click **Stop Test**, by selecting the **Infinite** radio button.
- Test area:

  Enter a custom test pattern to be used in tests that transfer data. The test pattern can be up to 8 hexadecimal bytes.
- Test Status area:

  The Test Status area displays how many completed cycles of each test ran, in addition to the number of errors for each test.

12.2.1 Running Loopback Tests

You can run the following loopback tests:

- PHY Loopback test – The PHY Loopback test connects the transmit output of the physical layer to the receive input of the physical layer. The data is transmitted, received, and checked for data miscompute errors.
- MAC Loopback test – The MAC Loopback test connects the transmit output of the MAC controller to the receive input of the MAC controller (bypassing the PHY).
- External Loopback test– A diagnostic test in which a random data pattern is sent down to an adapter link port. The data goes out the port and immediately returns by way of a loopback connector. The returned data is subsequently validated for integrity.

NOTE You cannot run the External Loopback test and Echo test concurrently. If you select **External Loopback**, the FoCE End-to-End (Echo) Test section is disabled, and if you select **Echo Test (FCoE Only)**, the External Loopback section is disabled.

To run loopback tests, perform these steps:

1. From the Loopback Test section of the **Diagnostics** tab, choose the type of loopback test you want to run and define the loopback test parameters.

NOTE You must insert a loopback plug in the selected adapter before running an External Loopback test. Also, you must ensure that the NIC function of the port goes to a link-up state. To verify the status, follow the instructions in Section 12.2.5, Creating Diagnostic Dumps.

2. Click **Start**. The following warning appears (Figure 90).

Figure 90 Run Diagnostic Tests Warning

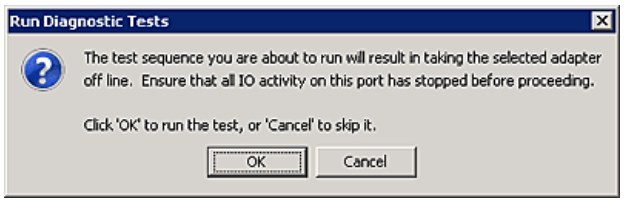

3. Click **OK**. If you choose to run an External Loopback test, the following window appears (Figure 91).

Figure 91 Advanced Diagnostic Tests Warning Window for External Loopback

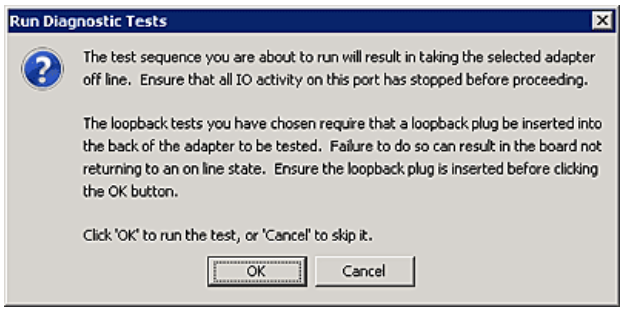

4. Click **OK**. The progress bar indicates that the test is running.

   Periodic test feedback, consisting of the current loopback test/cycle plus the completion status of each type of test, is displayed in the **Test Status** section of the dialog. Click **Show Test Log** to view and save the log file. After starting the tests, you can click **Stop Tests** to stop the tests before they complete. Depending upon the tests being run, it may take some time before they stop.

12.2.2 Running FCoE End-to-End (Echo) Tests

The FoCE End-to-End test enables you send an Echo command and response sequence between an adapter port and a target port.

NOTES

- Not all remote devices respond to an echo command.
- You cannot run the Echo test and the External Loopback test concurrently. If you select the ECHO Test, the External Loopback test is disabled.

To run end-to-end echo tests, perform these steps:

1. From the FCoE End-to-End (Echo) Test area, check **Echo Test (FCoE Only)**.
2. Enter the WWPN for the target. If you do not know the WWPN of the test target, click **Select From List**. The **Select Echo Test Target** dialog appears (Figure 86). Select the port to test from the tree-view and click **Select**. All relevant information for the selected port is automatically added to the Target Identifier section of the **Diagnostics** dialog.
3. Define the other parameters you want to use and click **Start Test**. The warning shown in Figure 87 appears.
4. Click **OK**. A result screen appears, and the test results appear in the Test Log. Click **Clear** to erase the contents of the log display, or click **Save to File** to save the log file.

12.2.3 Using Beaconing

NOTE This option is not available in read-only mode.

The beaconing capability enables you to force a specific adapter's LEDs to blink in a particular sequence. The blinking pattern acts as a beacon, making it easier to locate a specific adapter among racks of other adapters.

When you enable beaconing, the two LEDs blink rapidly in unison for 24 seconds, after which the LEDs report the adapter health status for 8 seconds. When the 8 seconds are up, the adapter returns to Beaconing mode. This cycle repeats indefinitely until you disable beaconing or you reset the adapter.

NOTE The beaconing buttons are disabled if the selected adapter does not support beaconing.

To enable or disable beaconing, perform these steps:

1. From the discovery-tree (Figure 4), select the adapter port whose LEDs you want to set.
2. Select the **Diagnostics** tab (Figure 89) and click either **Beacon On** or **Beacon Off**.

12.2.4 Saving the Log File

NOTE This option is not available in read-only mode.

You can save the test log to a log file for later viewing or printing. When new data is written to a saved file, the data is appended to the end of the file. Each entry has a two-line header that contains the identifier of the adapter being tested and the date and time of the test. Over time, the data accumulates to form a chronological history of the diagnostics performed on the adapter.

The default location is:

- In Windows: The OneCommand Manager application installation directory on your local drive.
- In Solaris: `/opt/ocmanager/Dump`.
- In Linux: `/var/opt/emulex/ocmanager/logs`.
- In VMware Server: A default directory does not exist for VMware.

After writing an entry into the log, you are prompted to clear the display. The default name of the saved file is `diagtest.log`. An example of a saved log file is shown in Figure 92.

Figure 92 Example of a DiagTest.log Window

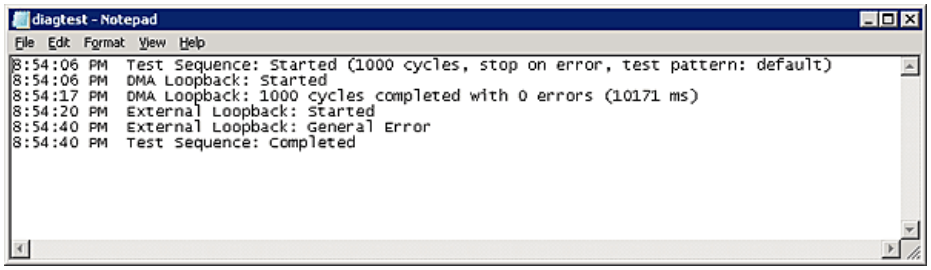

To save the log file, perform these steps:

1. After running a test from the **Diagnostic** tab (Figure 89), click **Save Test Log**. The **Diagnostic Test Log** dialog appears (Figure 92). The default name of a saved file is `DiagTest.log`.
2. Browse to the desired directory, change the log file name if you want, and click **Save to file**.

12.2.5 Creating Diagnostic Dumps

NOTE This option is not available in read-only mode.

The diagnostic dump capability enables you to create an EFD dump file for a selected adapter. Dump files contain various information, such as firmware version and driver version, that is particularly useful when troubleshooting an adapter. You can also retrieve dump files from remote hosts.

To start a diagnostic dump, perform these steps:

1. From the discovery-tree (Figure 4), select an adapter whose diagnostic information you want to dump.
2. Select the **Firmware** tab (Figure 24), and click **Diagnostic Dump**. The **Diagnostic Dump** dialog appears (Figure 93).

   For hosts being managed through the CIM interface, the **Set Dump Directory** button enables you to set the dump directory for ESXi host dumps (VMware only).
3. Specify how many files you want to retain using the Dumps Retained counter. Click **Delete Existing Dump Files** to remove existing dump files for the selected adapter from your system.

Figure 93 Diagnostic Dump Dialog

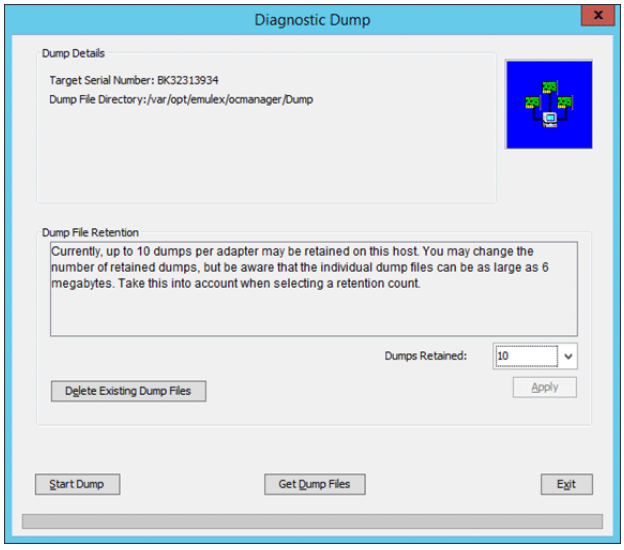

4. Click **Start Dump**. Dump files are created. Where these files are created depends upon your operating system:

   NOTE For VMware systems, you must set a dump directory before initiating a dump. The dump directory must be a storage partition (a datastore) under the directory `/vmfs/volumes`.

   — Windows – `%ProgramFiles%Util\Dump\`.
   — Solaris – `/opt/ELXocm/Dump`.
   — Linux – `/var/log/emulex/ocmanager/Dump`.
   — VMware – A dump directory you create under `/vmfs/volumes`.

   Two files are created:

   - <*Hostname_WWPN_Date-Time*>.efd
   - <*Hostname_WWPN_Date-Time*>.txt
5. To obtain remote host dump files and copy them to your local system, click **Get Dump Files**. The Diagnostic **Dump File Transfer** dialog appears (Figure 94).

NOTE The **Get Dump Files** button is disabled if a local adapter port is selected.

Figure 94 Diagnostic Dump File Transfer Dialog

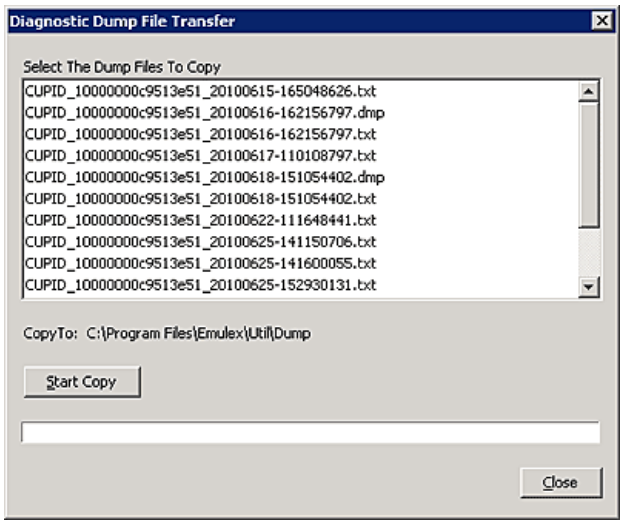

6. Select the files you want to copy (multiple selections are available) and click **Start Copy**. The remote dump files are copied to your local dump folder. The local dump folder locations are described in step 4.

Appendix A: License Notices

A.1 Secure Hash Algorithm (SHA-1) Notice

A.2 OpenPegasus Licensing Notice

persons to whom the Software is furnished to do so, subject to the following conditions:
The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.
THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

A.3 OpenSSL Notice

This is a copy of the current LICENSE file inside the CVS repository.
LICENSE ISSUES
==========
The OpenSSL toolkit stays under a dual license, i.e. both the conditions of the OpenSSL License and the original SSLeay license apply to the toolkit. below for the actual license texts. Actually both licenses are BSD-style Open Source licenses. In case of any license issues related to OpenSSL please contact openssl-core@openssl.org.
OpenSSL License

/* ==
* Copyright (c) 1998-2011 The OpenSSL Project. All rights reserved.
*
* Redistribution and use in source and binary forms, with or without
* modification, are permitted provided that the following conditions
* are met:
*
* 1. Redistributions of source code must retain the above copyright
* notice, this list of conditions and the following disclaimer.
*
* 2. Redistributions in binary form must reproduce the above copyright
* notice, this list of conditions and the following disclaimer in
* the documentation and/or other materials provided with the
* distribution.
*
* 3. All advertising materials mentioning features or use of this
* software must display the following acknowledgment:
* "This product includes software developed by the OpenSSL Project
* for use in the OpenSSL Toolkit. (http://www.openssl.org/)"
*
* 4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to
* endorse or promote products derived from this software without
* prior written permission. For written permission, please contact
* openssl-core@openssl.org.

```
*
* 5. Products derived from this software may not be called "OpenSSL"
* nor may "OpenSSL" appear in their names without prior written
* permission of the OpenSSL Project.
*
* 6. Redistributions of any form whatsoever must retain the following
* acknowledgment:
* "This product includes software developed by the OpenSSL Project
* for use in the OpenSSL Toolkit (http://www.openssl.org/)"
*
* THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT ``AS IS'' AND ANY
* EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
* IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
* PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR
* ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL,
* SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT
* NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES;
* LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
* HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
* STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
* ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
* OF THE POSSIBILITY OF SUCH DAMAGE.
* ====================================================================
*
* This product includes cryptographic software written by Eric Young
* (eay@cryptsoft.com). This product includes software written by Tim
* Hudson (tjh@cryptsoft.com).
*
*/
```

Original SSLeay License

```
/* Copyright (C) 1995-1998 Eric Young (eay@cryptsoft.com)
* All rights reserved.
*
* This package is an SSL implementation written
* by Eric Young (eay@cryptsoft.com).
* The implementation was written so as to conform with Netscapes SSL.
*
* This library is free for commercial and non-commercial use as long as
* the following conditions are aheared to. The following conditions
* apply to all code found in this distribution, be it the RC4, RSA,
* lhash, DES, etc., code; not just the SSL code. The SSL documentation
* included with this distribution is covered by the same copyright terms
* except that the holder is Tim Hudson (tjh@cryptsoft.com).
```

```
*
* Copyright remains Eric Young's, and as such any Copyright notices in
* the code are not to be removed.
* If this package is used in a product, Eric Young should be given attribution
* as the author of the parts of the library used.
* This can be in the form of a textual message at program startup or
* in documentation (online or textual) provided with the package.
*
* Redistribution and use in source and binary forms, with or without
* modification, are permitted provided that the following conditions
* are met:
* 1. Redistributions of source code must retain the copyright
* notice, this list of conditions and the following disclaimer.
* 2. Redistributions in binary form must reproduce the above copyright
* notice, this list of conditions and the following disclaimer in the
* documentation and/or other materials provided with the distribution.
* 3. All advertising materials mentioning features or use of this software
* must display the following acknowledgement:
* "This product includes cryptographic software written by
* Eric Young (eay@cryptsoft.com)"
* The word 'cryptographic' can be left out if the rouines from the library
* being used are not cryptographic related :-).
* 4. If you include any Windows specific code (or a derivative thereof) from
* the apps directory (application code) you must include an acknowledgement:
* "This product includes software written by Tim Hudson (tjh@cryptsoft.com)"
*
* THIS SOFTWARE IS PROVIDED BY ERIC YOUNG ``AS IS'' AND
* ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
* IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
* ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE
* FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL
* DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS
* OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
* HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT
* LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY
* OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF
* SUCH DAMAGE.
*
* The licence and distribution terms for any publically available version or
* derivative of this code cannot be changed. i.e. this code cannot simply be
* copied and put under another distribution licence
* [including the GNU Public Licence.]
*/
```

A.4 Java Notice

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistribution of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistribution in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

Neither the name of Sun Microsystems, Inc. or the names of contributors may be used to endorse or promote products derived from this software without specific prior written permission.

This software is provided “AS IS,” without a warranty of any kind. ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE HEREBY EXCLUDED. SUN MICROSYSTEMS, INC. ("SUN") AND ITS LICENSORS SHALL NOT BE LIABLE FOR ANY DAMAGES SUFFERED BY LICENSEE AS A RESULT OF USING, MODIFYING OR DISTRIBUTING THIS SOFTWARE OR ITS DERIVATIVES. IN NO EVENT WILL SUN OR ITS LICENSORS BE LIABLE FOR ANY LOST REVENUE, PROFIT OR DATA, OR FOR DIRECT, INDIRECT, SPECIAL, CONSEQUENTIAL, INCIDENTAL OR PUNITIVE DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF THE USE OF OR INABILITY TO USE THIS SOFTWARE, EVEN IF SUN HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

You acknowledge that this software is not designed, licensed or intended for use in the design, construction, operation or maintenance of any nuclear facility.


BROADCOM®

Emulex® OneCommand® Manager Command Line Interface for LightPulse® Adapters

User Guide

Version 11.4
September 6, 2017

OCM-CLI-LPE-UG114-100

Corporate Headquarters
San Jose, CA

Website
www.broadcom.com

Chapter 1: Introduction

The OneCommand® Manager command line interface (CLI) is a comprehensive management utility for Emulex® adapters. The CLI provides support for commonly used commands without requiring the installation of the OneCommand Manager graphical user interface (GUI). The OneCommand Manager CLI console application name is HbaCmd. At the command line interface, a single operation is performed by entering `hbacmd`, followed by a CLI client command and its possible parameters.

The OneCommand Manager application can be installed on multiple operating systems: Windows, Linux, and Solaris. For VMware ESXi hosts, use the OneCommand Manager application for VMware vCenter. For details, refer to the *Emulex OneCommand Manager for VMware vCenter for LightPulse Adapters User Guide*. You can also manage adapters using the OneCommand Manager CLI on Windows, but you must install and use the appropriate Emulex CIM Provider on those VMware hosts.

NOTE The Solaris operating system is supported only on LPe16202/OCe15100 adapters.

NOTE For VMware ESXi hosts, when advanced adapter management capabilities are required (for example, port disablement), use the OneCommand Manager for VMware vCenter Server. For more details, refer to the *Emulex OneCommand Manager for VMware vCenter for LightPulse Adapters User Guide*.

This product supports the following Emulex LightPulse® host bus adapters (HBAs) and converged fabric adapters (CFAs):

- LPe12000-series adapters
- LPe15000-series adapters
- LPe16000-series adapters, including LPe16202/OCe15100 adapters
- LPe31000-series adapters
- LPe32000-series adapters

For supported versions of operating systems and platforms, go to http://www.broadcom.com.

1.1 Abbreviations

API	application programming interface
BIOS	basic input-output system
CFA	converged fabric adapter
CIMOM	CIM Model Object Manager
CLI	command line interface
CSV	comma separated value
DAC	direct-attach copper
D_ID	destination identifier
DCB	Data Center Bridging
DCBX	Data Center Bridging Capabilities Exchange
DH	Diffie-Hellman
DHCHAP	Diffie-Hellman Challenge Handshake Authentication Protocol
ETS	Enhanced Transmission Selection

FA-PWWN	Fabric Assigned WWN
FAT	file allocation table
FC	Fibre Channel
FCF	Fibre Channel over Ethernet Forwarder
FCoE	Fibre Channel over Ethernet
FEC	forward error correction
FIP	FCoE Initialization Protocol
GFO	Get Fabric Object
GUI	graphical user interface
HBA	host bus adapter
IP	internet protocol
LDAP	Lightweight Directory Access Protocol
LLDP	Link Layer Discovery Protocol
LUN	logical unit number
MAC	Media Access Control
NIC	network interface card
NVRAM	nonvolatile random access memory
OAS	Optimized Access Storage
OB	open boot
OS	operating system
PAM	pluggable authentication module
PCI	Peripheral Component Interconnect
PFC	priority flow control
PG	priority group
POST	power-on self-test
PXE	Pre-boot execution Environment
QSFP	Quad Small Form-factor Pluggable
RHEL	Red Hat Enterprise Linux
Rx	receive
SAN	storage area network
SCSI	Small Computer Systems Interface
SFCB	Small Footprint CIM Broker
SFP	small form-factor pluggable
SLES	SUSE Linux Enterprise Server
TCP	Transmission Control Protocol
Tx	transmit
UEFI	Unified Extensible Firmware Interface
VLAN	virtual local area network
VLAN ID	VLAN identifier
VPD	vital product data
vPort	virtual port
WWN	World Wide Name
WWNN	World Wide Node Name

WWPN	World Wide Port Name
XML	Extensible Markup Language

1.2 OneCommand Manager Secure Management

OneCommand Manager Secure Management enables system administrators to further enhance the active management security of their networks. Using Secure Management, administrators can define each user's privileges for managing both local and remote adapters. When running in Secure Management mode, users must specify their user name and password to run the OneCommand Manager CLI. When users are authenticated, only they can perform the functions allowed by the OneCommand Manager user group to which they belong. If your systems are running in an LDAP or Active Directory domain, the OneCommand Manager CLI will authenticate the user with those defined in that domain. For Linux and Solaris systems, this is accomplished using PAM.

NOTE OneCommand Manager Secure Management is supported on Linux, Solaris, and Windows, but it is not supported on VMware hosts. For VMware hosts, the CIM credentials are used.

Administrators set up user accounts such that a user belongs to one of the OneCommand Manager user groups. The user groups define the management capabilities for the user. Table 1 defines the OneCommand Manager user groups and each group's management capabilities.

Table 1 Secure Management User Privileges

Group Name	OneCommand Manager Capability
ocmadmin	Allows full active management of local and remote adapters
ocmlocaladmin	Permits full active management of local adapters only
ocmuser	Permits read-only access of local and remote adapters
ocmlocaluser	Permits read-only access of local adapters

On Linux or Solaris systems, the *unix getent group* utility can be run on the target host system's command shell to verify the correct configuration of the groups. The groups, and users within the groups, appear in the output of this command.

NOTE Although users can belong to the administrator group or be a root user, they will not have full privileges to run the OneCommand Manager unless they are also a member of the ocmadmin group. Otherwise, when Secure Management is enabled, a root user or an administrator can manage only local adapters (similar to the ocmlocaladmin user).

Remote management operations between two machines is allowed or denied depending on the Secure Management status of the machines, and the domains to which the machines belong. The following tables (Table 2, Table 3, and Table 4) list the expected behavior for each machine domain condition (assuming appropriate user credentials are used).

Table 2 Active Commands: Machines on Same Domain

	Remote Server (Secure)	Remote Server (Not Secure)
Client (Secure)	Allowed	Denied [1]
Client (Not Secure)	Denied	Allowed

1. Informs you of an unsecured server that you might want to secure.

Table 3 Active Commands: Machines on Different Domains

	Remote Server (Secure)	Remote Server (Not Secure)
Client (Secure)	Denied[1]	Denied [2]
Client (Not Secure)	Denied	Allowed

1. Allowed if the user name and password are the same on both domains.
2. Informs you of an unsecured server that you might want to secure.

Table 4 Passive Commands: Machines on Any Domain

	Remote Server (Secure)	Remote Server (Not Secure)
Client (Secure)	Allowed	Allowed
Client (Not Secure)	Allowed	Allowed

1.2.1 OneCommand Manager Secure Management Configuration Requirements

For systems to run in the OneCommand Manager Secure Management environment, they must be configured to provide the following two capabilities:

- Authentication – On Linux and Solaris, this is accomplished by using the PAM interface and must be configured as follows:
  - On Solaris, place the correct value in the auth section of the `/etc/pam.d/other` file, or its earlier equivalent, `/etc/pam.conf`.

    NOTE For Solaris systems, you must use `useradd -G groupname` for authentication to work. You cannot use a lowercase **g**.
  - On Linux, it is the `/etc/pam.d/passwd` file auth section, or the equivalent.
- User Group Membership – From the host machine, OneCommand Manager Secure Management must be able to access the OneCommand Manager group to which the user belongs. For Linux and Solaris systems, it uses the `getgrnam` and `getgrid` C-library API calls. The equivalent to the API calls can be obtained by typing `getent group` from the shell command line. If the four OneCommand Manager group names are listed with their member users, the system is ready to use OneCommand Manager Secure Management.

1.3 Secure Management Installation

The enabling or disabling of the Secure Management feature is specified at OneCommand Manager installation time. This can be accomplished either interactively or by using dedicated installation switches on Windows, Linux, and Solaris. On Linux and Solaris, if the OneCommand Manager groups described in Table 1 are not configured on the machine at the time of the OneCommand Manager installation, the installation will fail when the Secure Management feature is selected.

NOTE Only a user with administrator or root privileges can enable or disable the Secure Management feature on a local host machine.

Management mode cannot be used if Secure Management is enabled.

1.3.1 Linux and Solaris

This section describes the Secure Management installation options for the Linux and Solaris operating systems.

1.3.1.1 Interactive Installation

Enterprise OneCommand Manager installations performed in Interactive mode ask if OneCommand Manager Secure Management mode should be enabled. If the answer is **yes**, the other management mode questions are skipped. If the answer is **no** to the OneCommand Manager Secure Management mode question, the management mode installation questions follow.

1.3.1.2 Unattended Installation with Install Script Switch Option Support

Enterprise OneCommand Manager installations performed in Unattended mode provide a switch option to enable OneCommand Manager Secure Management. If the OneCommand Manager Secure Management switch is not used with the installation, Secure Management is disabled.

1.3.2 Windows

During OneCommand Manager installations performed in Interactive mode, you are presented with a management mode window where you can select **Secure Management** as the management mode.

1.4 Setting Secure Management Mode for Linux and Solaris

To set the Secure Management mode for the Linux and Solaris operating systems, perform these steps:

1. Log on as root.
2. Set Secure Management:
   - To set Secure Management mode for Linux, type the following command:
     ```
     # /usr/sbin/ocmanager/set_operating_mode
     ```
   - To set Secure Management mode for Ubuntu 14, type the following command:
     ```
     # /opt/emulex/ocmanager/scripts/set_operating_mode.sh
     ```
   - To set Secure Management mode for Solaris, type the following command:
     ```
     # /opt/ELXocm/ocmanager/set_operating_mode
     ```

   Example

   The following example text is displayed:

   ```
   Do you want to enable Secure Management feature for OneCommand? (s/u)
   The secure management feature requires OneCommand groups be configured on
   the LDAP network or the local host machine to provide for OneCommand
   operation.
   Enter 's' to select secure management. (LDAP/NIS OCM group configuration
   required)
   Enter 'u' to run without secure management (default.
   Enter the letter 's' or 'u': s
   ```

1.4.1 Using OneCommand Manager with Secure Management Enabled

To run the OneCommand Manager CLI when Secure Management mode is enabled, you must include your user name and password each time you type a command.

The syntax for entering your user name and password is the following:

`hbacmd <m=sec> <u=userid> <p=password> <command>`

For example:

```
>hbacmd m=sec u=jsmith p=password download 00-12-34-56-78-9A
oc11-4.6.96.2.ufi
```

User names and passwords authenticate the commands. After the credentials are authenticated, the OneCommand Manager CLI will determine to which one of the four user groups you belong and will allow command usage as appropriate.

Chapter 2: Installing and Uninstalling the CLI

This chapter details prerequisites and procedures for installing and uninstalling the OneCommand Manager CLI in the following operating systems: Linux, Solaris, and Windows. It also describes the Secure Management capability and the procedure for starting and stopping daemon processes.

2.1 Linux

The following instructions are for installing and uninstalling the OneCommand Manager CLI on Linux operating systems. You can install Linux with or without an existing OneCommand CLI kit. Additionally, you can install the OneCommand Manager CLI for Citrix-based operating systems.

2.1.1 Citrix

Citrix is based on CentOS Linux; however, for the OneCommand Manager CLI, Citrix is more comparable to VMware – a hypervisor-style server for managing virtual machines. Citrix XenServer 6.5 and 7.0 operating systems require the OneCommand Manager CLI installation.

2.1.2 Installing in Linux without an Existing OneCommand CLI Kit

NOTE For Secure Management, prior to installation, OneCommand groups must be configured on the LDAP network or the local host machine for Secure Management operation. See Section 1.2.1, OneCommand Manager Secure Management Configuration Requirements, for configuration instructions.

2.1.2.1 Linux OneCommand Manager Requirements

For new systems, install the specific Linux driver rpm files before installing the OneCommand Manager CLI.

2.1.2.1.1 Libnl Library

On RHEL 6.x and 7.x, the OneCommand Core rpm file requires the Libnl library. This library is not installed by default, but it can be obtained from the operating system distribution media.

- For i386 RHEL, use the 32-bit libnl library.
- For x86_64 RHEL, use the 64-bit libnl library.
- For PowerPC RHEL, use the 64-bit libnl library.

2.1.2.1.2 libhbaapi Library

To install the OneCommand Manager CLI in Linux without an existing OneCommand CLI, perform these steps:

1. Copy the applications kit tar file to a directory on the installation machine.
2. Change to the directory where you copied the tar file.
3. Untar the file:

   ```
   tar zxvf elxocmcore-<supported_os>-<app_ver>-<rel>.tgz
   ```

4. Change to the Core kit directory created in step 3.

   ```
   cd elxocmcore-<supported_os>-<app_ver>-<rel>
   ```

5. Run the `install.sh` script.

```
./install.sh
```

The Core kit consists of three or four rpm files for each supported architecture and each supported version of Linux. For example:

- `elxocmlibhbaapi-*.rpm` (on 64-bit platforms that support 32-bit applications, there are two of these files)
- `elxocmcore-*.rpm`
- `elxocmcorelibs-*.rpm`

6. When you are prompted, choose whether to enable Secure Management for OneCommand:

```
Do you want to enable Secure Management feature for OneCommand? (s/u)
Enter 's' to select secure management. (LDAP/NIS OCM group configuration
required)
Enter 'u' to run without secure management (default).
Enter the letter 's' or 'u'.
```

If you enter `u`, an additional prompt is given for the management mode:

```
You selected: Secure Management Disabled
Select desired mode of operation for OneCommand Manager:
Enter the number 1, 2, 3, or 4: 1
You selected: 'Local Only Mode'

1 Strictly Local Management:  Only manage the adapters on this host. Management of
                              adapters on this host from other hosts is not allowed.
2 Local Management Plus:      Only manage the adapters on this host. Management of
                              adapters on this host from other hosts is allowed.
3 Full Management:            Manage the adapters on this host and other hosts that
                              allow it. Management of the adapters on this host from
                              another host is allowed.
4 Management Host:            Manage the adapters on this host and other hosts that
                              allow it. Management of the adapters on this host from
                              another host is not allowed.
```

2.1.2.1.3 Unattended Installation

The `install.sh` script can be run in Noninteractive (unattended or quiet) mode. Enter the following command to view the syntax:

```
./install.sh --help
```

To perform an unattended, silent installation, enter the following command, perform these steps:

```
#./install.sh -q2
```

NOTE The management mode default for unattended installation is Local Management Plus.

2.1.3 Installing in Linux with an Existing OneCommand CLI Kit

NOTE The OneCommand Manager Core kit cannot be installed if a previous version of the HBAnyware utility is installed.

Two options are available for installing the OneCommand Manager CLI on a Linux system with an existing OneCommand CLI kit:

- Updating an existing installation – Preserve existing settings

- Performing a clean install – Overwrite existing settings

2.1.3.1 Updating the CLI (Preserving Existing Settings)

To update the OneCommand Manager CLI and preserve settings, you must install the current Core kit as detailed in Section 2.1.2, Installing in Linux without an Existing OneCommand CLI Kit. The `.rpm` file handles the configuration file update. The install script executes an rpm file update (`rpm -U *.rpm`) to update the installed version of the core kit to the current version.

NOTE There is no update path from an HBAnyware 4.x or 3.x core kit to a OneCommand Manager 5.1 or later core kit. You must uninstall previous versions of the HBAnyware utility before installing a OneCommand Manager core kit. For information on uninstalling older versions of HBAnyware, see Section 2.1.5, Uninstalling Older HBAnyware Kits in Linux.

2.1.3.2 Performing a Clean Install (Removing Existing Settings)

1. Uninstall the existing OneCommand Manager CLI using the uninstall script included in the tar file or in the `/usr/sbin/ocmanager/scripts` directory. The configuration files are backed up by rpm with a `.rpmsave` extension.

   For Ubuntu 14 (LPe16202/OCe15100 adapters only), use the uninstall script in the following location:

   ```
   /opt/emulex/ocmanager/scripts/uninstall.sh
   ```

   NOTE If an HBAnyware CLI or Enterprise kit is installed, follow the procedure in Section 2.1.5, Uninstalling Older HBAnyware Kits in Linux.
2. Install the specific rpm file for your driver for Linux version. For information on installing the rpm file, see Section 2.1.2, Installing in Linux without an Existing OneCommand CLI Kit.

2.1.4 Uninstalling in Linux

To uninstall the OneCommand Manager CLI in Linux, perform these steps:

1. Log on as root.
2. Perform one of the following tasks:
   - Run the `uninstall_ocmanager.sh` script located in `/usr/sbin/ocmanager/scripts`.
   - Run the `uninstall.sh` script located in the installation tar file.
   - For Ubuntu 14 (LPe16202/OCe15100 adapters only), use the uninstall script in the following location:
     `/opt/emulex/ocmanager/scripts/uninstall.sh`

2.1.5 Uninstalling Older HBAnyware Kits in Linux

2.1.5.1 Uninstalling an Older HBAnyware Core Kit

Run the following command to remove the Core kit.

```
rpm -e elxlinuxcorekit
```

2.1.5.1.1 Uninstalling an Older HBAnyware Enterprise Kit

1. Perform one of the following tasks:
   - Run the uninstall script located in `/usr/sbin/hbanyware/scripts` to remove the Enterprise kit.
   - Run the uninstall script located in the tar file to remove the Enterprise kit.

If the HBAnyware Security Configurator is installed, you must uninstall it before uninstalling the HBAnyware configuration utility. You must use the uninstall script that shipped with the version of OneCommand Security Configurator that you want to remove and proceed to step 2. If the Security Configurator is not installed, proceed to step 3.

2. If the HBAnyware Security Configurator is installed, follow these steps:
   a. Log on as root.
   b. Change to the directory containing the tar file.
   c. Extract the tar file using the `tar -xvf` command.
   d. Change to the newly created directory.
   e. Type the following uninstall script with the `ssc` parameter specified:
      ```
      ./uninstall ssc
      ```
3. Uninstall the HBAnyware utility and the Application Helper module:
   a. Log on as root.
   b. Change to the directory containing the tar file.
   c. Extract the tar file using the `tar -xvf` command.
   d. Change to the newly created directory.
   e. Uninstall any previously installed versions. Type the following command:
      ```
      ./uninstall
      ```

2.2 Solaris

The following instructions are for installing and uninstalling the OneCommand Manager CLI on Solaris operating systems.

2.2.1 Installing in Solaris

NOTE For Secure Management, prior to installation, OneCommand groups must be configured on the LDAP network or the local host machine for Secure Management operation. See Section 1.2.1, OneCommand Manager Secure Management Configuration Requirements, for configuration instructions.

To install the OneCommand Manager CLI in Solaris, perform these steps:

1. Copy the OneCommand Manager core kit to a temporary directory on the system.
2. Untar the core kit by typing the following command:
   ```
   tar xvf elxocmcore-solaris-<kit version>.tar
   ```
3. Change to the newly created elxocmcore-solaris-<*kit version*> directory:
   ```
   cd ./elxocmcore-solaris-<kit version>/
   ```
4. Run the `install` script and follow the instructions.
   ```
   ./install
   ```

   NOTE The `install` script can also be run in Noninteractive (unattended, quiet) mode. Enter the following command to view the syntax:

   ```
   ./install --help
   ```

If any of the following are already present on the system, the `install` script attempts to remove them first:

- HBAnyware utility
- OneCommand Manager Core kit
- OneCommand Manager Enterprise kit
- Solaris driver utilities

5. When you are prompted, choose whether to enable Secure Management for OneCommand:

```
Do you want to enable Secure Management feature for OneCommand? (s/u)
Enter 's' to select secure management. (LDAP/NIS OCM group configuration
required)
Enter 'u' to run without secure management (default).
Enter the letter 's' or 'u'.
```

If you enter `u` here, an additional prompt is given for the management mode:

```
You selected: Secure Management Disabled
Select desired mode of operation for OneCommand Manager:
Enter the number 1, 2, 3, or 4: 1
You selected: 'Local Only Mode'

1 Strictly Local Management: Only manage the adapters on this host. Management of
                             adapters on this host from other hosts is not allowed.
2 Local Management Plus:     Only manage the adapters on this host. Management of
                             adapters on this host from other hosts is allowed.
3 Full Management:           Manage the adapters on this host and other hosts that
                             allow it. Management of the adapters on this host from
                             another host is allowed.
4 Management Host:           Manage the adapters on this host and other hosts that
                             allow it. Management of the adapters on this host from
                             another host is not allowed.
```

2.2.2 Uninstalling in Solaris

To uninstall the OneCommand Manager CLI in Solaris, perform these steps:

1. Log on as root.
2. Perform one of the following tasks:
   - Run `/opt/ELXocm/scripts/uninstall`.
   - Run the `uninstall` script located in the installation tar file.
   - Enter the command `pkgrm ELXocmcore`.

NOTE The `uninstall` script can also be run in Noninteractive (quiet) mode. Enter the following command to view the syntax:

```
./uninstall --help
```

2.3 VMware ESXi

The OneCommand Manager CLI cannot be run on a VMware ESXi operating system. However, a VMware ESXi host can be accessed remotely from the Windows OneCommand Manager CLI if the Broadcom® Emulex CIM Provider is installed on the ESXi host. For instructions on installing Broadcom Emulex CIM Provider on VMware ESXi operating systems, refer to the *Emulex CIM Provider Package for LightPulse Adapters Installation Guide*.

2.4 Windows

The following instructions are for installing and uninstalling the OneCommand Manager CLI on Windows operating systems. Install the OneCommand Manager CLI in Windows in one of two ways:

- Attended installation – You are present during the installation. You are prompted for more information for the installation to continue.
- Unattended installation – You do not need to be present during the installation. Installation will complete on its own. Installation progress can be displayed as an option.

2.4.1 Installing in Windows by Attended Installation

To install the OneCommand Manager CLI, run the `installation.exe` file for a Windows Core driver kit that does not include the OneCommand Manager GUI, and follow the installation instructions.

Use the following syntax for the installation executable file:

```
elxocmcore-windows-<arch>-<kit version>.exe
```

- *<arch>* is either x64 or x86.
- *<kit version>* represents the complete kit version.

For example, at the command prompt, type the following command:

```
elxocmcore-windows-x64-5.0.2.14-1.exe
```

2.4.2 Installing in Windows by Unattended Installation

To install the OneCommand Manager CLI in Windows unattended, perform these steps:

1. From http://www.broadcom.com, download the x64 or x86 OneCommand Manager Core kit installation file to your system.
2. Use the following syntax for the installation executable file:

   ```
   elxocmcore-windows-<arch>-<kit version>.exe <option>
   ```
3. Activate the kit with switch `/q` or `/q2`.
   - The `/q` switch displays progress reports.
   - The `/q2` switch does not display progress reports.
4. Either enable Secure Management mode by adding the `sec=1` argument or disable it by adding `sec=0`. If the `sec` argument is not entered, Secure Management is disabled by default. See Section 1.2, OneCommand Manager Secure Management, for more information.

   To enable Secure Management mode, at the command prompt, type the following command:

   ```
   elxocm-windows-<arch>-<kit version>.exe sec=1 /q2
   ```

   To disable Secure Management mode, at the command prompt, type the following command:

   ```
   elxocm-windows-<arch>-<kit version>.exe sec=0 /q2
   ```

   NOTE Two management mode defaults are available for unattended installation:
   - `mmode=3` (full management mode)
   - `achange=1`
5. Select a management mode by adding the `mmode` argument, and select the ability to change the management mode by adding the `achange` argument with selected values as in the following example.

NOTE If you enabled Secure Management mode in step 4 and attempt to enter an `mmode` value, a `conflicting parameters` error can occur.

For example, at the command prompt type the following command:

```
elxocm-windows-x64-5.01.00.10-4.exe mmode=3 achange=1 /q2
```

The following are the possible `mmode` values:

`1` – Local Only Management mode

`2` – Local Plus Management mode

`3` – Full Management mode

`4` – Local Plus Management mode and Read Only

`5` – Full Management mode and Read Only

`6` – Management host

The following are the possible `achange` values:

`0` – Do not allow management mode to change

`1` – Allow management mode to change

You can also set the following optional parameters:

— `MHost` – This optional switch allows a nonmanagement-host user to select a management host with which to register. If this switch is not specified, the default value of `0` is used, and the capability will be disabled. If the switch is specified, the value can be a host name or an IP address, which is validated by the installer. An error message appears if `mmode` is set as Local Only management mode or Management Host mode.

— `excl` – This optional switch allows the nonmanagement-host user to select whether the OneCommand Manager application processes requests exclusively from the management host specified by the `MHost` switch. This option is only accepted if accompanied by a valid `MHost` value; otherwise, an error message appears. If this switch is not specified, the default value of `0` is used. If the switch is specified, the valid values are:

`0` – Remotely managed by other hosts.

`1` – Remotely managed by management host only.

— `Mtcp` – This optional switch allows you to enable or disable remote management and to specify the TCP/IP port number over which management occurs. If this switch is not specified, the default TCP/IP port number 23333 is used.

If the management host option is selected, you must select the default port number or enter a valid TCP/IP port number on the command line. A value of `0` will not be accepted.

If one of the nonmanagement host options is selected, you can enter the TCP/IP port number on the command line.

2.4.3 Uninstalling in Windows

You can uninstall the OneCommand Manager CLI in Windows in one of two ways:

- Through the Control Panel
- Through the command line

2.4.3.1 Uninstalling through the Control Panel

To uninstall the OneCommand Manager CLI in Windows through the Control Panel, perform these steps:

1. In the Control Panel, select **Programs and Features**.
2. If present, select **Emulex OCManager CLI** [*version*], and click **Uninstall/Change;** you are prompted to continue. Click **Yes**.

The OneCommand Manager CLI components are removed from the system.

2.4.3.2 Uninstalling through the Command Line

To uninstall the OneCommand Manager CLI in Windows through the command line, perform these steps:

1. Change to the appropriate uninstall directory:

   ```
   cd <Installation Location>\Emulex\Util\Uninstall
   ```

2. Type the following command:

   ```
   uninstall_OCManager_Core.bat
   ```

2.5 Starting and Stopping Daemon Processes for Linux and Solaris Installations

On Linux and Solaris machines, you can stop and start the OneCommand Manager daemon processes using the `stop_ocmanager` and `start_ocmanager` scripts, respectively. These are found in the following OneCommand Manager installation directories:

- Linux – `/usr/sbin/ocmanager`
- Ubuntu 14 – `/opt/emulex/ocmanager/scripts`:
  - `stop_ocmanger.sh`
  - `start_ocmanager.sh`
- Solaris – `/opt/ELXocm`

The `elxhbamgrd` daemon process (included with OneCommand Manager CLI) is affected by these scripts. It is a remote management daemon that services requests from OneCommand Manager clients running on remote host machines.

The daemon processes start at system boot time.

Chapter 3: Updating the OneCommand Manager CLI to the OneCommand Manager Enterprise Kit

NOTE The full-featured OneCommand Manager Enterprise kit is not supported on Citrix XenServer 6.x, Citrix XenServer 7.x, or VMware ESXi server.

This chapter details procedures for updating the OneCommand Manager CLI to the OneCommand Manager Enterprise kit in Linux, Solaris, and Windows operating systems. An update can be performed only if the version of the OneCommand Manager Enterprise kit is the same or later than the OneCommand Manager CLI version.

NOTE You cannot update a OneCommand Manager CLI with a previous version of the OneCommand Manager Enterprise kit.

3.1 Linux

To update from the OneCommand Manager CLI to the full-featured OneCommand Manager Enterprise kit in Linux, run the `install.sh` script of the OneCommand Manager Enterprise kit.

The install script executes an rpm file update (`rpm -U *.rpm`) to update the installed core kit to an enterprise kit.

3.2 Solaris

To update from the OneCommand Manager CLI to the full-featured OneCommand Manager Enterprise kit in Solaris, perform these steps:

1. Download the OneCommand Manager Enterprise kit to a temporary directory on your system.
2. Untar the OneCommand Manager Enterprise kit tar file:

   `tar xvf elxocm-solaris-<kit version>.tar`
3. Change to the newly created `elxocm-solaris-<kit version>` directory:

   `cd ./elxocm-solaris-<kit version>/`
4. Run the `install` script and follow the instructions:

   `./install`

The `install` script can also be run in Noninteractive (quiet) mode. To view the syntax, type the following command:

```
/install --help
```

3.3 Windows

To update from the OneCommand Manager CLI to the full-featured OneCommand Manager Enterprise kit in Windows:

From the desktop, run the `elxocm-windows-<kit version>.exe` file that contains the full application kit. Running this executable file removes the OneCommand Manager CLI and installs a full-featured version of the OneCommand Manager application that includes the CLI and the GUI.

Chapter 4: CLI Client Command Usage

The CLI Client component of the OneCommand Manager application provides access to the capabilities of the Remote Management library or the CIM interface from a console command prompt to get the management information.

4.1 Overview

The CLI Client is intended for use in command shells or scripted operations from within shell scripts or batch files. The CLI Client is a console application named `HbaCmd`. A single operation is performed by typing `hbacmd` at the command line, followed by a CLI client command and its possible parameters. For example:

```
hbacmd [cli options] <command> [parameters]
```

The CLI options are specified for running the CLI commands on remote hosts or with Secure Management.

When the specified operation is completed, the command prompt is displayed. For a majority of commands, the first parameter following the command is the WWPN or MAC address of the port that the command is to act upon.

4.1.1 CLI in Read-Only Mode

The CLI does not allow the execution of some commands if it is configured for Read-Only mode. The following error message is returned if such a command is attempted:

```
Error: Read-only management mode is currently set on this host. The
requested command is not permitted in this mode.
```

4.2 HbaCmd Syntax Usage

The following syntax rules and usage apply to the `HbaCmd` application:

- Parameters denoted within angle brackets < > are required.
- Parameters denoted within square brackets [ ] are optional.
- For Linux and Solaris (which are case-sensitive), program names must be in lowercase letters. Therefore, the command line must begin with `hbacmd` (rather than `HbaCmd`). Windows is not case-sensitive, so the program name is not required to be in all lowercase letters.
- To run the command on a remote host, an IP address or a host name must be specified using the `h` option with the following syntax:

  ```
  hbacmd [h=IP_Address[:port] | Hostname[:port]] <command> [parameters]
  ```

  — If the `h` option is omitted, the command is run on the local host.

  — If the `h` option is specified, the command is sent to the specified remote host (assuming it is specified correctly, the remote host is up, and the remote host is running the OneCommand Manager remote management agent.

  — The *`:port`* option is optional. If it is omitted, the OneCommand Manager remote management protocol uses the default TCP port. If it is specified, it uses the user-specified TCP port.

  — **Examples**

  Using the IP address:

  ```
  hbacmd h=138.239.91.121 ListHBAs
  ```

Using the host name:

```
hbacmd h=cp-hp5670 ListHBAs
```

— The `h` option is available for all commands except for the `AddHost`, `RemoveHost`, and `Version` commands.

- For FC and FCoE functions, the WWPN of the adapter must be specified. Where the WWPN is specified, each pair of numbers within the WWPN is separated by colons (`:`) or spaces ( ). If space separators are used, the entire WWPN must be enclosed in quotation marks (" ").

  For example, the following command displays the port attributes for the adapter with the specified WWPN:

  ```
  hbacmd PortAttributes 10:00:00:00:c9:20:20:20
  ```

- For NIC functions, the MAC address must be specified. Where a MAC address is specified, each pair of numbers within the MAC address is separated by a dash (–).

  For example, the following command displays the server attributes for the server where the NIC function is running the NIC port with the specified MAC address:

  ```
  hbacmd ServerAttributes 00-11-22-33-44-55
  ```

- For NIC functions, only the permanent MAC address is supported for the port address parameter on an `HbaCmd` command line.

  Normally, for a NIC function, the function's permanent MAC address and current MAC address parameters are equal. However, it is possible to set a user-specified (current) MAC address that is different from the permanent MAC address. Also, for some implementations, it is possible to have multiple NIC functions with the same current MAC addresses, but with unique permanent MAC addresses. Therefore, to be sure that OneCommand Manager can access the correct function, only the permanent MAC address is supported.

  NOTE Both the permanent MAC address and the current MAC address are displayed by using the `ListHBAs` command. See Section 5.13.7, ListHBAs.

4.3 Secure Management CLI Interface

The Secure Management CLI interface is supported by the Linux, Solaris, and Windows operating systems.

NOTE Users with root or administrator privileges on the local machine will retain full configuration capability in the OneCommand Manager CLI without the use of credentials (local machine only).

4.3.1 Device Management Using the Secure Management Interface

To run the `HbaCmd` CLI client application when the Secure Management feature is enabled, each invocation must include a user name and password. The user name and password options are added to the existing `HbaCmd` command in the same way as they are for CIM commands, except the `<m=cim>` option is replaced by the `<m=sec>` option (to distinguish it from a CIM command). For example:

Without Secure Management (or if running as root or administrator):

```
hbacmd <cmd>
```

With Secure Management (as non-root or non-administrator user):

```
hbacmd <m=sec> <u=userid> <p=password> <cmd>
```

4.3.2 Syntax Rules for the Secure Management Interface

For the Secure Management interface, all of the syntax rules in Section 4.2, HbaCmd Syntax Usage, apply.

Example

In Windows, to download firmware on an LPe16202/OCe15100 adapter managed on a remote host at IP address 192.168.1.122 using the Secure Management interface, run the following command:

```
hbacmd h=192.168.1.122 m=sec u=jsmith p=password download 00-12-34-56-78-9A
lancer_a11.2.123.45.grp
```

4.4 CIM Client Interface

NOTE In Linux and Solaris, you cannot use `HbaCmd` as a CIM client.

4.4.1 Device Management Using the CIM Interface

VMware on the hypervisor-based ESXi platforms use the CIM as the only standard management mechanism for device management.

For VMware ESXi hosts, you can manage adapters using the OneCommand Manager CLI on Windows, but you must install and use the appropriate Broadcom Emulex CIM Provider on the VMware ESXi host. For installation, refer to the *CIM Provider Package for LightPulse Adapters Installation Guide*.

NOTE For VMware ESXi hosts, if advanced adapter management capabilities are required, use the OneCommand Manager for VMware vCenter Server. For more details, refer to the *OneCommand Manager for VMware vCenter for LightPulse Adapters User Guide*.

4.4.2 Syntax Rules for the CIM Interface

For the CIM interface, all the syntax rules in Section 4.2, HbaCmd Syntax Usage, apply, except that the `h` option is required. Additionally, the `m=cim` parameter is required in the command line for getting the data from the ESXi host. For example:

```
hbacmd h=192.168.1.110 m=cim u=root p=password n=root/emulex listhbas
```

4.4.2.1 Syntax Options and Setting CIM Credentials

For issuing CIM-based commands, two main syntax options are available.

Option A

hbacmd <h=*IP_Address*[:*port*]> m=cim [u=*userid*] [p=*password*] [n=*root*/emulex] <*command*> <*WWPN*>

Option B

```
hbacmd <h=IP_Address[:port]> <m=cim> <command>
```

Before using the option B syntax, you must set the CIM credentials. Perform one of the following tasks:

- Set the default CIM credentials using the `SetCimCred` command (see Section 5.13.10, SetCimCred). This command sets only the CIM credentials. After you have set them, subsequent `HbaCmd` commands do not require you to specify the CIM credentials on the command line.

  Command syntax:

```
hbacmd setcimcred <username> <password> <namespace> <portnum>
```

- Add the host IP address with CIM credentials using the `AddHost` command.

  Command syntax:

```
hbacmd <m=cim> [u=userid] [p=password] [n=namespace] addhost <IP_Address>
```

4.4.2.1.1 Default CIM Credentials

If you specify the command with the CIM method `m=cim` without specifying the CIM credentials (*`userid`*, *`password`*, or *`namespace`*), the default value for the missing CIM credential is obtained in the following order:

1. The information entered using the `addhost` command is looked up.
2. If no values exist, the information entered using the `setcimcred` command is used.
3. If no values exist, the following defaults are used:

```
username=root
password=root
namespace=root/emulex
portnum=5988
```

4.4.2.2 Example of Using the CIM Interface to Display Adapters

In Windows, to display a list of adapters managed for a specified host using the CIM interface, run the following command:

```
hbacmd h=10.192.113.128 m=cim u=root p=root n=root/emulex listhbas
```

For a list of `HbaCmd` commands supported through the CIM interface, see Table 2, CLI Client Command Reference.

Chapter 5: CLI Client Command Descriptions

CLI Client commands are organized by command groups. Two tables are presented for your convenience; a table organized by command group and another by alphabetically listing CLI Client commands.

The following table shows each command group with a short description and the commands in each group. After you determine the command group of interest, click the command link and go directly to the command you selected.

Table 5 CLI Client Command Reference Functional Groups

Command Group	Description	Commands
Attributes Commands	This group manages the display of adapter, port, server attributes, and port statistics for each adapter specified. You can also set the port speed on LPe16202/OCe15100 adapters.	HbaAttributes ServerAttributes SetPhyPortSpeed SetPortEnabled
Authentication Commands	These commands configure a DHCHAP connection between an FC port and a switch port. Authentication commands apply to only LPe12000-series adapters.	AuthConfigList DeleteAuthConfig GetAuthConfig GetAuthStatus InitiateAuth SetAuthConfig SetPassword
Boot Commands	This group manages the commands that enable or disable network boot for NIC ports or the boot code for FC adapter ports. You can also show and change FC and FCoE boot parameters.	EnableBootCode GetBootParams SetBootParam SetLinkSpeed
DCB Commands	These commands display and set the DCB and LLDP parameters for FCoE and NIC ports on LPe16202/OCe15100 adapters.	GetDCBParams GetPGInfo SetCnaPGBW SetDCBParam SetDCBPriority
Diagnostic Commands	This group provides commands that enable you to detect cabling problems, to examine transceiver data, and to flash memory load lists. Additionally, you can run specific diagnostic tests, such as the Loopback test and the POST.	D_PortTest EchoTest GetBeacon GetXcvrData LoadList LoopBackTest LoopMap PciData PostTest SetBeacon Wakeup
Driver Parameter Commands	Use the driver parameter commands to show, set, and save the driver parameter values. You can also change the parameters back to factory default values.	DriverConfig GetDriverParams GetDriverParamsGlobal SaveConfig SetDriverParam SetDriverParamDefaults

Table 5 CLI Client Command Reference Functional Groups (Continued)

Command Group	Description	Commands
Dump Commands	Use the diagnostic dump feature to create a dump file for a selected adapter. Dump files contain information, such as firmware version, driver version, and operating system information. This information is useful for troubleshooting an adapter, but it is unavailable in Read-Only mode.	DeleteDumpFiles Dump GetDumpDirectory GetDumpFile GetDumpFileNames GetRetentionCount SetDumpDirectory SetRetentionCount
FCoE Commands	This group of commands manages the FIP parameters and displays the FCF for an LPe16202/OCe15100 adapter in NIC+FCoE mode.	GetFCFInfo GetFIPParams SetFIPParam
Firmware Commands	These commands enable you to view and set firmware parameters.	getfwparams setfwparam
LUN Masking Commands	The commands in this group manage LUN masking activities. LUN masking is supported only for FC and FCoE ports.	GetLunList GetLunUnMaskByHBA GetLunUnMaskByTarget RescanLuns SetLunMask
LUN ExpressLane Commands	This group of commands enables, disables, and displays the ExpressLane™ status on a particular LUN. You can also assign a frame priority to an ExpressLane LUN if the adapter and the switch support it. LUN ExpressLane commands do not apply to LPe16202/OCE15100 adapters in NIC+FCoE mode.	GetExpressLaneLunList SetExpressLaneLunState GetLunXLaneConfig SetLunXLaneConfig
Miscellaneous Commands	This group contains commands that do not belong in other groups.	AddHost Download ExportSANInfo FecEnable GetCimCred GetVPD ListHBAs Reset SetCimCred TargetMapping Version

Table 5 CLI Client Command Reference Functional Groups (Continued)

Command Group	Description	Commands
Persistent Binding Commands	This group of commands facilitates persistent binding operations. These commands are supported only for FC and FCoE ports.	AllNodeInfo BindingCapabilities BindingSupport PersistentBinding RemoveAllPersistentBinding RemovePersistentBinding SetBindingSupport SetPersistentBinding
vPort Commands	vPort commands manage virtual ports and functions only on FC and FCoE adapters. In Linux, VPorts do not persist across system reboots.	CreateVPort DeleteVPort ListVPorts VPortTargets
WWN Management Commands	WWN management validates WWNs to avoid WWPN duplication; however, WWNN duplication is acceptable. You might see error and warning messages if a name duplication is detected. Make sure that the activation requirement is fulfilled after each WWN is changed or restored. If pending changes exist, some diagnostic and maintenance features are not available.	ChangeWWN GetWWNCap ReadWWN RestoreWWN

Table 2 lists each command alphabetically and shows the operating system and CIM Interface support for each command. A linked page number for each command is provided for your convenience. A check mark (✓) designates a supported command for a particular operating system and CIM interface.

NOTE For VMware ESXi, two options support the CLI:

- Using the OneCommand Manager CLI on Windows with the appropriate Emulex CIM Provider installed on a VMware host. These commands are covered in this section.
- Using the OneCommand Manager for VMware vCenter command line interface (`elxvcpcmd`). Although the available commands are listed in Table 2 for your convenience, refer to the *OneCommand Manager for VMware vCenter for LightPulse Adapters User Guide* for specific information.

Table 6 CLI Client Command Reference

Command	Linux		Solaris	Windows	CIM Interface Support	elxvcpcmd	Section
	RHEL, SLES, Ubuntu, and Oracle	Citrix					
`AddHost`	✓	✓	✓	✓	✓		Section 5.13.1
`AllNodeInfo`	✓	✓	✓	✓	✓	✓	Section 5.14.1
`AuthConfigList`			✓	✓			Section 5.3.1
`BindingCapabilities`			✓	✓			Section 5.14.2
`BindingSupport`			✓	✓			Section 5.14.3
`ChangeWWN`	✓	✓	✓	✓	✓	✓	Section 5.16.1
`CreateVPort`	✓		✓	✓			Section 5.15.1
`DPortTest`	✓	✓		✓	✓	✓	Section 5.6.1

Table 6 CLI Client Command Reference (Continued)

Command	Linux		Solaris	Windows	CIM Interface Support	elxvcpcmd	Section
	RHEL, SLES, Ubuntu, and Oracle	Citrix					
DeleteAuthConfig			✓	✓			Section 5.3.2
DeleteDumpFiles	✓	✓	✓	✓	✓	✓	Section 5.8.1
DeleteVPort	✓		✓	✓			Section 5.15.2
Download	✓	✓	✓	✓	✓	✓	Section 5.13.2
DriverConfig	✓	✓		✓		✓	Section 5.7.1
Dump	✓	✓	✓	✓	✓	✓	Section 5.8.2
EchoTest	✓	✓	✓	✓	✓	✓	Section 5.6.2
EnableBootCode	✓	✓	✓	✓	✓	✓	Section 5.4.1
ExportSANInfo	✓	✓	✓	✓			Section 5.13.3
FecEnable	✓	✓	✓	✓	✓	✓	Section 5.13.4
GetBeacon	✓	✓	✓	✓	✓	✓	Section 5.6.3
GetAuthConfig			✓	✓			Section 5.3.3
GetAuthStatus			✓	✓			Section 5.3.4
GetBootParams	✓	✓	✓	✓		✓	Section 5.4.2
GetCimCred				✓			Section 5.13.5
GetDCBParams	✓	✓	✓	✓	✓	✓	Section 5.5.1
GetDriverParams	✓	✓	✓	✓	✓	✓	Section 5.7.2
GetDriverParamsGlobal	✓	✓	✓	✓	✓	✓	Section 5.7.3
GetDumpDirectory	✓	✓	✓	✓	✓	✓	Section 5.8.3
GetDumpFile	✓	✓	✓	✓	✓	✓	Section 5.8.4
GetDumpFileNames	✓	✓	✓	✓	✓	✓	Section 5.8.5
GetExpressLaneLUNList	✓	✓		✓	✓	✓	Section 5.12.1
GetFCFInfo	✓	✓	✓	✓	✓	✓	Section 5.9.1
GetFIPParams	✓	✓	✓	✓	✓	✓	Section 5.9.2
GetFwParams	✓	✓		✓			Section 5.10.1
GetLunList	✓			✓	✓	✓	Section 5.11.1
GetLunUnMaskByHBA				✓			Section 5.11.2
GetLunUnMaskByTarget				✓			Section 5.11.3
GetPGInfo	✓	✓	✓	✓	✓	✓	Section 5.5.2
GetRetentionCount	✓	✓	✓	✓	✓	✓	Section 5.8.6
GetVPD	✓	✓	✓	✓	✓	✓	Section 5.13.6
GetWWNCap	✓	✓	✓	✓	✓	✓	Section 5.16.2
GetXcvrData	✓	✓	✓	✓	✓	✓	Section 5.6.4
HbaAttributes	✓	✓	✓	✓	✓	✓	Section 5.2.1
Help	✓	✓	✓	✓	N/A	✓	Section 5.1
InitiateAuth			✓	✓			Section 5.3.5
ListHBAs	✓	✓	✓	✓	✓	✓	Section 5.13.7
ListVPorts	✓	✓	✓	✓	✓	✓	Section 5.15.3

Table 6 CLI Client Command Reference (Continued)

Command	Linux		Solaris	Windows	CIM Interface Support	elxvcpcmd	Section
	RHEL, SLES, Ubuntu, and Oracle	Citrix					
LoadList	✓	✓	✓	✓	✓		Section 5.6.5
LoopBackTest	✓	✓	✓	✓	✓	✓	Section 5.6.6
LoopMap	✓	✓	✓	✓		✓	Section 5.6.7
PciData	✓	✓	✓	✓	✓		Section 5.6.8
PersistentBinding		✓	✓	✓		✓	Section 5.14.4
PortAttributes	✓	✓	✓	✓	✓		Section 5.2.2
PortStatistics	✓	✓	✓	✓	✓	✓	Section 5.2.3
PostTest	✓	✓	✓	✓	✓	✓	Section 5.6.9
ReadWWN	✓	✓	✓	✓	✓	✓	Section 5.16.3
RemoveAllPersistent Binding			✓	✓		✓	Section 5.14.5
RemovePersistentBinding			✓	✓			Section 5.14.6
RemoveHost	✓	✓	✓	✓	✓		Section 5.13.8
RescanLuns				✓		✓	Section 5.11.4
Reset	✓	✓	✓	✓	✓		Section 5.13.9
RestoreWWN	✓	✓	✓	✓	✓	✓	Section 5.16.4
SaveConfig	✓	✓	✓	✓		✓	Section 5.7.4
ServerAttributes	✓	✓	✓	✓	✓	✓	Section 5.2.4
SetAuthConfig			✓	✓		✓	Section 5.3.6
SetBeacon	✓	✓	✓	✓	✓		Section 5.6.10
SetBindingSupport			✓	✓		✓	Section 5.14.7
SetBootParam	✓	✓	✓	✓			Section 5.4.3
SetCnaPGBW	✓	✓	✓	✓	✓	✓	Section 5.5.3
SetCimCred				✓		✓	Section 5.13.10
SetDCBParam	✓	✓	✓	✓	✓		Section 5.5.4
SetDCBPriority	✓	✓	✓	✓	✓	✓	Section 5.5.5
SetDriverParam	✓	✓	✓	✓	✓	✓	Section 5.7.5
SetDriverParamDefaults	✓	✓	✓	✓		✓	Section 5.7.6
SetDumpDirectory					✓	✓	Section 5.8.7
SetExpressLaneLUNState	✓	✓		✓	✓	✓	Section 5.12.2
SetFIPParam	✓	✓	✓	✓	✓	✓	Section 5.9.3
SetFwParam	✓	✓		✓			Section 5.10.2
SetLinkSpeed	✓	✓	✓	✓	✓	✓	Section 5.4.4
SetLunMask						✓	Section 5.11.5
SetPassword			✓	✓		✓	Section 5.3.7
SetPersistentBinding			✓	✓			Section 5.14.8
SetPhyPortSpeed	✓	✓	✓	✓	✓	✓	Section 5.2.5
SetPortEnabled	✓	✓	✓	✓	✓	✓	Section 5.2.6

Table 6 CLI Client Command Reference (Continued)

Command	Linux		Solaris	Windows	CIM Interface Support	elxvcpcmd	Section
	RHEL, SLES, Ubuntu, and Oracle	Citrix					
SetRetentionCount	✓	✓	✓	✓	✓	✓	Section 5.8.8
TargetMapping	✓	✓	✓	✓	✓	✓	Section 5.13.11
Version	✓	✓	✓	✓	✓	✓	Section 5.13.12
VPortTargets	✓	✓	✓	✓	✓		Section 5.15.4
Wakeup	✓	✓	✓	✓			Section 5.6.11

5.1 Help

This command displays command information for the HbaCmd application. Without using its optional parameters, the Help command lists all the commands in their respective groups. Using the optional parameter *GroupName*, it lists the commands in a group. Using the optional parameter *CmdName*, it shows the details for a specific command.

Supported By

Linux, Solaris, and Windows

Syntax

Help [*GroupName*] [*CmdName*]

Parameters

GroupName	This optional parameter lists the commands in a particular group.
CmdName	This optional parameter shows the details for a particular CLI command.

Examples

This Help command example lists all the commands in their respective groups:

```
hbacmd help
```

This Help command example shows the details for the SetDCBParam command:

```
hbacmd help setdcbparam
```

5.2 Attributes Commands

The Attributes Command group manages the display of adapter, port, function, server attributes, and port statistics for each adapter specified. You can also set the port speed on LPe16202/OCe15100 adapters in NIC+FCoE mode.

5.2.1 HbaAttributes

This command shows a list of all adapter attributes for the adapter. The type of information listed might vary according to the adapter model.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
HbaAttributes <WWPN|MAC>
```

Parameters

WWPN The WWPN of an FC or FCoE function.

MAC The MAC address of a NIC function.

Example

```
HBA Attributes for 10:00:00:90:fa:94:26:af

Host Name:WIN-9ILVRDLR7JC
Manufacturer:Emulex Corporation
Serial Number:VA53900065
Model:LPe32002-M2
Model Desc:Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
Node WWN:20 00 00 90 fa 94 26 af
Node Symname:Emulex LPe32002-M2 FV11.1.38.61 DV11.0.247.0
HN:   WIN-9ILVRDLR7JC
OS:   Windows 2012 R2
HW Version:0000000C
FW Version:11.1.38.61
Vendor Spec ID:10DF
Number of Ports:1
Driver Name:elxfc
Driver Version:11.0.247.0
Device ID:E300
HBA Type:LPe32002-M2
Operational FW:11.1.38.61
IEEE Address:00 90 fa 94 26 af
Boot Code:Enabled
Boot Version:11.1.38.56
Board Temperature:Normal
Function Type:FC
Sub Device ID:E300
PCI Bus Number:32
PCI Func Number:0
Sub Vendor ID:10DF
IPL Filename:H62LEX1
Service Processor FW Name:11.1.38.61
ULP FW Name:11.1.38.61
FC Universal BIOS Version:11.1.38.56
FC x86 BIOS Version:11.1.38.56
FC EFI BIOS Version:11.1.38.48
FC FCODE Version:11.0.150.0
Flash Firmware Version:11.1.38.61
```

5.2.2 PortAttributes

This command shows a list of attributes for the adapter-specified function. The type of information listed might vary according to the adapter model.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
PortAttributes <WWPN|MAC>
```

Parameters

WWPN The WWPN of an FC or FCoE function.

MAC The MAC address of a NIC function.

Example

```
hbacmd h=10.192.32.197 portattributes 10:00:00:00:c9:88:88:89
Port Attributes for 10:00:00:00:c9:88:88:89
Node WWN: 20 00 00 00 c9 88 88 89
Port WWN: 10 00 00 00 c9 88 88 89
Port Symname:
Port FCID: 0000
Port Type: Unknown
Port State: Link Down
Port Service Type: 8
Port Supported FC4: 00 00 01 00 00 00 00 01
        00 00 00 00 00 00 00 00
        00 00 00 00 00 00 00 00
        00 00 00 00 00 00 00 00
Port Active FC          : 00 00 01 00 00 00 00 01
        00 00 00 00 00 00 00 00
        00 00 00 00 00 00 00 00
        00 00 00 00 00 00 00 00
Port Supported Speed: 4 8 16 Gbit/sec
Configured Port Speed: Auto Detect
Port Speed: Not Available
Max Frame Size: 2048
OS Device Name: \\.\Scsi5:
Num Discovered Ports: 0
Fabric Name: 00 00 00 00 00 00 00 00
Function Type: FC
FEC   : Disabled
```

5.2.3 PortStatistics

This command shows all function statistics for the specified function. The type of information listed may vary according to the adapter model.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
PortStatistics <WWPN>
```

Parameters

WWPN The WWPN of an FC or FCoE function.

5.2.4 ServerAttributes

This command shows a list of server attributes for the server where the specified function is running. The type of information listed may vary according to the adapter model.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
ServerAttributes <WWPN|MAC>
```

Parameters

WWPN	The WWPN of an FC or FCoE function.
MAC	The MAC address of a NIC function.

5.2.5 SetPhyPortSpeed

This command sets the port speed on LPe16202/OCe15100 adapters in NIC+FCoE mode.

LPe16202/OCe15100 adapters in NIC+FCoE mode have configurable physical port speeds. Depending on the port module or transceiver installed in the physical port, the speed settings can be forced to a specific value, for instance 1 Gb, or to a range of values for auto-negotiation with the switch; for example, 10 Gb/1 Gb/100 Mb. Three values can be configured: port speed mode, speed values, and the DAC cable length.

The configurable port speeds are based on the port module type and the mode defined by the port speed *Mode* parameter. For the default port speed mode, the speed setting is not required.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
SetPhyPortSpeed <WWPN|MAC> <Mode> [Speed [Length]]
```

Parameters

WWPN	The WWPN of an FCoE function.
MAC	The MAC address of a NIC function.
Mode	The Mode number: 0 = Default 1 = Auto-negotiate; requires the *Speed* parameter 2 = Force; requires the *Speed* and *Length* parameters If the adapter's port speed value and the switch's port speed value conflict, the link will not be brought up.

Speed The speed string of the PHY port. Some valid string values include `100Mb`, `1Gb`, and `10Gb`. The `PortAttributes` command lists all the valid speeds in Auto-negotiate and Force modes.

Auto-negotiated Speeds

A comma-separated list of available auto-negotiated speeds is displayed by the `PortAttributes` command for LPe16202/OCe15100 adapters in NIC+FCoE mode. For combinations of speeds, each speed is separated by a slash, for example, `10Gb/1Gb/100Mb`. If the port does not support auto-negotiated speeds, this property is displayed as `Not Supported`.

Forced Speeds

A comma-separated list of available forced speeds is displayed by the `PortAttributes` command for LPe16202/OCe15100 adapters in NIC+FCoE mode. Combinations of speeds for forced speeds are not available. If the port does not support forced speeds, this property is displayed as `Not Supported`.

If the *Mode* parameter is `1` or `2`, the *Speed* parameter is required. If the *Mode* parameter is `0`, the *Speed* parameter is ignored.

Length The length of the DAC cable in meters. Valid values are `0` to `10`. A length value of `0` indicates an optical cable. A Length value is required if you are using a 10Gb SFP and QSFP transceiver interface type.

If the *Mode* parameter is `0`, *Speed* and *Length* parameters are ignored, and if the `Mode` parameter is `1`, the `Length` parameter is ignored.

Examples

The following example configures the PHY port to a forced speed of 1 Gb/s with a cable length of 10 meters:

```
hbacmd setphyportspeed 00-00-c9-ad-ad-ac 2 1Gb 10
```

The following example tries to configure the PHY port to a forced speed of 100 Mb/s:

```
hbacmd setphyportspeed 00-00-c9-a9-41-88 2 100Mb
```

If the command is successful, the following is displayed:

```
Successfully changed speed settings on port.
```

If the *Mode* parameter is `2`, this command results in the following error because you must include a value for the *Length* parameter:

```
ERROR: <431>: Cable length required for force mode and interface type
```

5.2.6 SetPortEnabled

This command enables or disables a port. When a port is disabled, packets are not transmitted or received on the port.

NOTE Make sure that all I/O traffic on the port is stopped before disabling the port.

NOTE If the `SetPortEnabled` command disables an FC port, the adapter must be reset to activate the new value.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

`SetPortEnabled <WWPN|MAC> <PortEnable>`

Parameters

WWPN The WWPN of an FC or FCoE function on the port.

MAC The MAC address of a NIC function on the port.

PortEnable	The port-enabled state: 0 = Disabled 1 = Enabled

5.3 Authentication Commands

These commands configure a DHCHAP connection between an FC function and a switch port.

NOTE Authentication commands are supported only on LPe12000-series adapters.

5.3.1 AuthConfigList

This command returns the list of WWPNs that have an authentication connection configuration with the specified adapter.

Supported By

Linux, Solaris, and Windows

Syntax

`AuthConfigList <WWPN>`

Parameters

WWPN	The WWPN of an FC function.

5.3.2 DeleteAuthConfig

This command deletes the authentication configuration on the adapter.

Supported By

Linux, Solaris, and Windows

Syntax

`DeleteAuthConfig <WWPN1> <WWPN2> <PasswordType> <Password>`

Parameters

WWPN1	The WWPN of an FC function.
WWPN2	Either use `ff:ff:ff:ff:ff:ff:ff:ff` for a switch, or use the WWPN for a target.
PasswordType	1 = ASCII 2 = Hexadecimal (binary) 3 = Password not yet defined
Password	The current password value.

5.3.3 GetAuthConfig

This command retrieves the authentication configuration for the adapter.

Supported By

Linux, Solaris, and Windows

Syntax

`GetAuthConfig <WWPN1> <WWPN2>`

Parameters

WWPN1 The WWPN of an FC function.

WWPN2 Either use `ff:ff:ff:ff:ff:ff:ff:ff` for a switch, or use the WWPN for a target.

5.3.4 GetAuthStatus

This command returns the current status for the authentication connection specified by WWPN1 and WWPN2 (adapter and the switch). It includes the current authentication state (connected or failed). Currently authenticated connections specify the hash algorithm and DH group used in the DHCHAP associated with this connection. Failed status includes the failure reason.

Supported By

Linux, Solaris, and Windows

Syntax

`GetAuthStatus <WWPN1> <WWPN2>`

Parameters

WWPN1 The WWPN of an FC function.

WWPN2 Either use `ff:ff:ff:ff:ff:ff:ff:ff` for a switch, or use the WWPN for a target.

5.3.5 InitiateAuth

This command initiates the authentication configuration on the adapter.

Supported By

Linux, Solaris, and Windows

Syntax

`InitiateAuth <WWPN1> <WWPN2>`

Parameters

WWPN1 The WWPN of an FC function.

WWPN2 Either use `ff:ff:ff:ff:ff:ff:ff:ff` for a switch, or use the WWPN for a target.

5.3.6 SetAuthConfig

This command sets the authentication configuration for the adapter.

Supported By

Linux, Solaris, and Windows

Syntax

`SetAuthConfig <WWPN1> <WWPN2> <PasswordType> <Password> <Param> <Value>`

NOTE Where multiple parameters and values are used, separate them using commas.

Parameters

WWPN1 The WWPN of an FC function.

WWPN2 Either use `ff:ff:ff:ff:ff:ff:ff:ff` for a switch, or use the WWPN for a target.

PasswordType
`1` = ASCII
`2` = Hexadecimal (binary)
`3` = Password not yet defined

Password The current password value.

Param The parameter names:
- *Mode*
- *Timeout*
- *Bidirectional*
- *Hash-priority*
- *DH-priority*
- *Re-authentication*
- *Re-authentication-interval*

Value The value is based on the type of <*Param*>:
- *Mode*: `disabled`, `enabled`, or `passive`
- *Timeout*: time in seconds
- *Bi-directional*: `disabled` or `enabled`
- *Hash-priority*: `md5` or `sha1` (`md5` = first md5, then sha1; `sha1` = first sha1, then md5)
- *DH-priority*: `1`, `2`, `3`, `4`, `5`; any combination up to 5 digits
- *Re-authentication*: `disabled` or `enabled`
- *Re-authentication-interval*: `0`, `10` to `3600`, in seconds

5.3.7 SetPassword

This command sets the password on the adapter for an authenticated connection to the switch.

Supported By

Linux, Solaris, and Windows

Syntax

`SetPassword <WWPN1> <WWPN2> <Flag> <Cpt> <Cpw> <Npt> <Npw>`

Parameters

WWPN1 The WWPN of an FC function.

WWPN2	Either use `ff:ff:ff:ff:ff:ff:ff:ff` for a switch, or use the actual WWPN for a target.
Flag	`1` = Local (password used by the adapter when the adapter authenticates to the switch) `2` = Remote (password used by the adapter when the switch authenticates to the adapter)
Cpt	Current password type. `1` = ASCII `2` = Hexadecimal (binary) `3` = Password not yet defined
Cpw	Current password value.
Npt	New password type. `1` = ASCII `2` = Hexadecimal (binary)
Npw	New password value.

5.4 Boot Commands

The Boot Commands group manages the commands that enable or disable the network boot for NIC ports or the boot code for FC and FCoE adapter ports. You can also show and change FC and FCoE boot parameters.

CAUTION Using the EnableBootCode or SetBootParam commands on an older FC adapter (for example, LPe12000) that is being used to boot from SAN is not advisable. After the command has completed, the system performs an adapter reset, which may cause a loss of connectivity to the SAN and possible loss of data. To perform these commands on an older FC adapter, you must make sure that the adapter is not currently being used to boot from SAN.

Do one of the following:

- Move the target adapter to a non-boot from SAN host.
- If the host with the target adapter is also hosting other boot from SAN adapters, perform a boot from SAN using one of the other boot from SAN adapters. The target adapter can now be used.

5.4.1 EnableBootCode

This command performs either of the following:

- Enables or disables network boot for a NIC function (LPe16202/OCe15100 adapters only). If network boot is enabled, it is necessary to select the specific network boot type. The supported network boot type is PXE.
- Enables or disables the boot code for an FC function (all other LightPulse adapters). If the boot code is disabled on the FC function, the adapter does not boot from the SAN, regardless of the value for the `EnableBootFromSan` boot parameter. If the boot code is enabled on the FC function, the adapter boots from SAN if the `EnableBootFromSan` parameter is also enabled.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

`EnableBootCode <WWPN|MAC> <Flag> <NetworkBootMethod>`

Parameters

WWPN	The WWPN of an FC function.
MAC	The MAC address of a NIC function. For an LPe16202/OCe15100 adapter in NIC+FCoE mode, you must specify the MAC address.
Flag	`D` = Disable the boot code. `E` = Enable the boot code.
NetworkBootMethod	Network boot method to be used by the NIC (PXE). The acceptable values are in the Available Network Boot Methods string displayed by the `HbaAttributes` command. This argument is not required for FC ports, and must only be specified when enabling network boot.

Examples

The following example enables boot code for an FC adapter:

```
hbacmd EnableBootCode 10:00:00:00:c9:20:20:20 e
```

The following example disables network boot on an LPe16202/OCe15100 adapter in NIC+FCoE mode:

```
hbacmd EnableBootCode 00-00-c9-11-22-33 d
```

The following example enables network boot on an LPe16202/OCe15100 adapter in NIC+FCoE mode:

```
hbacmd EnableBootCode 00-00-c9-11-22-33 e PXE
```

5.4.2 GetBootParams

This command shows the FC or FCoE boot parameters. If any arguments are missing or invalid, an error is reported. If all arguments are correct, the data is displayed in tabular form.

Supported By

Linux, Solaris, and Windows

Syntax

`GetBootParams <`*`WWPN`*`> <`*`Type`*`>`

Parameters

WWPN	The WWPN of an FC function.
Type	`X86`, `OpenBoot`, or `UEFI`.

5.4.3 SetBootParam

This command changes the FC and FCoE boot parameters. You can change function parameters and boot device parameters for x86, OpenBoot, and EFI boot.

- If you change adapter parameters, omit the `BootDev` keyword and value; otherwise, an error is reported.
- If you change boot device parameters for OpenBoot, omit the `BootDev` keyword and value; otherwise, an error is reported.
- For boot device parameters for x86 or EFI, you must provide the `BootDev` keyword and value.

Supported By

Linux, Solaris, and Windows

Syntax

`SetBootParam <WWPN> <Type> <AdapterParam> <Value1> [BootDev <Value2>]`

Parameters

WWPN	The WWPN of an FC or FCoE port.
Type	{x86, EFI, OB}
AdapterParam	The parameter name.
Value1	The parameter value.
BootDev	The boot device.
Value2	The boot device entry number: {0 to 7}.

Adapter Parameters	Boot Type	Value
DefaultAlpa	All	{ Value }
EnableAdapterBoot	All	{ State } (0=Disable, 1=Enable)
EnableBootFromSan	All	{ State } (0=Disable, 1=Enable)
LinkSpeed	All	{ 0, 1, 2, 4, 8 } This parameter is available only for the following: LPe12000-series adapters All adapters on Windows + CIM Provider on a VMware ESXi 5.5 host Use the SetLinkSpeed command instead for all other adapters and operating systems.
PlogiRetryTimer	All	{ 0, 1, 2, 3 }
Topology	All	{ 0, 1, 2, 3 }
AutoScan	x86	{ 0, 1, 2, 3 }
AutoBootSectorEnable	x86	{ State } (0=Disable, 1=Enable)
EDD30Enable	x86	{ State } (0=Disable, 1=Enable)
EnvVarEnable	x86	{ State } (0=Disable, 1=Enable)
SpinupDelayEnable	x86	{ State } (0=Disable, 1=Enable)
StartUnitCommandEnable	x86	{ State } (0=Disable, 1=Enable)
BootTargetScan	EFI	{ 0, 1, 2 }
DevicePathSelection	EFI	{ 0, 1 }
MaxLunsPerTarget	EFI	{ Value }
ResetDelayTimer	EFI	{ Value }
SfsFlag	OB	{ State } (0=Disable, 1=Enable)
Boot Device Parameters		
D_ID	All	{ Value [*BootDev* <*Value2*>] }
LUN	All	{ Value [*BootDev* <*Value2*>] }
TargetWwpn	All	{ Value [*BootDev* <*Value2*>] }
TargetID	OB	{ Value }

5.4.4 SetLinkSpeed

This command sets the link speed for a specific port on an LPe16000-series, LPe3100-series, or LPe32000-series adapter.

NOTE This command is not supported on Windows + CIM Provider on a VMware ESXi 5.5 host on any adapter. For ESXi 5.5, and for LPe12000-series adapters, use the `LinkSpeed` parameter of the `SetBootParam` command instead.

NOTE This command is not supported on LPe16202/OCe15100 adapters in NIC+FCoE mode. Use the `SetPhyPortSpeed` command instead.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host (except ESXi 5.5)

Syntax

`SetPortSpeed <WWPN> <LinkSpeed>`

Parameters

WWPN	The WWPN of an FC or FCoE port.
LinkSpeed	Numeric value representing a supported link speed. For a list of port speeds supported by the adapter, use the `PortAttributes` command to display `Port Supported Speed`. Specify a value of `0` to configure Auto Detect mode.

NOTE A port reset is required to activate the new settings.

5.5 DCB Commands

This command group controls the DCB and LLDP parameters for FCoE and NIC adapter ports on LPe16202/OCe15100 adapters in NIC+FCoE mode.

5.5.1 GetDCBParams

This command shows the active and configured DCB and LLDP settings on a port of an LPe16202/OCe15100 adapter. The active parameters show what the adapter port is currently running, and the configured parameters show the value to which the adapter port's DCB parameter is set.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

`GetDCBParams <WWPN|MAC>`

Parameters

WWPN	The WWPN of an FCoE function on the port.
MAC	The MAC address of a NIC function on the port.

Example

```
hbacmd GetDCBParams 00-00-c9-93-2f-d8
```

5.5.2 GetPGInfo

This command shows the ETS priority group bandwidth percentages for the port of an LPe16202/OCe15100 adapter in NIC+FCoE mode. Additionally, this command displays the number of priority groups supported by an adapter.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

`GetPGInfo <WWPN|MAC>`

Parameters

WWPN	The WWPN address of an FCoE function on the port.
MAC	The MAC address of a NIC function on the port.

Example

```
hbacmd getpginfo 00-00-c9-93-2f-d8
```

5.5.3 SetCnaPGBW

This command sets the ETS priority group bandwidth percentages on a port of an LPe16202/OCe15100 adapter in NIC+FCoE mode according to the following rules:

- Bandwidths (`BW0–BW7`) for priority groups 0 to 7 (`PG0` to `PG7`) must total 100 (for 100 percent).
- Bandwidth can be assigned to a priority group that has priorities.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

`SetCnaPGBW <WWPN|MAC> <BW0–BW7>`

Parameters

WWPN	The WWPN of an FCoE function on the port.
MAC	The MAC address of a NIC function on the port.
BW0–BW7	The bandwidths allocated for the priority groups `0` to `7`.

Example

This command sets the bandwidth of PG0 to 50%, PG1 to 50%, and PG2 to PG7 to 0%.

```
hbacmd SetCnaPGBW 10:00:00:00:c9:3c:f7:88 50 50 0 0 0 0 0 0
```

5.5.4 SetDCBParam

This command configures the DCB and LLDP settings on an LPe16202/OCe15100 adapter in NIC+FCoE mode. Use the `GetDCBParams` command to obtain valid parameter names for use in this command.

NOTE You cannot set DCBX mode. If you attempt to specify a *dcbxmode* parameter, an error message is displayed.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
SetDCBParam <WWPN|MAC> <Param> <Value>
```

Parameters

WWPN	The WWPN of an FCoE function on the port.
MAC	The MAC address of a NIC function on the port.
Param	The parameter name. See DCB Settings for <Param> and <Value> and LLDP Settings for <Param> and <Value>.
Value	The parameter value. See DCB Settings for <Param> and <Value> and LLDP Settings for <Param> and <Value>.

DCB Settings for <*Param*> and <*Value*>

<*Param*>	**Description and <*Value*>**
DCBXState	The DCBX protocol state. 0 = Disabled 1 = Enabled
PFCEnable	Flow control in both directions (transmit and receive). 0 = Disabled 1 = Enabled
FCoEPriority	This parameter is only applicable for ports running FCoE. A single priority must be specified. The range of valid values is 0 to 7. Only one priority can be specified for each invocation of this command and must be for a protocol running on the port. If more than one protocol priority can be set, they must be unique values.
PFCPriority	A list of comma-separated values where multiple PFC priorities are supported. The comma-separated list can contain up to seven values ranging from 0 to 7.
defaults	Use to set the DCB parameters (including priority groups) to their default values. For example: `hbacmd SetDCBParam <WWPN|MAC> defaults`

LLDP Settings for <*Param*> and <*Value*>

<*Param*>	**Description and <*Value*>**
TxState	Transmit State: DCBX uses LLDP to exchange parameters between two link peers. For the DCBX protocol to operate correctly, both LLDP Rx and Tx must be enabled. If either Rx or Tx is disabled, DCBX is disabled. 0 = Disabled 1 = Enabled
RxState	Receive State: DCBX uses LLDP to exchange parameters between two link peers. For the DCBX protocol to operate correctly, both LLDP Rx and Tx must be enabled. If either Rx or Tx is disabled, DCBX is disabled. 0 = Disabled 1 = Enabled
TxPortDesc	Transmit Port Description: Provides a description of the port in an alphanumeric format. 0 = Disabled 1 = Enabled
TxSysDesc	Transmit System Description: Provides a description of the network entity in an alphanumeric format. 0 = Disabled 1 = Enabled
TxSysName	Transmit System Name: Provides the system's assigned name in an alphanumeric format. 0 = Disabled 1 = Enabled
TxSysCap	Transmit System Capabilities: 0 = Disabled 1 = Enabled

Example

```
hbacmd SetDCBParam 00-00-c9-3c-f7-88 fcoepriority 3
```

5.5.5 SetDCBPriority

This command sets the PFC priorities and the ETS priority groups, priorities. The values must be set according to the following rules:

- The priorities range from 0 to 7.
- A priority (0 to 7) must exist in only one priority group.
- All priorities must appear once in any of the eight (PG0 to PG7) priority groups, or if available, PG15.
- To not specify priorities for a priority group, use a dash (–).
- Any assigned PFC priority must be assigned as the single priority in a priority group (for example, no other priorities allowed in a group assigned the PFC priority).
- Any PG assigned one or more priorities must also be assigned a nonzero bandwidth value (see Section 5.5.3, SetCnaPGBW).

The following rules are specific to LPe16202/OCe15100 in NIC+FCoE mode:

- A maximum of two PFC priorities can be assigned.
- If FCoE is running on the port, one of the PFC priorities must match the FCoE priority.

The following rules are specific to NIC-only adapters:

- Only one PFC priority can be assigned.
- In NIC-Only mode, PFC is disabled by default. To enable PFC, NIC ETS must be enabled.

  To enable NIC ETS:

  — In Windows, enable `Enhanced Transmission Selection` in the driver properties (for example, in the Device Manager property page for the NIC driver).

  — In Linux, load the NIC driver with the `tx_prio` driver parameter set to `1`.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

`SetDcbPriority <`*`WWPN`*`|`*`MAC`*`> <`*`PFC`*`> <`*`PG0`*`> <`*`PG1`*`> <`*`PG2`*`> <`*`PG3`*`> <`*`PG4`*`> <`*`PG5`*`> <`*`PG6`*`> <`*`PG7`*`> [`*`PG15`*`]`

Parameters

WWPN	The WWPN of an FCoE function on a port.
MAC	The MAC address of a NIC function on a port.
PFC	The PFC priority that is a comma-separated list of up to eight values, ranging from `0` to `7`.
PG0–PG7	Priority group membership that is a comma-separated list of priorities ranging from `0` to `7`. Each set of priorities for a group must be separated by a space. All priorities (`0` to `7`) must be assigned to a PGID.

Example

```
hbacmd SetDCBPriority 10:00:00:00:c9:3c:f7:88 3 0,1,2,4,5,6,7 3 0 0 0 0 0 0
```

5.6 Diagnostic Commands

The Diagnostic Commands group provides commands that enable you to detect cabling problems, examine transceiver data, and to flash memory load lists. Additionally, you can run specific diagnostic tests, such as the Loopback test and POST test.

CAUTION Using the `LoopBackTest`, `PciData`, or `Post` commands on an older FC adapter (for example, LPe12000) that is being used to boot from SAN is not advisable. After the command has completed, the system performs an adapter reset, which may cause a loss of connectivity to the SAN and possible loss of data. To perform these commands on an older FC adapter, you must make sure that the adapter is not currently being used to boot from SAN.

Do one of the following:

- Move the target adapter to a non-boot from SAN host.
- If the host with the target adapter is also hosting other boot from SAN adapters, perform a boot from SAN using one of the other boot from SAN adapters. The target adapter can now be used.

5.6.1 D_PortTest

D_Port, also called ClearLink, is a set of diagnostic tests that detects physical cabling issues that result in increased error rates and intermittent behavior.

This command is only supported for LPe16000-series, LPe31000-series, and LPe32000-series FC adapters connected to D_Port-enabled Brocade switches.

NOTE Do not enable D_Port on the switch port.

The `DPortTest` command runs a series of tests, including local electrical loopback, loopback to the remote optics, loopback from the remote port to the local optics, and a full-device loopback test with data integrity checks. It also provides an estimate of cable length, from the switch to the adapter, to validate that a proper buffering scheme is in place.

The various tests allow some fault isolation, so you can distinguish faults that are the result of marginal cables, optics modules, and connectors or optics seating. If the adapter, firmware, SFP, or switch do not support D_Port testing, an error is generated.

These tests run with the physical connection in an offline diagnostic state, so normal I/O cannot be sent through the physical port while the test is in progress. While the port is in D_Port mode, the link will appear down on that port; similar to an unplugged cable.

NOTE If you are using D_Port in a boot from SAN configuration, the configuration must have redundant paths to the boot LUN, and only one of the redundant adapter ports should be set to D_Port.

Supported By

Linux, Windows, and Windows + CIM Provider on a VMware host

Syntax

`DPortTest <WWPN>`

Parameters

WWPN The WWPN of the FC function on the adapter.

Considerations when using D_Port

- A test failure occurs if the `DPortTest` command is run with a switch that does not support D_Port testing.
- Typing **CTRL+C** while the D_Port tests are running terminates the tests and the completed results are displayed.
- If the Overall Test Result is FAILED, you must either rerun the tests successfully, or reset the adapter port to bring up the link.
- If a test phase fails, the D_Port diagnostics are stopped automatically. As a result, test phases that would have occurred after the failure are not displayed.
- If more than one error is reported in a single test phase, multiple lines are displayed showing each error.

Examples

Successful test and test failure examples are below.

Successful Test

```
>hbacmd DPortTest 10:00:00:00:c9:d1:a2:d0

Running D_Port Tests. Please wait. Polling for results.........................

D_Port Test Status:     Passed
Buffers Required:       1
Frame Size:             2112
Round Trip Latency:     1898 nanoseconds
Estimated Cable Length: 172  meters

================================================================================
Test Phase                Result    Latency   Local Errors   Remote Errors
================================================================================
Electrical Loopback       Passed    122
Optical Loopback          Passed    1898
Reverse Optical Loopback  Skipped   0
Link Traffic              Passed    0
```

Test Failures

```
>hbacmd DPortTest 10:00:00:00:c9:d1:a2:d0

Running D_Port Tests. Please wait. Polling for results.......................

D_Port Test Status:     Failed
Buffers Required:       0
Frame Size:             0
Round Trip Latency:     0    nanoseconds
Estimated Cable Length: 0    meters

================================================================================
Test Phase                Result    Latency   Local Errors   Remote Errors
================================================================================
Electrical Loopback       Failed    n/a

>hbacmd DPortTest 10:00:00:00:c9:d1:a2:d0
```

```
Running D_Port Tests. Please wait. Polling for results.........................

D_Port Test Status:     Failed
Buffers Required:       0
Frame Size:             0
Round Trip Latency:     0    nanoseconds
Estimated Cable Length: 0    meters

================================================================================
Test Phase                Result    Latency   Local Errors   Remote Errors
================================================================================
Electrical Loopback       Passed    0
Optical Loopback          Failed    n/a

>hbacmd DPortTest 10:00:00:00:c9:d1:a2:d0

Running D_Port Tests. Please wait. Polling for results.........................

D_Port Test Status:     Failed
Buffers Required:       1
Frame Size:             2112
Round Trip Latency:     1898 nanoseconds
Estimated Cable Length: 172  meters

================================================================================
Test Phase                Result    Latency   Local Errors   Remote Errors
================================================================================
Electrical Loopback       Passed    127
Optical Loopback          Passed    1898
Reverse Optical Loopback  Skipped   0
Link Traffic              Failed    n/a
```

5.6.2 EchoTest

This command runs the Echo test on FC functions. It is supported only on LPe12000-series adapters.

NOTE The `EchoTest` command fails if the target WWPN does not support the `ECHO ELS` command.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

`EchoTest <WWPN Source> <WWPN Destination> <Count> <StopOnError> <Pattern>`

Parameters

WWPN Source	The WWPN of the originating FC function.
WWPN Destination	The WWPN of the destination (echoing) FC functions.
Count	The number of times to run the test. Use `0` to run the test indefinitely.

StopOnError	Checks if the test must be halted on error: 0 = No halt 1 = Halt on error
Pattern	Hexadecimal data pattern to transmit (up to 8 characters).

5.6.3 GetBeacon

This command shows the current beacon state (either on or off).

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
GetBeacon <WWPN|MAC>
```

Parameters

WWPN	The WWPN of the FC or FCoE function on the port.
MAC	The MAC address of the NIC function on the port.

5.6.4 GetXcvrData

This command shows transceiver data for a port on an adapter.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
GetXcvrData <WWPN|MAC> [Type]
```

Parameters

WWPN	The WWPN of an FC or FCoE function on the port.
MAC	The MAC address of a NIC function on the port.
Type	The type of SFP data to display: 1 = Formatted SFS data (default) 2 = Raw SFS data (not supported by Windows + CIM Provider on a VMware host)

Example

```
hbacmd GetXcvrData 00-00-c9-93-2f-d6
```

5.6.5 LoadList

This command shows the flash memory load list data for the FC function on the adapter. It is supported only on LPe12000-series adapters.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
LoadList <WWPN>
```

Parameters

WWPN The WWPN of the FC function on the adapter.

5.6.6 LoopBackTest

This command runs one of the loopback tests available on the adapter port specified by the WWPN or MAC address.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

`LoopBackTest <WWPN|MAC> <Type> <Count> <StopOnError> [Pattern]`

Parameters

WWPN	The WWPN of an FC or FCoE function on the port.
MAC	The MAC address of a NIC function on the port.
Type	The type of Loopback test to run: 0 = PCI Loopback test; not supported on LPe16202/OCe15100 adapters 1 = Internal Loopback test; not supported on LPe16202/OCe15100 adapters 2 = External Loopback test (requires loopback plug) 4 = PHY Loopback test; supported only on LPe16202/OCe15100 adapters in NIC+FCoE mode 5 = MAC Loopback test; supported only on LPe16202/OCe15100 adapters in NIC+FCoE mode
Count	Number of times to run the test. Possible values are 1 to 99999. To run the test infinitely, use 0.
StopOnError	Checks if the test must be halted on error. 0 = No halt 1 = Halt
Pattern	An optional parameter that specifies 1–8 hexadecimal bytes to use for loopback data (for example, 1a2b3c4d).

Example

```
hbacmd LoopBackTest 10:00:00:00:c9:20:20:20 1 120 0
```

5.6.7 LoopMap

This command shows the arbitrated loop map data on an FC function.

Supported By

Linux, Solaris, and Windows

Syntax

`LoopMap <WWPN>`

Parameters

WWPN The WWPN of the FC function.

5.6.8 PciData

This command shows the PCI configuration data (if available).

The PCI registers displayed are specific to the function referenced in the OneCommand Manager CLI. For example, if you specify the WWPN for the FCoE function, the PCI registers for that FCoE function are returned. If you specify the MAC address for the NIC function on that same physical port, the PCI registers for that NIC function are returned.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

`PciData <WWPN|MAC>`

Parameters

WWPN The WWPN of an FC or FCoE function.

MAC The MAC address of a NIC function.

Example

```
hbacmd pcidata 10:00:B0:5A:DA:01:B1:0D
```

The example output:

```
Vendor ID:            0x10DF       Device ID:            0xE300
Command:              0x0546       Status:               0x0010
Revision ID:          0x01         Prog If:              0x00
Subclass:             0x04         Base Class:           0x0C
Cache Line Size:      0x10         Latency Timer:        0x00
Header Type:          0x80         Built In Self Test:   0x00
Base Address 0:       0x92C0800C   Base Address 1:       0x00000000
Base Address 2:       0x00000000   Base Address 3:       0x00000000
Base Address 4:       0x00000000   Base Address 5:       0x00000000
CIS:                  0x00000000   SubVendor ID:         0x1590
SubSystem ID:         0x0214       ROM Base Address:     0x00000000
Interrupt Line:       0x00         Interrupt Pin:        0x02
Minimum Grant:        0x00         Maximum Latency:      0x00
Capabilities Ptr:     0x54

FeatureEnable:                     0x00000000
PwrMgt Caps/Nxt/ID:                0x00036001
PwrMgt Ctl/Stat:                   0x00000008
MSI Ctl/Nxt/ID:                    0x018A7805
MSI Lo Address:                    0x00000000
MSI High Address:                  0x00000000
MSI Data:                          0x00000000
MSI Mask Bits:                     0x00000000
MSI Pending Bits:                  0x00000000
MSI-X Ctl/Nxt/ID:                  0x81FF9411
MSI-X Table Offset:                0x00004000
MSI-X PBA Offset:                  0x00003400
VPD Address/Nxt/ID:                0x00000003
VPD Data:                          0x51000D820000000000000000
PCIE Capabilities/Nxt/ID:          0x0002F810
PCIE Device Cap. Reg:              0x10008724
PCIE Device Status & Control:      0x00095136
PCIE Link Capabilities:            0x0041DC83
PCIE Link Status & Control:        0x10830040
Slot Capabilities Register:        0x00000000
```

```
Slot Status & Control Register:   0x00000000
Root Capabilities & Ctl Register: 0x00000000
Root Status Register:             0x00000000
Device Capabilities 2 Register:   0x0010001F
Device StatusControl 2 Register:  0x00000000
Link Capabilities 2 Register:     0x0000000E
Link Status 2 & Ctl 2 Register:   0x00000000
Slot Capabilities 2 Register:     0x00000000
Slot Status Control 2 Register:   0x00000000
Enhanced Cap Header AER:          0x00000000
Uncorrectable Error Status:       0x00000000
Uncorrectable Error Mask:         0x00000000
Uncorrectable Error Severity:     0x00000000
Correctable Error Status:         0x00000000
Correctable Error Mask:           0x00000000
Adv. Error Cap & Control:         0x00000000
Header Log 0x0:                   0x00000000
Header Log 0x4:                   0x00000000
Header Log 0x8:                   0x00000000
Header Log 0xC:                   0x00000000
Enhanced Cap Header PBUDG:        0x00000000
Data Select Reg:                  0x00000000
Data Register:                    0x00000000
Power Budget:                     0x00000000
```

5.6.9 PostTest

This command runs the POST on the adapter.

NOTE This command is supported only on LPe12000-series adapters.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

`PostTest <WWPN>`

Parameters

WWPN The WWPN of the FC port.

5.6.10 SetBeacon

This command turns the beacon on or off on the adapter port.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

`SetBeacon <WWPN|MAC> <BeaconState>`

Parameters

WWPN	The WWPN of an FC or FCoE function on the port.
MAC	The MAC address of a NIC function on the port.
BeaconState	Indicates the state of the beacon. 0 = Off 1 = On

5.6.11 Wakeup

This command shows the firmware's wakeup parameters for the FC function on the adapter.

NOTE This command is supported only on LPe12000-series adapters.

Supported By

Linux, Solaris, and Windows

Syntax

`Wakeup <WWPN>`

Parameters

WWPN	The WWPN of an FC function.

5.7 Driver Parameter Commands

The Driver Parameter Commands group controls the driver parameters. You also can change the parameters back to factory default values.

NOTE Driver Parameter commands are supported only for FC and FCoE ports.

The `DriverConfig` and `SetDriverParamDefaults` commands are not supported for Solaris.

Considerations

- Driver parameters set to temporary or global values (using the `T` and `G` flags, respectively) must be read using the GetDriverParams command to view the current value of the parameter. The `GetDriverParamsGlobal` command returns only permanently set driver parameter values.

  Additionally, if temporary and global values are set for one or more driver parameters, the `SaveConfig` command must be run with the `N` flag (using the `N` flag is analogous to using the `GetDriverParams` command) to force the driver parameter values for the specified adapter to be saved. Inaccurate values can be saved if the `G` flag is used for this command.
- The list of available driver parameters that can be configured are different depending on the operating system and protocol (FC or FCoE).

5.7.1 DriverConfig

This command sets all driver parameters to the values in the `.dpv` file type. The `.dpv` file's driver type must match the driver type of the host operating system adapter.

Supported By

Linux and Windows

Syntax

`DriverConfig <WWPN> <FileName> <Flag>`

Parameters

WWPN	The WWPN of an FC or FCoE function.
FileName	The name of the `.dpv` file, which is stored in the Emulex Repository directory.
Flag	`G` = Make the change global (all FC or FCoE functions on this host). `N` = Make the change non-global (function-specific).

NOTE On Windows hosts, depending on the `.dpv` file specified, the settings are applied to either the FC or FCoE functions, but not both at the same time, because FC and FCoE have separate drivers. The `.dpv` file has settings for only one of the driver types.

5.7.2 GetDriverParams

This command shows the name and value of each parameter.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

`GetDriverParams <WWPN>`

Parameters

WWPN	The WWPN of an FC or FCoE function.

NOTE On Windows, if an FC function is specified, the global parameters for the FC driver are displayed. If an FCoE function is specified, the global parameters for the FCoE driver are displayed.

5.7.3 GetDriverParamsGlobal

This command shows the name and the global value of each driver parameter.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

`GetDriverParamsGlobal <WWPN>`

Parameters

WWPN	The WWPN of an FC or FCoE function.

NOTE On Windows, if an FC function is specified, the global parameters for the FC driver are displayed. If an FCoE function is specified, the global parameters for the FCoE driver are displayed.

5.7.4 SaveConfig

This command saves the specified adapter's driver parameters to a file. The resulting file contains a list of driver parameter definitions in ASCII file format with definitions delimited by a comma. Each definition has the following syntax:

<parameter-name>=<parameter-value>

The command saves either the values of the global set, or those specific to the adapter in the Emulex Repository directory.

Supported By

Linux, Solaris, and Windows

Syntax

`SaveConfig <WWPN> <FileName> <Flag>`

Parameters

WWPN	The WWPN of an FC or FCoE function.
FileName	The name of the file that contains the driver parameters list.
Flag	G = Save the global parameter set. N = Save the local (function-specific) parameter set.

NOTE On Windows hosts, depending on the WWPN specified, the settings are saved for either the FC or FCoE functions, but not both at the same time, because FC and FCoE have separate drivers. The `.dpv` file will be saved with the settings for only one of the driver types.

5.7.5 SetDriverParam

This command changes a driver parameter and designates the scope of the change.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

`SetDriverParam <WWPN> <Flag1> <Flag2> <Param> <Value>`

Parameters

WWPN	The WWPN of an FC or FCoE function.
Flag1	L = Make the change local for only this function. G = Make the change global (all FC or FCoE functions on this host).
Flag2	P = Make the change permanent (persists across reboot). For Linux, to make a permanent change that persists across reboots, you must set `Flag1` to G (Global). T = Make the change temporary.
Param	Name of the parameter to modify.
Value	New parameter value, decimal or hexadecimal (`0xNNN`).

NOTE On Windows hosts, if the G (global) flag is specified, the changes depend on the WWPN specified. The settings are applied to either the FC or FCoE functions, but not both at the same time, because FC and FCoE have separate drivers.

Example

To enable dynamic target mode:

```
hbacmd SetDriverParam 10:00:00:00:c9:ff:ff:ff L P enable-dtm 1
```

To disable dynamic target mode, set the flags to `0`.

5.7.6 SetDriverParamDefaults

This command changes all values to the default for the adapter.

Supported By

Linux and Windows

Syntax

`SetDriverParamDefaults <WWPN> <Flag1> <Flag2>`

Parameters

WWPN	The WWPN of an FC or FCoE function.
Flag1	`L` = Make the change local for only this function. `G` = Make the change global (applies to all FC or FCoE functions on this host).
Flag2	`P` = Make the change permanent (the change persists across reboot). `T` = Make the change temporary.

NOTE On Windows hosts, if the `G` (global) flag is specified, the changes depend on the WWPN specified. The defaults are applied to either the FC or FCoE functions, but not both at the same time, because FC and FCoE have separate drivers.

5.8 Dump Commands

The Diagnostic Dump feature enables you to create a dump file for a selected adapter. Dump files contain information, such as firmware version, driver version, and operating system information. This information is useful when you are troubleshooting an adapter, but it is unavailable in Read-Only mode.

CAUTION Disruption of service can occur if a diagnostic dump is run during I/O activity.

The dump files created are text files (`.txt` extension) and binary files. The extension for binary files depends on the following adapter types:

- LPe16202/OCe15100 adapters (Enhanced FAT Dump) – `.edf` extension
- LPe16000-series (except LPe16202/OCe15100 adapters), LPe31000-series, and LPe32000-series adapters – `.bin` extension
- LPe12000-series adapters – `.dmp` extension

5.8.1 DeleteDumpFiles

This command deletes all diagnostic dump files for an adapter.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

`DeleteDumpFiles <WWPN|MAC>`

Parameters

WWPN	The WWPN of an FC or FCoE function on the adapter.
MAC	The MAC address of a NIC port function on the adapter.

5.8.2 Dump

This command creates a diagnostic dump file in the `HbaCmd` dump file directory.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

`Dump <WWPN|MAC>`

Parameters

WWPN	The WWPN of an FC or FCoE port.
MAC	The MAC address of a NIC port.

5.8.3 GetDumpDirectory

This command shows the dump file directory for the adapters in the host.

NOTE The dump directory can be set only on VMware ESXi hosts.
The dump directory applies to all adapters in the server. A separate dump directory for each adapter does not exist.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

`GetDumpDirectory [WWPN|MAC]`

Parameters

WWPN	Obsolete; ignored if specified.
MAC	Obsolete; ignored if specified.

5.8.4 GetDumpFile

This command gets the user-specified dump file to the local client's dump directory. The dump directory (local and remote) is named `Dump`. The dump files are copied from the dump directory of the remote host to the dump directory of the local host. Therefore, if the remote host option is not specified (`h=IP_Address[:port]`), this command returns an error because the source and destination directories are the same.

Dump directory:

- Windows – *SystemDrive_Letter*:\Program Files\Emulex\Util\Dump
- Linux – /var/opt/emulex/ocmanager/Dump
- Solaris – /opt/ELXocm/Dump
- VMware ESXi – The dump directory set using the SetDumpDirectory command.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

GetDumpFile <h=*IP_Address*[:*port*]>[*WWPN*|*MAC*] <*filename*>

Parameters

WWPN	Obsolete; ignored if specified.
MAC	Obsolete; ignored if specified.
filename	The name of the dump file to be copied from the remote host.

Example

```
hbacmd h=10.192.193.154  GetDumpFile
BG-HBANYWARE-15_10000000c97d1314_20100120-032820421.dmp
```

5.8.5 GetDumpFileNames

This command gets the names of the files in the host's dump directory.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

GetDumpFileNames [*WWPN*|*MAC*]

Parameters

WWPN	Obsolete; ignored if specified.
MAC	Obsolete; ignored if specified.

Example

```
hbacmd  GetDumpFileNames
```

5.8.6 GetRetentionCount

This command shows the maximum number of diagnostic dump files to keep.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

GetRetentionCount [*WWPN*|*MAC*] <*value*>

Parameters

WWPN	Obsolete; ignored if specified.
MAC	Obsolete; ignored if specified.
Value	The number of dump files to retain.

Example

```
hbacmd getRetentionCount 6
```

5.8.7 SetDumpDirectory

This command sets the dump directory (valid only on VMware ESXi hosts).

Supported By

Windows + CIM Provider on a VMware host

To use the `SetDumpDirectory` command, you must have a directory (which must be a storage partition) mapped under `/vmfs/volumes` where the files are dumped. This directory points to the internal hard disk or an external storage area and can also be mapped using the vSphere Client utility from VMware.

The application checks for the dump directory and creates the dump files in that location.

In a remote environment, you can use the `SetDumpDirectory` command from a host running any operating system (including Linux, Solaris, and Windows), but only to a remote host that is running VMware ESXi.

NOTE The dump directory applies to all adapters in the server. A separate dump directory for each adapter does not exist.

Syntax

```
SetDumpDirectory <DumpDirectoryName>
```

Parameters

DumpDirectoryName	The directory under `/vmfs/volumes` that you created to store the dump files.

Example

This example shows the dump directory set to `/vmfs/volumes/ocm-datastore`:

```
hbacmd h=10.192.203.173 m=cim u=root p=Swamiji001 n=root/emulex SetDumpDirectory
10:00:00:00:c9:61:f2:64 ocm-datastore
```

5.8.8 SetRetentionCount

This command specifies the maximum number of diagnostic dump files for the adapter. When the count reaches the limit, the next dump operation deletes the oldest file.

NOTE The retention count applies to all adapters in the server.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
SetRetentionCount [WWPN|MAC] <Value>
```

Parameters

WWPN Obsolete; ignored if specified.

MAC Obsolete; ignored if specified.

Value The number of dump files to retain.

Example

```
hbacmd  SetRetentionCount 6
```

5.9 FCoE Commands

The FCoE Commands group manages the FIP parameters and displays the FCF for an FCoE function.

NOTE These commands are supported only on LPe16202/OCe15100 adapters in NIC+FCoE mode.

5.9.1 GetFCFInfo

This command shows the FCF information of the FCoE function.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
GetFCFInfo <WWPN>
```

Parameters

WWPN The WWPN of an FCoE function.

Example

```
hbacmd GetFCFInfo 10:00:00:00:c9:3c:f7:88
Number of FCFs:   1
Active FCFs:      1
Entry 0:
      State:            1
      Priority:         133
      Fabric Name:      10:00:00:05:1E:0C:54:49
      Switch Name:      10:00:00:05:1E:0C:54:49
      MAC:              00:05:9B:71:3D:71
      FC Map:           0x0EFC00
      VLAN IDs:
      LKA Period:       8
```

5.9.2 GetFIPParams

This command gets the FIP parameters of an FCoE function.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

`GetFIPParams <WWPN>`

Parameters

WWPN The WWPN of an FCoE function.

Example

```
hbacmd h=10.231.140.83 getfipparams 10:00:00:00:c9:bc:a9:31
Param Description       Param Name  Value
-----------------       ----------  -----
Primary Fabric Name     pfabric     FF:FF:FF:FF:FF:FF:FF:FF
Primary Switch Name     pswitch     FF:FF:FF:FF:FF:FF:FF:FF
DCB VLan ID             vlanid      Any VLAN ID is valid
```

5.9.3 SetFIPParam

This command sets the FIP parameters of an FCoE function.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

`SetFIPParam <WWPN> <Param> <Value>`

Parameters

WWPN The WWPN of an FCoE function.

Param The FIP parameter name:
- pfabric
- pswitch
- vlanid
- fcmap
- cinvlanid

Value The value based on the FIP parameter name:
- pfabric: 8-byte fabric name (format XX:XX:XX:XX:XX:XX:XX:XX)
- pswitch: 8-byte switch name (format XX:XX:XX:XX:XX:XX:XX:XX)
- vlanid: 2-byte VLAN ID [0–4095] or *any* for any VLANID
- fcmap: 3-byte FC_map, 0x0EFCxx
- cinvlanid: 2-byte VLAN_ID [0–4095]

Example

```
hbacmd SetFIPParam 10:00:00:00:c9:5b:3a:6d fcmap 0x0efc99
```

5.10 Firmware Commands

These commands enables you to view and configure firmware parameters.

5.10.1 getfwparams

This command displays a list of all configurable firmware parameters.

Supported By

Windows and Linux

Syntax

`getfwparams <WWPN>`

Parameters

WWPN The WWPN of the adapter port.

Example

```
hbacmd getfwparams 10:00:00:90:fa:94:2e:ca
FW Params for 10:00:00:90:fa:94:2e:ca. Values in HEX format.
DX      Param       Low      High       Def       Cur          Dyn
00:     FA-PWWN       0         1         0         1            5
01      FEC           0         1         1         1            1
```

5.10.2 setfwparam

This command enables you to configure firmware parameters.

Supported By

Windows and Linux

Syntax

`setfwparam <WWPN> <param> <value>`

Parameters

WWPN The WWPN of the adapter port.

param The parameter to modify.

FA-PWWN- Enables or disables the Fabric Assigned WWN (FA-PWWN) parameter. Disabled is the default setting.

NOTE You must reset the adapter port to activate new firmware parameter setting, and you must reload the OneCommand Manager application to display the new setting.

NOTE When a new WWPN is assigned using FA-PWWN, persistently stored configuration information associated with the original WWPN, such as driver parameters and LUN frame priority settings, is not applied to the newly assigned WWPN. The configuration information associated with the original WWPN must be reconfigured for the new WWPN.

Value 0 = Disables the parameter.
1 = Enables the parameter.

Example

```
hbacmd setfwparam 10:00:00:90:fa:94:2e:ca fa-pwwn 1
Set FW Parameter FA-PWWN=1 for 10:00:00:90:fa:94:2e:ca
Reset adapter port to activate new firmware parameter setting.
```

5.11 LUN Masking Commands

The LUN Masking Commands group manages LUN masking activities. LUN masking commands are supported only for FC and FCoE functions.

NOTE Linux and Solaris do not support LUN masking commands.

NOTE Windows + CIM Provider on a VMware host do not support the following commands:

- `GetLunUnMaskByHBA`
- `GetLunUnMaskByTarget`
- `RescanLuns`
- `SetLunMask`

5.11.1 GetLunList

This command queries for the presence of any masked LUNs.

Supported By

Windows and Windows + CIM Provider on a VMware host

Syntax

`GetLunList <HBA WWPN> <Target WWPN> <Option>`

Parameters

HBA WWPN	The WWPN of an FC or FCoE function on the adapter.
Target WWPN	The WWPN of the target.
Option	0 = Get information from the driver 1 = Get information from the configuration

5.11.2 GetLunUnMaskByHBA

This command queries for the presence of any unmasked LUNs by FC or FCoE functions.

Supported By

Windows

Syntax

`GetLunUnMaskByHBA <HBA WWPN> <Option>`

Parameters

HBA WWPN	The WWPN of an FC or FCoE port.
Option	0 = Get information from the driver 1 = Get information from the configuration

5.11.3 GetLunUnMaskByTarget

This command queries for any unmasked LUNs by target.

Supported By

Windows

Syntax

`GetLunUnMaskByTarget <HBA WWPN> <Target WWPN> <Option>`

Parameters

HBA WWPN	The WWPN of an FC or FCoE function.
Target WWPN	The WWPN of the target.
Option	0 = Get information from the driver 1 = Get information from the configuration

5.11.4 RescanLuns

This command rescans LUNs to find any new LUNs.

Supported By

Windows

Syntax

`RescanLuns <HBA WWPN> <Target WWPN>`

Parameters

HBA WWPN	The WWPN of an FC or FCoE function.
Target WWPN	The WWPN of the target.

5.11.5 SetLunMask

This command masks the specified LUNs.

Supported By

Windows

Syntax

`SetLunMask <HBA WWPN> <Target WWPN> <Option> <Lun> <LunCount> <MaskOp>`

Parameters

HBA WWPN	The WWPN of an FC or FCoE function.
Target WWPN	The WWPN of the target.
Option	0 = Get information from the driver 1 = Get information from the configuration (make persistent) 2 = Send information to both
Lun	The starting LUN number.
LunCount	The number of LUNs.

MaskOp
A = Mask LUN
B = Clear unmask target level
C = Clear unmask HBA level
D = Unmask LUN
E = Unmask target level
F = Unmask HBA level

5.12 LUN ExpressLane Commands

The LUN ExpressLane Commands group enables, disables, and displays the ExpressLane status on a particular LUN.

The OneCommand Manager application allows you set special priority queuing for selected LUNs by making them ExpressLane LUNs. ExpressLane LUN performance is superior to that of regular LUNs. You can enable ExpressLane LUNs attached to both physical and virtual ports. ExpressLane LUN assignments persist across system reboots.

NOTE ExpressLane is not supported on LPe12000-series adapters or on LPe16202/OCe15100 adapters in NIC+FCoE mode.

For Linux operating systems, if ExpressLane LUNs are created, the vPort needs to be re-created after a system boot because the ExpressLane LUNs do not persist across system reboots. If the vPort is re-created with the same WWPN to which the ExpressLane LUN was previously assigned and that same LUN is then detected, it becomes an ExpressLane LUN again.

5.12.1 GetExpressLaneLunList

This command displays LUNs on a target and their respective ExpressLane status.

Supported By

Linux, Windows, and Windows + CIM Provider on a VMware host

NOTE For Linux and VMware operating systems, only ExpressLane-enabled LUNs are shown by this command. LUNs without ExpressLane-enabled support are not shown.

Syntax

```
GetExpressLaneLunList <WWPN> [vport=<vPort WWPN>] <Target WWPN> <Option>
```

Parameters

WWPN The WWPN of the FC function connected to the target or physical WWPN if virtual ports are selected.

vPort WWPN The WWPN of an optional vPort allowing you to get the ExpressLane LUNs of a vPort.

Target WWPN The WWPN of the target LUNs.

Option
0 = Get information from driver
1 = Get information from configuration

Example

```
hbacmd h=10.192.87.198 GetExpressLaneLunList 10:00:00:00:00:87:01:98
20:22:d4:ae:52:6e:6f:08 0

Number of LUNs: 4
```

```
      FCP_LUNOS_LUN           ExpressLane
------------------  ------   -------------
0000 0000 0000 00000          No
0001 0000 0000 00001          Yes
0002 0000 0000 00002          No
0003 0000 0000 00003          Yes
```

5.12.2 SetExpressLaneLunState

This command enables or disables ExpressLane on a particular LUN.

NOTE ExpressLane cannot be enabled for masked LUNs.

Supported By

Linux, Windows, and Windows + CIM Provider on a VMware host

Syntax

`SetExpressLaneLunState <WWPN> [vport=<vPort WWPN>] <Target WWPN> <Lun> <State> <Option>`

Parameters

WWPN	The WWPN of the FC function connected to the target or physical WWPN if virtual ports are selected.
vPort WWPN	The WWPN of an optional vPort allowing you to set the state of a vPort LUN.
Target WWPN	The WWPN of the target LUNs.
LUN	The LUN number on which to set the ExpressLane status. Obtain the LUN number from the output of the `GetExpressLaneLunList` command under the OS LUN column.
State	0 = Disable ExpressLane 1= Enable ExpressLane
Option	0 = Set ExpressLane LUN state in driver to temporary, until reboot 1 = Set ExpressLane LUN state in the configuration to persist across reboots 2 = Set ExpressLane LUN state in both driver and in the configuration to persist across reboots

Example

```
hbacmd h=10.192.87.198 SetExpressLaneLUNState 10:00:00:00:00:87:01:98
20:22:d4:ae:52:6e:6f:08 2 1 2
```

5.12.3 GetLunXLaneConfig

This command displays the frame priority value for ExpressLane LUNs on the specified target.

Supported By

Linux, Windows, and Windows + CIM Provider on a VMware host

Syntax

`hbacmd GetLunXLaneConfig <WWPN> [vport=<vPort WWPN>] <Target WWPN> <Option>`

Parameters

WWPN	World-wide port name of any FC/FCoE function on the adapter.
vPort WWPN	The WWPN of an optional vPort.
Target WWPN	The WWPN of the target connected to the LUNs.
Option	0 = Get the information from the driver. 1 = Get the information from the configuration.

Example

```
hbacmd GetLunXLaneConfig 10:00:00:00:c9:55:55:56 50:06:01:60:10:20:5C:38 0

Supported Priority Levels: High, Medium, Low

                                Priority   Priority
    FCP_LUN           OS_LUN    Level      Value
------------------- -------- ---------- ----------
0000 0000 0000 0000    0         High       113
0001 0000 0000 0000    1         High       113
0002 0000 0000 0000    2         Medium     92
0009 0000 0000 0000    9         Low        53
```

5.12.4 SetLunXLaneConfig

If the adapter and switch support it, the `SetLunXLaneConfig` command enables you to configure the ExpressLane Optimized Access Storage (OAS) state and the frame priority levels, or values, for ExpressLane LUNs. Switches can provide up to three priority levels, Low, Medium, and High, but they might provide fewer options.

NOTE If the switch connected to the FC or FCoE initiator does not support LUN specific frame priority levels using the Get Fabric Object (GFO), you must manually enter the frame priority values in the range of 0–127 for all ExpressLane-enabled LUNs.

You can also use the `SetLunXLaneConfig` command to disable ExpressLane on all LUNs in a single operation.

Use the `GetLunXLaneConfig` command to determine if frame priority levels are supported.

The following rules apply when using the `SetLunXLaneConfig` command:

- The ExpressLane (OAS) state and priority levels, or values, will be saved automatically to both the driver and configuration settings. You cannot specify where to save the configuration.
- The priority parameter is only required if the ExpressLane state parameter is set to `1` (enable).
  - — You cannot disable the ExpressLane OAS state with the priority level, or value, parameter set.
  - — You cannot disable all LUNs with the priority level, or value, parameter set.
- The `EnableXLane` driver parameter must be enabled.

NOTE `EnableXLane` will be disabled when `vmid_priority_tagging` and `max_vmid` parameters are enabled.

Supported By

Linux, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
hbacmd SetLunXLaneConfig <WWPN> [vport=<vPort WWPN>] <Target WWPN> <LUN> <State>
<PriorityLevel|PriorityValue>
```

Parameters

WWPN	The World Wide Port Name of any FC/FCoE function on the adapter.
vPort WWPN	The WWPN of an optional vPort.
Target WWPN	The WWPN of the target connected to the LUNs.
LUN	The LUN number. (The OS_LUN from the `GetExpressLaneLunList` command to set.) Use `ALL` to disable ExpressLane for all LUNs.)
State	`0`=disable ExpressLane `1`=enable ExpressLane
PriorityLevel	The levels are high, medium, or low. (Use the `GetSmartSanPriorities` command to get an accurate list of the supported priority levels). This value must be omitted if the *State*=0.
PriorityValue	A value within the range of 0– 127. This value is only allowed if priority levels are not supported by the switch. This value must be omitted if the *State*=0.

Example

```
hbacmd SetLunXLaneConfig 10:00:00:00:c9:55:55:56 50:06:01:60:10:20:5C:38 0 1 low

ExpressLane configuration successfully changed on the specified LUN(s)
```

5.13 Miscellaneous Commands

Commands in the Miscellaneous Command group do not fit in other groups. See specific commands for adapter limitations.

5.13.1 AddHost

This command adds a host to the hosts file for remote TCP/IP management in the OneCommand Manager application. The adapters for these hosts are also presented by the `ListHBAs` command (see Section 5.13.7, ListHBAs).

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

To add non-VMware ESXi hosts:

```
AddHost <hostname|IP_address>[:Port_Number]
```

Parameters

`HostName`	The name of the host to add to the hosts file.
IP_address	The IP address (IPv4 or IPv6) of the host to add to the hosts file. Example IPv4: `10.192.80.102` Example IPv6: `fe80::6445:80e9:9878:a527`

Port_Number The optional IP port number used to access remote host.

Example: `10.192.80.102:9876`

NOTE When specifying IPv6 address with `Port_Number`, it must be enclosed in `[ ]`.

Example: `[fe80::6445:80e9:9878:a527]:9876`

- An attempt is made to contact the host to confirm remote access before adding it to the host list. If the attempt fails, the host is not added.
- The `h` option (for specifying an optional IP address or host name) after `hbacmd` is not available for the `AddHost` command.

To add VMware ESXi hosts to Windows using the OneCommand Manager application:

`m=cim [u=<username>] [p=<password>] [n=<namespace>] AddHost <IP_Address>`

If the *username*, *password*, and *namespace* are not specified, see Section 4.4.2.1.1, Default CIM Credentials.

Parameters

host_address The IP address (using the IPv4 or IPv6 format) or the host name.

5.13.2 Download

This command downloads a firmware image to the port function or adapter specified by the WWPN or MAC address.

Considerations

- If you attempt to update unauthenticated firmware for an LPe31000-series or LPe32000-series adapter, the following error message is displayed:

```
ERROR: Download Failed due to invalid firmware digital signature. Please
contact customer support for additional help.
ERROR: <203>: Failed validating firmware digital signature
```

- If you attempt to update unsecured firmware for an LPe31000-series or LPe32000-series adapter, the following error message is displayed:

```
ERROR: Download Failed due to missing digital signature in firmware file.
Please contact customer support for additional help.
ERROR: <209>: Firmware digital signature missing
```

- For LPe16000-series, LPe31000-series, and LPe32000-series firmware downloads, the OneCommand Manager application accepts only `.grp` files.
- For LPe16000-series, LPe31000-series, and LPe32000-series adapters, the WWPN (and MAC address for LPe16202/OCe15100 adapters) identifies the adapter, and the updated firmware applies to the entire adapter.
- If you attempt to download firmware that is not compatible with the adapter, the following error message is displayed:

```
ERROR: <24>: This firmware version is not supported on this board model.
```

- For LPe12000-series adapters, you update the firmware and boot code on each FC port/function. The firmware and boot code are two separate binaries. You must flash both the firmware and boot binaries to update LPe12000-series adapters.

CAUTION Updating firmware or boot code on an LPe12000-series adapter that is being used to boot from SAN is not advisable. After the update has completed, the system performs an adapter reset, which might cause a loss of connectivity to the SAN and a possible loss of data. To update firmware on an LPe12000-series adapter, you must make sure that the adapter is not currently being used to boot from SAN. Perform one of the following:

- Move the adapter to be updated to a non-boot from SAN host, and perform the update from that location.
- If the host with the adapter that needs to be updated is also hosting other boot from SAN adapters, perform a boot from SAN using one of the other boot from SAN adapters. The target adapter can now be updated, because it is no longer being used for boot from SAN.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

`Download <WWPN|MAC> <FileName>`

Parameters

WWPN	The WWPN of an FC or FCoE function on the adapter.
MAC	The MAC address of a NIC function on the adapter.
FileName	The name and location of the firmware image (any file accessible to the CLI client).

5.13.3 ExportSANInfo

For reporting purposes, this command captures the SAN information in `.xml` for XML-formatted files and `.csv` for CSV-formatted files.

NOTE This command can take a long time on large SAN configurations because of the large amount of information that must be obtained and reported. The output can also be redirected to a file if required.

Supported By

Linux, Solaris, and Windows

Syntax

`ExportSANInfo [format]`

NOTE The `h` option (for specifying an optional IP address or host name) after `hbacmd` is not available for the `ExportSANInfo` command.

Parameters

`format`	An optional parameter that specifies the format of the adapter information: `csv` = CSV-formatted files `xml` = XML-formatted files Leaving the format blank shows the data in xml format (default).

5.13.4 FecEnable

This command enables or disables FEC on LPe16000-series, LPe31000-series, and LPe32000-series FC adapters.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

FecEnable *WWPN <0|1>*

Parameters

WWPN	The WWPN of the FC function.
0	Disables FEC on the function
1	Enables FEC on the function

5.13.5 GetCimCred

This command shows the default credentials set for the CIM client.

NOTE The password is encrypted.

Supported By

Windows

Syntax

`GetCimCred`

Parameters

None.

5.13.6 GetVPD

This command shows the port's VPD.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

`GetVPD <WWPN|MAC>`

Parameters

WWPN	The WWPN of an FC or FCoE function.
MAC	The MAC address of a NIC function.

5.13.7 ListHBAs

This command shows a list of the manageable Broadcom Emulex adapters found by local discovery. For an LPe16202/OCe15100 adapter in NIC+FCoE mode, the MAC address is displayed instead of the WWPN. The node WWN and fabric WWN are not displayed. The type of information listed may vary according to the adapter model.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

`ListHBAs [local] [m=model] [pt=type] [down]`

Parameters

local	Displays only local adapters.
m=model	Model filter. Append * to the end of the model name for a wildcard match. For example: LPe16*
pt=type	The port type filter. Valid types are NIC, FC, and FCoE.
down	Displays only the NIC functions of LPe16202/OCe15100 adapters on the local system in which the adapter's ARM processor has stopped. This parameter detects adapters that might not respond to commands from the OneCommand Manager CLI or application.

5.13.8 RemoveHost

This command removes a host from the hosts file used for TCP/IP management in the OneCommand Manager application GUI. The <*host_address*> can be an IP address that uses the IPv4 or IPv6 format, or it can be a host name.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

For the remote management interface:

```
RemoveHost host_address
```

For VMware ESXi using the CIM interface:

```
m=cim RemoveHost <IP_Address>
```

NOTE The h option (for specifying an optional IP address or host name) after hbacmd is not available for the RemoveHost command.

Parameters

host_address	The host to remove.
IP_Address	The IP address of the host to remove.

5.13.9 Reset

This command resets an FC or FCoE function. A reset can require several seconds to complete, especially for remote devices. When the reset is completed, the system command prompt is displayed.

NOTE This command applies only to FC and FCoE functions.

For LPe16202/OCe15100 FCoE functions, this command only resets the driver to update changed driver parameters that require a driver reset. It does not cause a hardware reset of the FCoE function.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
Reset <WWPN>
```

Parameters

WWPN The WWPN of an FC or FCoE function.

5.13.10 SetCimCred

This command sets the default CIM credentials. You must specify all four credentials: *username*, *password*, *namespace*, and *portnum*. Default credentials are used if any credential is not in the `hbacmd` command argument. After the default credentials for a host are set, any other command can be issued by specifying `m=cim`.

Supported By

Windows

Syntax

`SetCimCred <username> <password> <namespace> <portnum>`

NOTE Use this command to set only the CIM credentials. After this is finished, subsequent `hbacmd` commands do not require you to specify the CIM credentials in the command line.

Parameters

username The logon user ID of the VMware ESXi.

password The logon password of the VMware ESXi.

namespace The namespace where the Emulex CIM provider is registered in the SFCB CIMOM of VMware ESXi, specifically `root/emulex`.

portnum The port number of the SFCB that CIMOM is listening to, that is, `5988` (HTTP) or `5989` (HTTPS).

5.13.11 TargetMapping

This command shows a list of mapped targets and the LUNs for an FC or FCoE function on a port.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

`TargetMapping <WWPN>`

Parameters

WWPN The WWPN of an FC or FoE adapter.

5.13.12 Version

This command shows the current version of the OneCommand Manager CLI Client.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

For the remote management interface:

`Version`

NOTE The `h` option (for specifying an optional IP address or host name) after `hbacmd` is not available for the `Version` command.

For VMware ESXi using the CIM interface:

`h=<IP address> m=<cim Version>`

Parameters

h	The IP address of the VMware ESXi.
m	The CIM version of the VMware ESXi.

5.14 Persistent Binding Commands

The Persistent Binding Commands group facilitates persistent binding operations.

In a remote environment, you can perform persistent bindings operations from a host running any operating system (including Linux or VMware ESXi), but only to a remote host that is running Windows or Solaris.

For a binding to take effect immediately (that is, `SetPersistentBinding parameter: Scope = I` or `B`), the <*SCSIBus*> and <*SCSITarget*> parameters must match the SCSI bus and SCSI target to which the FC or FCoE target is already automapped. If automapping is disabled, the binding takes effect immediately if the FC or FCoE target is not already persistently bound, and the specified <*SCSIBus*> and <*SCSITarget*> parameters are available to be persistently bound. Also, the <*BindType*> parameter must match the currently active bind type. Otherwise, you are notified that you must reboot the system to cause the persistent binding to become active.

These commands are supported only for FC and FCoE ports.

The following persistent binding commands are not supported on Linux or on Windows + CIM Provider on a VMware host:

- `BindingCapabilities`
- `BindingSupport`
- `PersistentBinding`
- `RemoveAllPersistentBinding`
- `RemovePersistentBinding`
- `SetBindingSupport`
- `SetPersistentBinding`

5.14.1 AllNodeInfo

This command shows target node information for each target accessible by the adapter.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

`AllNodeInfo <WWPN>`

Parameters

WWPN	The WWPN of an FC or FCoE function.

5.14.2 BindingCapabilities

This command shows the binding capabilities of the adapter. If a binding is configured, it is maintained across reboots.

Supported By

Solaris and Windows

Syntax

`BindingCapabilities <WWPN>`

Parameters

WWPN	The WWPN of an FC or FCoE function.

5.14.3 BindingSupport

This command shows the binding support for an FC or FCoE function.

Supported By

Solaris and Windows

Syntax

`BindingSupport <WWPN> <Source>`

Parameters

WWPN	The WWPN of an FC or FCoE function.
Source	C = Configuration support L = Live support

5.14.4 PersistentBinding

This command specifies the set of persistent binding information (configuration or live state) that is requested.

Supported By

Citrix, Solaris and Windows

Syntax

`PersistentBinding <WWPN> <Source>`

Parameters

WWPN	The WWPN of an FC or FCoE function.
Source	C = Configuration support L = Live support

5.14.5 RemoveAllPersistentBinding

This command removes all persisting bindings for an FC or FCoE function.

Supported By

Solaris and Windows

Syntax

`RemoveAllPersistentBinding <WWPN>`

Parameters

WWPN	The WWPN of an FC or FCoE function.

5.14.6 RemovePersistentBinding

This command removes persistent binding between an FC or FCoE target and a SCSI bus and target. The binding to be removed can be to a target WWPN, a target WWNN, or a target D_ID.

Supported By

Solaris and Windows

Syntax

`RemovePersistentBinding <WWPN> <BindType> <ID> <SCSIBus> <SCSITarget>`

Parameters

WWPN	The WWPN of an FC or FCoE function.
BindType	P = Remove binding by WWPN N = Remove binding by WWNN D = Remove binding by D_ID
ID	The type of ID based on <*BindType*>: Target WWPN if <*BindType*> = P Target WWNN if <*BindType*> = N Target D_ID if <*BindType*> = D
SCSIBus	The bus number of the SCSI device.
SCSITarget	The target number of the SCSI device.

5.14.7 SetBindingSupport

This command enables and sets the binding support for an FC or FCoE function.

Supported By

Solaris and Windows

Syntax

`SetBindingSupport <WWPN> <BindFlag>`

Parameters

WWPN	The WWPN of an FC or FCoE function.

BindFlag	The type of binding support for the adapter: D = Binding by D_ID (not available for Windows driver) P = Binding by WWPN N = Binding by WWNN (not available for Windows driver) A = Binding by automap (not available for Windows driver) DA = Binding by D_ID and automap PA = Binding by WWPN and automap NA = Binding by WWNN and automap

5.14.8 SetPersistentBinding

This command sets a persistent binding between an FC or FCoE target and a SCSI bus target. The binding can be to a target WWPN, a target WWNN, or a target D_ID.

Supported By

Solaris and Windows

Syntax

`SetPersistentBinding <WWPN> <Scope> <BindType> <TargetId> <SCSIBus> <SCSITarget>`

Parameters

WWPN	The WWPN of an FC or FCoE function.
Scope	P = Permanent binding (survives reboot) I = Immediate binding B = Binding is both permanent and immediate
BindType	P = Enable binding by WWPN N = Enable binding by WWNN D = Enable binding by D_ID
TargetId	If <*BindType*> = P, Target WWPN If <*BindType*> = N, Target WWNN If <*BindType*> = D, Target D_ID
SCSIBus	The bus number of the SCSI device.
SCSITarget	The target number of the SCSI device.

5.15 vPort Commands

The vPort Commands group manages virtual ports and functions on FC and FCoE adapters.

NOTE In Linux, vPorts do not persist across system reboots. vPorts must be re-created after a system reboot.

5.15.1 CreateVPort

This command creates a virtual port with an automatically generated WWPN or a user-specified virtual WWPN on the specified physical port. If you specify `auto`, the virtual WWPN is generated automatically. Otherwise, you must specify the virtual WWPN for this parameter. If creation is successful, the WWPN is displayed as part of the output from the command. The *vname* optional parameter can be specified for the virtual port's name.

Supported By

Linux, Solaris, and Windows

Syntax

`CreateVPort <physical WWPN> auto [vname]`

-or-

`CreateVPort <physical WWPN> <virtual WWPN> <virtual WWNN> [vname]`

Parameters

physical WWPN	The WWPN of an FC or FCoE function.
auto	The virtual WWPN is automatically generated for the virtual port.
vname	The virtual port's name (optional).
virtual WWPN	The virtual WWPN to create.
virtual WWNN	The virtual WWNN to create.

5.15.2 DeleteVPort

This command deletes the virtual port specified by a physical and virtual WWPN.

Supported By

Linux, Solaris, and Windows

Syntax

`DeleteVPort <physical WWPN> <virtual WWPN>`

Parameters

physical WWPN	The WWPN of an FC or FCoE function.
virtual WWPN	The WWPN of the virtual port.

5.15.3 ListVPorts

This command lists virtual ports on the specified physical FC or FCoE function. Leaving the physical WWPN parameter blank lists all virtual ports on all manageable hosts that support virtual ports.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

`ListVPorts <physical WWPN>`

Parameters

physical WWPN	The WWPN of an FC or FCoE function.

5.15.4 VPortTargets

This command lists targets visible to the specified virtual port.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

`VPortTargets <physical WWPN> <virtual WWPN>`

Parameters

physical WWPN	The WWPN of an FC or FCoE function.
virtual WWPN	The WWPN of the virtual port.

5.16 WWN Management Commands

NOTE These commands are supported only for FC or FCoE functions.

The WWN Management Commands group validates WWNs carefully to avoid WWPN duplication, but WWNN duplication is acceptable. You might see error and warning messages if a name duplication is detected. Fulfill the activation requirement after each WWN is changed or restored. If pending changes exist, some diagnostic and maintenance features are not available.

CAUTION Using the `ChangeWWN` or `RestoreWWN` commands on an LPe12000-series adapter that is being used to boot from SAN is not advisable. After the command is completed, the system performs an adapter reset, which might cause a loss of connectivity to the SAN and possible loss of data. To perform these commands, you must make sure that the adapter is not currently being used to boot from SAN.

Do one of the following:

- Move the target adapter to a non-boot from SAN host.
- If the host with the target adapter is also hosting other boot from SAN adapters, perform a boot from SAN using one of the other boot from SAN adapters. The target adapter can now be used.

5.16.1 ChangeWWN

This command allows you to change WWPNs and WWNNs, and it allows you to change the WWN to volatile or nonvolatile. If you attempt to select volatile on an adapter that does not support volatile WWNs, a `Not Supported` error is returned.

NOTE This command is disabled when FA-PWWN is enabled on the adapter port.

When a volatile change is supported, a reboot is required to activate the new value. Volatile names are active until system power-down or adapter power-cycle.

For VMware ESXi:

- After changing the WWN of a function, update your zoning settings before you reboot your ESXi server. If the zoning is not updated before your reboot, the subsequent boot could take a long time.
- After changing the WWN of a function, you must reboot the ESXi system before trying to access the adapter on that system. For information on rebooting the ESXi system, refer to the VMware documentation.
- If you are using the CIM interface to access functions, after changing the WWN of a function, you must restart the CIMOM (that is, SFCB) on the ESXi system before trying to access the function on that system. For information on restarting the CIMOM, refer to the VMware documentation.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

`ChangeWWN <WWPN> <New WWPN> <New WWNN> <Type>`

Parameters

WWPN	The WWPN of an FC or FCoE function.
New WWPN	The WWPN of the FC or FCoE function.
New WWNN	The WWNN of an FC or FCoE function.
Type	0 = Volatile 1 = Nonvolatile

5.16.2 GetWWNCap

This command shows if volatile change is supported for the WWPN.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

`GetWWNCap <WWPN>`

Parameters

WWPN	The WWPN of an FC or FCoE function.

5.16.3 ReadWWN

This command reads different types of WWNs.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

`ReadWWN <WWPN> <Type>`

Parameters

WWPN	The WWPN of an FC or FCoE function.
Type	0 = Volatile 1 = Nonvolatile 2 = Factory default 3 = Current 4 = Configured

5.16.4 RestoreWWN

This command changes the WWNs to the factory default or nonvolatile values. The change is nonvolatile.

NOTES
- A reboot is required to activate the new value.

- This command is disabled when FA-PWWN is enabled on the adapter port.

For VMware ESXi:

- After changing the WWN of an function, you must reboot the ESXi system before trying to access the adapter on that system. For information on rebooting the ESXi system, refer to the VMware documentation available from the VMware website.
- If you are using the CIM interface to access adapters, after changing the WWN of a function, you must restart the CIMOM (that is, SFCB) on the ESXi system before trying to access the function on that system. For information on restarting the CIMOM, refer to the VMware documentation available from the VMware website.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

`RestoreWWN <WWPN> <Type>`

Parameters

WWPN	The WWPN of an FC or FCoE function.
Type	0 = Restore default WWNs 1 = Restore NVRAM WWNs

Appendix A: OneCommand Manager Error and Return Messages

Table 1 contains a list of some of the error messages that may be encountered during a OneCommand Manager session.

Table 7 OneCommand Manager Error and Warning Messages

Error Message	Commands	Description
`Error: Read-only management mode is currently set on this host. The requested command is not permitted in this mode.`	Active management commands that change a property on an adapter or host.	This message is returned when some commands are attempted when the CLI is configured for read-only mode. See Section 4.1.1, CLI in Read-Only Mode.
`Not supported.`	`ChangeWWN`	If a volatile change is requested on an adapter that does not support volatile WWNs, it returns a `Not Supported` error. See Section 5.16.1, ChangeWWN.
`RETURN CODE: <0>: Success, no further action is needed.`	`Download`	The firmware download completed without errors.
`ERROR: HBACMD_GetDumpFile: RM_GetDumpFile call failed (2)` `ERROR: <2>: Not Supported`	`GetDumpFile`	Dump files are copied from the Dump directory of the remote host to the Dump directory of the local host. Specifying a local port identifier for this command returns an error, because the source and destination directory are the same. See Section 5.8.4, GetDumpFile.
`ERROR: <4>: Invalid argument`	`getfwparams` `setfwparam`	This message is returned when there is a problem with the command. See Section 5.10, Firmware Commands.
`ERROR: <5>: Illegal WWN format`	`getfwparams` `setfwparam`	This message is returned when the WWN format is incorrect. See Section 5.10, Firmware Commands.
`ERROR: <24>: This firmware version is not supported on this board model.`	`Download`	This message is returned when the firmware version is incompatible with the adapter. Download compatible firmware. See Section 5.13.2, Download.
`ERROR: <35>: Unable to allocate buffer`	`getfwparams`	This message is returned when the command cannot allocate a buffer. See Section 5.10, Firmware Commands.
`ERROR: <180>: Authentication: User unknown`	All	The specified user name is not valid or could not be authenticated by the system. See Section 1.2, OneCommand Manager Secure Management, for more information.
`ERROR: <181>: Authentication: Insufficient credentials`	All	The specified user name and password are valid and the user is a member of a OneCommand Manager group. However, the OneCommand Manager group does not have sufficient privileges to execute the specified command. See Section 1.2, OneCommand Manager Secure Management, for more information.
`ERROR: <183>: Secure Mgmt: user not a member of OCM group`	All	The specified user name and password could be authenticated, but the user is not a member of a OneCommand Manager group. See Section 1.2, OneCommand Manager Secure Management, for more information.

Table 7 OneCommand Manager Error and Warning Messages (Continued)

Error Message	Commands	Description
`ERROR: Download Failed due to invalid firmware digital signature. Please contact customer support for additional help.` `ERROR: <203>: Failed validating firmware digital signature`	`Download`	If you attempt to update unauthenticated firmware for a secure LPe31000-series or LPe32000-series adapter, this error message is displayed. See Section 5.13.2, Download.
`ERROR: <206>: Authentication Failed`	All	This indicates either a valid user name with an invalid password, or a general user authentication error. See Section 1.2, OneCommand Manager Secure Management, for more information.
`ERROR: Download Failed due to missing digital signature in firmware file. Please contact customer support for additional help.` `ERROR: <209>: Firmware digital signature missing`	`Download`	If you attempt to update unsecured firmware for a secure LPe31000-series or LPe32000-series adapter, this error message is displayed. See Section 5.13.2, Download.
`RETURN CODE: <247>: Download succeeded, but a reboot is required to activate the new firmware.`	`Download`	Reboot the system to activate the new firmware.
`ERROR: <251>: Hardware or firmware does not support command`	`getfwparams` `setfwparam`	This message is returned when the hardware or firmware does not support the command. See Section 5.10, Firmware Commands.
`Error: <431> Cable length required for force mode and interface type.`	`SetPhyPortSpeed`	This error is displayed when a length value is not included when the mode is set to 2. Example: `hbacmd setphyportspeed 00-00-c9-a9-41-88 2 100Mb` See Examples.

Appendix B: License Notices

B.1 Secure Hash Algorithm (SHA-1) Notice

B.2 OpenPegasus Licensing Notice

persons to whom the Software is furnished to do so, subject to the following conditions:
The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.
THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

B.3 OpenSSL Notice

This is a copy of the current LICENSE file inside the CVS repository.
LICENSE ISSUES
==========
The OpenSSL toolkit stays under a dual license, i.e. both the conditions of the OpenSSL License and the original SSLeay license apply to the toolkit. below for the actual license texts. Actually both licenses are BSD-style Open Source licenses. In case of any license issues related to OpenSSL please contact openssl-core@openssl.org.
OpenSSL License

/* ==
* Copyright (c) 1998-2011 The OpenSSL Project. All rights reserved.
*
* Redistribution and use in source and binary forms, with or without
* modification, are permitted provided that the following conditions
* are met:
*
* 1. Redistributions of source code must retain the above copyright
* notice, this list of conditions and the following disclaimer.
*
* 2. Redistributions in binary form must reproduce the above copyright
* notice, this list of conditions and the following disclaimer in
* the documentation and/or other materials provided with the
* distribution.
*
* 3. All advertising materials mentioning features or use of this
* software must display the following acknowledgment:
* "This product includes software developed by the OpenSSL Project
* for use in the OpenSSL Toolkit. (http://www.openssl.org/)"
*
* 4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to
* endorse or promote products derived from this software without
* prior written permission. For written permission, please contact
* openssl-core@openssl.org.


BROADCOM®

Emulex® Elxflash and LpCfg Utilities for LightPulse® Adapters

User Guide

Version 11.4
September 6, 2017

FlashUTIL-LPE-UG114-100

Corporate Headquarters
San Jose, CA

Website
www.broadcom.com

Chapter 1: Introduction

This document explains how to use the Emulex® Elxflash and LpCfg adapter management utility kits for Linux, Windows, and VMware pre-boot environments.

1.1 Elxflash Offline Kit

The Elxflash Offline kit allows you to configure Emulex adapters before you install or boot a server operating system. You can also use the Elxflash Stand-alone kit, which lets you update firmware on Emulex adapters without installing any tools. See Section 4.2, Running the LpCfg Utility Included with the Elxflash Stand-alone Kit, for more information.

The Elxflash Offline kit includes:

- linlpcfg and elxflash for Linux
- winlpcfg and elxflash for Windows

The linlpcfg and the winlpcfg are management utilities for adapter configuration and diagnostics for Linux and Windows, respectively. The LpCfg utility can also download firmware, but you must manually select the adapters to be updated.

Elxflash is a firmware download utility. When the `/auto` switch is used, Elxflash automatically chooses the best image to download. When the `fwmatrix.txt` file is used, you choose the images to apply to the adapters.

NOTE The Offline utility must be installed on your system before you can run it. While the stand-alone utility does not need to be installed before it is run, you must go to the directory into which the stand-alone utility was extracted and run it from the command prompt.

See Chapter 4, LpCfg Utility Command Line Interface, for further information about the offline utility.

1.2 Elxflash Utility

The Elxflash utility uses a command line interface allowing you to build scripts for automated and unattended firmware and boot code download solutions for Emulex LightPulse® Fibre Channel (FC) host bus adapters (HBAs), and Fibre Channel over Ethernet/network interface card (NIC+FCoE) in production systems. You can also download firmware and boot code on local and remote machines simultaneously.

The Elxflash Stand-alone kit, which allows you to update firmware on Emulex adapters without any tools, is also available. See Section 4.2, Running the LpCfg Utility Included with the Elxflash Stand-alone Kit, for more information.

The Elxflash utility uses auto-discovery (`/auto`), so you do not need to maintain the `fwmatrix.txt` configuration file. When auto-discovery is invoked, the utility automatically discovers local adapters and performs selected operations, such as: `/ff` (force firmware), `/fb` (force boot), `/update`, `/downgrade`, and `/rewrite`.

When using auto-discovery, you must populate the firmware and boot subdirectories with image files for specific adapter models. Auto-discovery uses Emulex-defined firmware and boot code file names in their respective directories.

If you do not want to use auto-discovery, the Elxflash utility can update the firmware or boot code of an adapter using the `fwmatrix.txt` configuration file. When using the `fwmatrix.txt` file, you must update each supported adapter's type, firmware, and boot code entries and place the corresponding firmware and boot code images into their respective directories. The `fwmatrix.txt` file enables all operations that are used with auto-discovery.

1.3 Command Syntax

In all cases, the commands are given with the prefix *os*. When you enter any command, you must replace *os* with either:

- `./lin` – For the Linux LpCfg utility commands, for use on Linux machines.
- `win` – For the WinPE LpCfg utility commands, for use on WinPE machines.

NOTE WinPE is supported on FC adapters only. It is not supported on LPe16202/OCe15100 adapters in NIC+FCoE mode.

The LpCfg utility commands are not case-sensitive, nor are their arguments. You can enter them in uppercase, lowercase, or mixed case. However, for Linux operating systems, use `linlpcfg` (all lowercase) for the utility commands; for example, `./linlpcfg listhba`.

The LpCfg utility commands use both of the following items:

- One space between the command name and the first argument.
- One space between additional arguments.

NOTE Do not put a space before or after the equal sign within an argument.

World Wide Name (WWN) values are reported and specified with two 4-byte hexadecimal words – WWN word 0 (w0) and WWN word 1 (w1). To make the full WWN, concatenate the w0 and w1 values.

For detailed information on the command line interface, see Chapter 3, Elxflash Utility Command Line Interface, and Chapter 4, LpCfg Utility Command Line Interface.

1.4 Supported Platforms

For information on supported adapters and supported versions of operating systems and platforms, see http://www.broadcom.com.

1.5 Abbreviations

BIOS	basic input/output system
CFA	converged fabric adapter
EFI	Extensible Firmware Interface
EROM	erasable read-only memory
FC	Fibre Channel
FCoE	Fibre Channel over Ethernet
HBA	host bus adapter
HEX	hexadecimal
IEEE	Institute of Electrical and Electronics Engineers
GUI	graphical user interface
ISO	International Organization for Standards
LUN	logical unit number
MAC	media access control
MILI	management interface library
ms	milliseconds
NIC	network interface card
PCI	Peripheral Controller Interconnect
PCI_ID	PCI device identification number
PE	Preinstallation Environment
POST	power-on self-test
RHEL	Red Hat Enterprise Linux
ROM	read-only memory
RPM	Red Hat Package Manager
SAN	storage area network
SLES	SuSE Linux Enterprise Server
UCNA	universal converged network adapter
VPD	Vital Product Data
VWWNN	Virtual World Wide Node Name
VWWPN	Virtual World Wide Port Name
WWN	World Wide Name
WWNN	World Wide Node Name
WWPN	World Wide Port Name
XML	Extensible Markup Language

Chapter 2: Installation

This chapter details prerequisites and procedures for installing and uninstalling the Elxflash Offline and Stand-alone kits for Linux, VMware, and Windows systems. It also describes the update procedure for each operating system.

2.1 Platform Prerequisites

The utilities included in the Elxflash Offline and Stand-alone kits for Linux have prerequisites that must be installed prior to installing the utilities. There are no prerequisites for installing the Elxflash Offline and Stand-alone kits on the VMware and Windows operating systems.

2.1.1 Linux

The Linux platform uses additional software, such as libraries, to run the different versions of the utility. This section identifies the prerequisites.

NOTE The Linux driver must be installed on the system for the applications to work properly.

The following software must be installed to run the utility.

- libnl
- ethtool
- lspci
- libsysfs (optional)

2.2 Installing the Elxflash Offline Kit for Linux

The Elxflash Offline kit for Linux uses an install script to install the elxflash and linlpcfg utilty RPM packages. The install script determines the correct architecture and distribution, and it updates the existing Elxflash Offline and Elxlinlpcfg RPM packages. If there are no existing Elxflash Offline or Elxlinlpcfg RPM packages, the install script installs the packaged Elxflash Offline and Elxlinlpcfg RPM packages.

To install the Elxflash Offline kit for Linux:

1. Untar the installation tarball.
2. Run the install script located on the root of the installation kit.

Example:

```
$ tar zxvf elxflashOffline-linux-<version>-<rel>.tgz
$ cd elxflashOffline-linux-<version>-<rel>
$ ./install.sh
```

2.2.1 Updating

To update an existing Elxflash Offline kit for a Linux installation, run the `install.sh` script to update the Elxflash Offline and Elxlinlpcfg RPM packages.

2.2.2 Uninstalling

The Elxflash Offline kit for Linux uses an uninstall script to uninstall the utility. The uninstall script performs specific actions depending on the switches that are used.

- `./uninstall.sh` – Uninstalls the Elxflash Offline and Elxlinlpcfg RPM packages.
- `./uninstall.sh -h` – Displays a summary of all available switches.

To uninstall the Elxflash Offline kit for Linux, run the uninstall script located on the root of the installation kit.

Example:

```
$ cd elxflashOffline-linux-<version>-<rel>
$ ./uninstall.sh
```

2.3 Installing the Elxflash Offline Kit for Windows PE

To install the Elxflash Offline kit for Windows PE:

1. Unzip the `Offline-WinPE-<version>-<rel>.zip`.
2. Change the directory to the correct architecture subdirectory.
3. Run `setupElxAll-<arch>.exe`.

The following components are installed:

- Storport CNA driver
- Storport Fibre Channel (FC) driver
- winLpCfg – Elxflash Offline

NOTE The `setupElxAll-<arch>.exe` file is compatible only with WinPE. It does not function with any version of Windows Server.

2.3.1 Updating

To update an existing Elxflash Offline kit installation, run the same steps that you used for installing the Elxflash Offline kit for Windows PE previously.

The installer uninstalls the existing version and then installs the updated version.

2.3.2 Uninstalling

To uninstall the Elxflash Offline kit for Windows PE, run the following commands:

```
cd \ProgramData\Tarma Installer\
setup.exe /remove
```

2.4 Extracting the Driver Files from the Elxflash Offline Kit for Windows

To extract the driver files from the Elxflash Offline kit, you can use the installer or the command line. When extracting these files, all files, both x64 and x86, are extracted to `<Drive>:\My Documents\Emulex\Drivers`. This location is hard-coded and cannot be changed.

NOTE Extracting all the drivers takes an additional 28 MB. By default, Windows PE allocates only 32 MB of writable memory. The size of the writable memory can be increased when building the WinPE ISO image.

You can extract the driver files using the GUI mode by running the installer, or you can use the command line in silent mode.

To extract the drivers from the Elxflash Offline kit for Windows, perform these steps:

1. In GUI mode:
   a. Run the installer.
   b. From the Installation Options screen, select **Extract All Drivers**, and deselect **Install elxApp_Drivers (64-bit)**.
   c. Click **Install** to continue.
2. In Silent mode, from the command line, type:

```
start/wait SetupElxAll-x64.exe/q2 extract=2
```

NOTE Values for the `extract` parameter are:

- `0` – Install the package; do not extract the drivers (default)
- `1` – Install the package; extract the drivers
- `2` – Do not install the package; extract the drivers

2.5 Extracting the Elxflash Stand-alone Kit for Linux, Windows, and VMware

The Elxflash Stand-alone kit is not installed; you must extract it from its zip file. After you extract the utility, the following directories are created. The Linux executable files are extracted to the `lx\` directory, the VMware executable files are extracted to the `esxi\` directory, and the Windows executable files are extracted to the `win\` directory.

- `boot\`
- `firmware\`
- `lx\`
- `esxi\`
- `win\`

In Windows, for example, three directories are visible in the kit's root directory after you extract the stand-alone utility. The top-level boot directory and firmware directory are the default directories in which Elxflash looks for boot code and firmware images. Ensure that firmware and boot code are located in these directories.

```
# dir

07/05/2012  07:02 PM    <DIR>          .
07/05/2012  07:02 PM    <DIR>          ..
07/02/2012  03:14 PM    <DIR>          boot
07/02/2012  03:14 PM    <DIR>          firmware
07/02/2012  03:15 PM    <DIR>          win
```

The Elxflash Stand-alone kit for Linux has the same dependencies as the Elxflash Offline kits for Linux. The Elxflash Stand-alone kits for Windows and VMware do not have any dependencies. See Section 2.1, Platform Prerequisites.

NOTE Update and uninstall procedures for the Elxflash Stand-alone kit are not needed. When a new version is available, you extract it, replacing the older version.

NOTE To manage LightPulse adapters, you must install the operating system-specific FC drivers. To manage LPe16202/OCe15100 adapters, you must install the operating system-specific NIC drivers and the operating system-specific FC/FCoE or FCoE drivers.

Chapter 3: Elxflash Utility Command Line Interface

Firmware images are available on the Broadcom® support site at http://www.broadcom.com. For all adapters, firmware image files are expected to be in the sub-directory named `firmware`. For all LPe16000-series, LPe31000-series, and LPe32000-series adapters, the firmware image file includes boot code. LPe12000-series adapters use a separate boot code image file, which is expected to be in the sub-directory named `boot`.

NOTE In this chapter, references to FCoE and NIC apply only to LPe16202/OCe15100 adapters in NIC+FCoE mode. References to FC apply to all other LightPulse adapters and to LPe16202 adapters in FC mode.

Two supported modes for updating firmware and boot code are available for each Elxflash utility switch. The `fwmatrix.txt` file mode allows you to specify the adapter models and firmware file (or firmware and boot files) to update. Auto-discovery mode automatically updates the discovered adapters with the contents of the `firmware` directory (and the `boot` directory, if applicable).

Regardless of the firmware update mode you choose, use the `/query` (`/q`) switch with all operational switches.

For example:

```
# ./elxflash /q
HBA=LPe32000, Port Type=FC, WWN=10:00:00:90:FA:94:2E:CA, PCI ID=E300, VID=10DF,
SSID=E321, SVID=10DF, Firmware=11.2.50.48, Boot Code=11.2.50.35, Boot Enabled=1
HBA=LPe12002, Port Type=FC, WWN=10:00:00:00:C9:A1:80:00, PCI ID=F100, VID=10DF,
SSID=F100, SVID=10DF, Firmware=UD202A1, Boot Code=UU700A2, Boot Enabled=1
HBA=LPe12002, Port Type=FC, WWN=10:00:00:00:C9:A1:80:01, PCI ID=F100, VID=10DF,
SSID=F100, SVID=10DF, Firmware=UD202A1, Boot Code=UU700A2, Boot Enabled=1
```

If the `/q` switch is used with an operational switch, the results include an additional field called `Supported Firmware` or `Supported Boot Code`.

Regardless of the mode and operational switch you choose, one of the following download summaries is displayed after the command is executed:

- FC adapter download summary:

  ```
  <date><time>
  HBA=<model>, Port Type=<port_type>, WWN=<wwn>,
  Update=<Boot Code|Firmware>, Image=<image>, New=<version>, Old=<version>,
  Status=<description>
  ```

- FCoE adapter download summary

  ```
  <date><time>
  HBA=<model>, Port Type=<port_type>, WWN=<wwn>,
  Update=Firmware, Image=<image>, New=<version>, Old=<version>,
  Status=<description>
  ```

- NIC adapter summary

  ```
  <date><time>
  HBA=<model>, Port Type=<port_type>, MAC=<mac_address>,
  Update=Firmware, Image=<image>, New=<version>, Old=<version>,
  Status=<description>
  Return Code=<n>
  ```

  where <*description*> is `Success` or `Error`, and <*n*> = `0` for completion with no errors or a nonzero error code for any error.

3.1 Using the fwmatrix.txt File Mode

To use the `fwmatrix.txt` file, you must update the firmware and boot code directories with the appropriate firmware and boot code images. The `fwmatrix.txt` file includes adapter entries followed by a firmware image file entry. An additional boot image entry is provided for LPe12000-series adapters.

NOTE To avoid a firmware update failure, ensure that there are no spaces in the model names in the `fwmatrix.txt` file.

When the Elxflash utility uses the `fwmatrix.txt` file, it locates the adapter model entry and then attempts to download the specified firmware image file and, for LPe12000-series adapters, the specified boot code image file.

On an LPe16202/LPe15100 adapter running in NIC+FCoE mode, the Elxflash model name must include the highest protocol being used on that model. You can identify an adapter's Elxflash model name by running the `/query` command.

NOTE On an LPe16202/LPe15100 adapter running in NIC+FCoE mode, the highest protocol is FCoE.

3.2 Using Auto-Discovery Mode

If the `/auto` switch is used with an operational switch, such as `/update`, the Elxflash utility automatically discovers adapters, and using the `firmware` and `boot` subdirectories, performs the operation specified by the switch on each adapter.

3.3 Auto-Discovery (/auto)

Usage: `/auto`

The `/auto` switch instructs the Elxflash utility to ignore the `fwmatrix.txt` file, automatically discover local adapters, and perform specified operations employing an additional switch using the `firmware` and `boot` directories.

The `/auto` switch must be used with an additional operational switch, such as:

- `/f`
- `/ff`
- `/fb`
- `/downgrade`
- `/rewrite`
- `/update`

Example usage:

`./elxflash /auto /update /q` – Updates the firmware and boot code using the `firmware` and `boot` directories.

- The `fwmatrix.txt` file is ignored. The desired versions of firmware must be in the `firmware` directory.
- Using the `firmware` subdirectory, the Elxflash utility automatically discovers the best matching firmware for each installed and supported adapter.

- If multiple versions of firmware or boot code are found for an adapter, the Elxflash utility uses the most recent version when performing the firmware update.

3.4 Downgrade (/downgrade or /g)

Usage: `/downgrade` or `/g`

The `/downgrade` switch downgrades the firmware or boot code of each adapter if the currently installed versions are more recent than the downgrade versions. This switch cannot be used with the `/update` or `/rewrite` commands.

Example usage:

`./elxflash /downgrade /auto /q` – Downgrades the firmware or boot code using auto-discovery.

- The `fwmatrix.txt` file is ignored. The desired downgrade versions of firmware or boot code must be in their respective directories.
- If the downgrade versions are older than the currently installed versions on the adapter, the downgrade versions are downloaded to the adapter.
- If multiple downgrade versions of firmware or boot code are found for an adapter, the next-previous downgrade versions are downloaded to the adapter.
- When performing the boot code downgrade operation, the Elxflash utility first tries to match by adapter family and boot type. If a match is not found, the Elxflash utility then tries to match by boot type. If the utility matches by boot type and multiple versions of boot code are detected, the utility chooses the downgrade file in the following order:
  a. Universal (U)
  b. Pair (P)
  c. Open (O)
  d. EFI (E)
  e. x86 (B)

`./elxflash /downgrade /q` – Downgrades the firmware or boot code using the `fwmatrix.txt` file.

- For each installed and supported adapter, the current firmware or boot code versions are compared with the versions specified in `fwmatrix.txt`.
- If the downgrade version in `fwmatrix.txt` is older than the currently installed version, the downgrade version of firmware or boot code is downloaded to that adapter.

3.5 Force Firmware and Boot Code (/f)

Usage: `/f`

The `/f` switch forces a firmware and boot code download to an adapter regardless of the current version on the adapter, and it is performed regardless of any additional operational switches given on the command line. Also, see Section 3.6, Force Boot Code (/fb), and Section 3.7, Force Firmware (/ff).

Example usage:

`./elxflash /f /auto /q` – Forces a firmware and boot code download using auto-discovery.

- The `fwmatrix.txt` file is ignored. The desired versions of firmware and boot code must be in their respective directories.

- If multiple versions of firmware or boot code are found for an adapter, the Elxflash utility uses the most recent version when performing the firmware and boot code downloads.

`./elxflash /f /q` – Forces a firmware and boot code download using the `fwmatrix.txt` file.

- For each installed and supported adapter, this command forces a download of firmware and boot code using the versions specified in the `fwmatrix.txt` file.

3.6 Force Boot Code (/fb)

Usage: `/fb`

The `/fb` switch forces a boot code download to an adapter regardless of the boot code version installed on the adapter, and it is performed regardless of any additional operational switches given on the command line.

Example usage:

`./elxflash /fb /auto /q` – Forces a boot code download using auto-discovery.

- The `fwmatrix.txt` file is ignored. The desired versions of boot code must be in the `boot` directory.
- If multiple versions of boot code are found for an adapter, the Elxflash utility uses the most recent version when performing the boot code download.

`./elxflash /fb /q` – Forces a boot code download using the `fwmatrix.txt` file.

- For each installed and supported adapter, this command forces a download of boot code using the boot code version specified in the `fwmatrix.txt` file.

NOTE The `/fb` switch applies only to LPe12000-series adapters.

3.7 Force Firmware (/ff)

Usage: `/ff`

The `/ff` switch forces a firmware download to an adapter regardless of the firmware version installed on the adapter or any additional operational switches given on the command line.

NOTE Because boot code is included in the firmware image for LPe16000-series, LPe31000-series, and LPe32000-series adapters, this command has the same result as the `/f` command, which forces a firmware and boot code download.

Example usage:

`./elxflash /ff /auto /q` – Forces a firmware download using auto-discovery.

- The `fwmatrix.txt` file is ignored. The desired versions of firmware must be in the `firmware` directory.
- If multiple versions of firmware are found for an adapter, the Elxflash utility uses the most recent version when performing the firmware download.

`./elxflash /ff /q` – Forces a firmware download using the `fwmatrix.txt` file.

- For each installed and supported adapter, this command forces a download of firmware using the firmware version specified in the `fwmatrix.txt` file.

3.8 Firmware Matrix Directory Change (/fmd)

Usage: `/fmd=<directory>`

The `/fmd` switch changes the location of the base directory in which `fwmatrix.txt` is located.

3.9 Help (/h or /?)

Usage: `/h or /?`

The `/h` switch displays a help message detailing instructions on how to use the Elxflash utility.

3.10 Image Directory Change (/id)

Usage: `/id=<image_directory>`

The `/id` switch is used to specify the location of the `firmware` and `boot` directories.

Example usage:

`./elxflash /f /auto /id=/tmp /q` – Forces a firmware and boot code download using auto-discovery.

- Elxflash looks for the `firmware` and `boot` directories in the `/tmp` directory.
- The `fwmatrix.txt` file is ignored. The desired versions of firmware and boot code must be in their respective directories.
- In this example, firmware must be placed in the `/tmp/firmware` directory, and boot code must be placed in the `/tmp/boot` directory.
- If multiple versions of firmware or boot code are found for an adapter, Elxflash uses the most recent versions when performing the firmware and boot code downloads.

`./elxflash /f /id/tmp /q` – Forces a firmware and boot code download using the `fwmatrix.txt` file.

- Elxflash looks for the `firmware` and `boot` directories in the `/tmp` directory.
- For each installed and supported adapter, a forced download of firmware and boot code occurs using the versions specified in the `fwmatrix.txt` file.
- In this example, firmware must be placed in the `/tmp/firmware` directory, and boot code must be placed in the `/tmp/boot` directory.

NOTE This switch was `/i` in previous versions of the Elxflash utility.

3.11 Image Version Display (/iv)

Usage: `/iv=<image file>`

The `/iv` switch shows the firmware image file's version number. You must specify a path to the firmware image file for the command to decode the image file's version.

Example usage:

```
./elxflash /iv=firmware/A11460.grp
1.1.46.0
```

3.12 Log (/log)

Usage: `/log=<logfile.txt>`

The `/log` switch appends the output of the Elxflash utility to a text file. `Log` can be used with any switch.

3.13 Preview (/p)

Usage: `/p`

The `/p` switch provides a download preview of all adapters the Elxflash utility can update using either auto-discovery or the `fwmatrix.txt` file. The preview switch can be used with any of the operational switches, such as:

- `/ff`
- `/fb`
- `/downgrade`
- `/rewrite`
- `/update`

When the preview switch is used, the Elxflash utility displays a download summary, but it does not actually perform the download.

NOTE If the `/p` switch is used, the `Status=<description>` field displays `Preview`.

Each adapter's download preview displays the adapter's old and new image versions. The old image version represents the image version that is currently on the adapter. The new image version represents the image version the Elxflash utility would use during a download.

Example usage:

`./elxflash /preview /auto /update` – Previews an upgrade of firmware or boot code using auto-discovery.

- The `fwmatrix.txt` file is ignored. The desired update versions of firmware or boot code must be in their respective directories.
- If the update versions are later than the currently installed versions on the adapter, the Elxflash utility provides a download preview for each adapter that can be updated.
- If multiple update versions of firmware or boot code are found for an adapter, the Elxflash utility provides a download preview using the most recent versions.

`./elxflash /preview` – Previews an upgrade of firmware or boot code using the `fwmatrix.txt` file.

- For each installed and supported adapter, the current firmware and boot code versions are compared with the versions specified in the `fwmatrix.txt` file.
- If the update versions in the `fwmatrix.txt` file are more recent than the currently installed versions, the Elxflash utility provides a download preview of firmware or boot code for each adapter that can be updated.

3.14 Process FC HBAs Only (/fc)

Usage: `/fc`

The `/fc` switch causes Elxflash to act only on FC HBAs, including LPe16202/OCe15100 adapters in FC mode. It does not act on LPe16202/OCe15100 adapters in NIC+FCoE mode.

Example usage:

`./elxflash /q /fc` – Only FC HBAs are displayed for a query.

`./elxflash /auto /up /fc` – This command applies auto update only to FC HBAs.

When a firmware or boot code update is performed and the `/fc` switch is used, only FC adapters are updated; LPe16202/OCe15100 adapters in NIC+FCoE mode are not displayed.

3.15 Process UCNAs (/ucna)

Usage: `/ucna`

The `/ucna` switch causes Elxflash to act only on LPe16202/OCe15100 adapters in NIC+FCoE mode.

Example usage:

`./elxflash /q /ucna` – Only LPe16202/OCe15100 adapters in NIC+ FCoE mode are displayed for a query.

`./elxflash /auto /up /ucna` – Apply auto update only to LPe16202/OCe15100 adapters in NIC+FCoE mode.

When a firmware/boot code update is performed and the `/ucna` switch is used, only LPe16202/OCe15100 adapters in NIC+FCoE mode are updated; FC HBAs are not displayed.

3.16 Query (/q)

Usage: `/q`

The `/q` switch displays an adapter's model, WWN or MAC address, PCI_ID, firmware version, and boot code version. Query can be used with any switch.

Example usage:

```
C:\elxflashStandalone-windows-10.0.567.22-1\win>elxflash.bat /q
```

NOTE On FC adapters, the query switch displays boot code version information only if the adapter has boot code installed.

3.17 Ramdrive (/ramdrive)

NOTE This command is available in the Windows Offline utility only.

Usage: `/ramdrive=<drive letter>[:]`

The `/ramdrive` switch specifies the drive on which to create temporary files and log files. The default is drive X; the default drive in Windows PE.

3.18 Rewrite (/rewrite or /e)

Usage: `/rewrite` -or- `/e`

The `/rewrite` switch updates the firmware or boot code of each adapter if the installed versions are earlier than, or the same as, the rewrite versions. This switch cannot be used with `/update` or `/downgrade`.

Example usage:

`./elxflash /rewrite /auto` – Rewrites the firmware or boot code using auto-discovery.

- The `fwmatrix.txt` file is ignored. The desired rewrite versions of firmware or boot code must be in their respective directories.
- If the rewrite versions are more recent than, or the same as, the versions installed on the adapter, the rewrite versions are downloaded to the adapter.
- If multiple rewrite versions of firmware or boot code are found for an adapter, the most recent versions are downloaded to the adapter.
- When performing the boot code rewrite operation, Elxflash first tries to match by adapter family and boot type. If a match is not found, Elxflash then tries to match by boot type. If the utility matches by boot type, and multiple versions of boot code are detected, rewrite always chooses the boot code in the following order:
  a. Universal (U)
  b. Pair (P)
  c. Open (O)
  d. EFI (E)
  e. x86 (B)

`./elxflash /rewrite` – Rewrites the firmware or boot code using the `fwmatrix.txt` file.

- For each installed and supported adapter, the current firmware or boot code versions are compared with the versions in `fwmatrix.txt` file.
- If the currently installed versions are less than or equal to the rewrite versions in `fwmatrix.txt` file, the rewrite versions of firmware or boot code are downloaded to the adapter.

3.19 Silent (/s)

Usage: `/s`

The `/s` switch prevents all output from being displayed.

Example usage:

`./elxflash /f /auto /s` – Forces a firmware and boot code download using auto-discovery and mutes all output to stdout.

- The `fwmatrix.txt` file is ignored. The desired versions of firmware and boot code must be in their respective directories.
- If multiple versions of firmware or boot code are found for an adapter, Elxflash uses the most recent versions when performing the firmware and boot code downloads.
- No output is printed to stdout.

`./elxflash /f /s` – Forces a firmware and boot code download using the `fwmatrix.txt` file and mutes all output to stdout.

- For each installed and supported adapter, a download of firmware and boot code is forced using the versions specified in the `fwmatrix.txt` file.
- No output is printed to stdout.

3.20 Discover CFAs Using SysFS (/sysfs) (LPe16202/OCe15100 Adapters Only)

NOTE The libsysfs library must be installed. If the libsysfs library is not found, the `/sysfs` switch is not available.

Usage: `/sysfs + operational switch`

`/sysfs` is an operational switch that discovers CFA (NIC) functions using the SysFS tool, and it discovers FC functions using the LpCfg tool. It also runs firmware downloads on CFAs using the ethtool utility or SysFS interface, and on FC adapters using the LpCfg utility. The `/sysfs` switch is used with any switch that displays adapter information; for example, `/q`, or any combination of switches that performs a download.

The following switches can be used:

- `/f` – Force firmware and boot code
- `/fb` – Force boot code
- `/ff` – Force firmware
- `/downgrade` or `/g`
- `/query` or `/q`
- `/rewrite` or `/e`
- `/update`

For example, `./elxflash /sysfs /auto /f` discovers CFA NIC functions using SysFS and discovers legacy FC functions using LpCfg. It forces a firmware download on CFA adapters using ethtool or SysFS and forces a firmware and boot code download on legacy FC adapters using LpCfg.

NOTE Only CFA NIC functions are displayed. Displaying FCoE functions is not supported.

When `/sysfs` is used, the VPD model name is not available.

The `/sysfs` switch is not supported by the Elxflash Stand-alone kit for VMware.

Example usage:

```
/elxflash /sysfs /q
```

3.21 Update (/update)

Usage: `/update`

The `/update` switch updates the firmware or boot code of each adapter if the currently installed versions are older than the update versions. This switch cannot be used with `/downgrade` or `/rewrite`.

Example usage:

`./elxflash /update /auto /q` – Upgrades the firmware or boot code using auto-discovery.

- The `fwmatrix.txt` file is ignored. The desired update versions of firmware or boot code must be in their respective directories.
- If the update versions are newer than the currently installed versions on the adapter, the update versions are downloaded to the adapter.
- If multiple update versions of firmware or boot code are found for an adapter, the most recent versions are downloaded to the adapter.
- When performing the boot code update operation, the Elxflash utility uses the most recent boot code version found. Because multiple compatible versions of boot code can exist, the most recent version is selected in the following order:
  a. Universal (U)
  b. Pair (P)
  c. Open (O)
  d. EFI (E)
  e. x86 (B)

NOTE The most recent boot code version is downloaded regardless of the installed boot code. This allows upgrading from one type of boot code to another type.

`./elxflash /update /q` – Upgrades the firmware or boot code using the `fwmatrix.txt` file.

- For each installed and supported adapter, the current firmware and boot code versions are compared with the versions specified in `fwmatrix.txt`.
- If the update versions in `fwmatrix.txt` are more recent than the currently installed versions, the update versions of firmware or boot code are downloaded to that adapter.

3.22 Verbose (/v)

Usage: `/v`

The `/v` switch displays progress messages, and it can be used with any switch. When this switch is used, the following information is displayed:

- A download summary for all adapters that had successful or failed downloads
- A summary of unsupported adapters, if applicable
- A per-adapter message for each adapter the Elxflash utility did not update

3.23 VPD (/vpd)

Usage: `/vpd`

The `/vpd` command displays the VPD model name for supported adapters. You must use the `vpd` command with one of the following switches:

- `/f` – Force firmware and boot
- `/fb` – Force boot code
- `/ff` – Force firmware

- `/xml` – XML output
- `/downgrade` or `/g`
- `/query` or `/q`
- `/rewrite` or `/e`
- `/update`

NOTE The `/vpd` command does not display any information if used without the switches listed previously. If you attempt to use the `/vpd` command without the switches listed previously, an error message is displayed.

If the `/vpd` command is used with the `fwmatrix.txt` file, the model name in the `fwmatrix.txt` file must match the VPD model name reported by Elxflash. Supported adapters include legacy FC adapters that support VPD.

3.24 XML Output (/xml)

Usage: `/xml`

The `/xml` switch displays utility output in XML format.

Chapter 4: LpCfg Utility Command Line Interface

The LpCfg utility allows you to configure Emulex adapters before you install or boot a server operating system.

The LpCfg utility includes:

- linlpcfg
- winlpcfg

NOTE Usage code and examples show *os*`lpcfg`. Note that the *os* designation changes to `./lin` if you are using LpCfg for Linux or VMware and to `win` if you are using LpCfg for Windows.

You can use the LpCfg utility to do the following:

- View information about an Emulex adapter
- Reset the adapter
- Return the adapter to its factory default settings
- Download firmware and boot code files
- Select a boot device
- Read and update WWNs
- Read MAC addresses
- Enable boot code
- Update configuration regions
- Set the adapter to use soft jumpers
- Run diagnostic tests
- Read and process script files
- Read, write, and reset VLAN IDs and VLAN Priorities

4.1 Running the LpCfg Utility from the Command Prompt

NOTE The LpCfg utility is a command line utility and does not include a graphical user interface.

To run the LpCfg utility from the command prompt:

1. Boot the system with a supported operating system.
2. Start the LpCfg utility with a valid command or a valid script file name.
   - To start the LpCfg utility from the command line, move to the directory where the executable file resides and type:

     os`lpcfg <`*valid command*`>`

     Replace *os* with `./lin` or `win` as appropriate.
     - For example, starting the Linux or VMware LpCfg utility with a `reset` command:

       `./linlpcfg reset n=2`
     - For example, starting the Windows LpCfg utility with a `reset` command:

       `winlpcfg reset n=2`

— To start the LpCfg utility with a script file name, move to the directory where the LpCfg utility resides and type:

`oslpcfg @<script file name>`

- For example, starting the Windows LpCfg utility with `script1.txt` in the `c:\test` directory:

```
winlpcfg @C:\test\script1.txt
```

NOTE To redirect screen output to a file, add `<filename>` at the end of each command.
For example:

`oslpcfg listboot n=1 >result.out`

For more information on script files, see Section 4.3.8, Using Script Files.

4.2 Running the LpCfg Utility Included with the Elxflash Stand-alone Kit

The Elxflash Stand-alone kit allows you to run Elxflash and LpCfg without installing the utilities. No drivers are included in the Stand-alone kit. No applications are installed, because the utility runs from inside the kit.

The Elxflash Stand-alone kit has the same capability as the Elxflash Offline kit, except you use the following scripts to run the utility:

- **Linux** and **VMware** – `linplcfg.sh` and `elxflash.sh`
- **Windows** – `winlpcfg.bat` and `elxflash.bat`

4.2.1 Linux and VMware

NOTE Running 32-bit applications on Linux x86_64 is not supported. The `sh` scripts run only native executable files.

1. Extract the kit contents.
2. Change directory (`cd`) to `ElxflashStandalone-linux-<version>`.

   The following directories must be present:

   — `boot\`
   — `firmware\`
   — `lx\`
3. For Elxflash, copy the firmware images to the `firmware` directory. Copy the boot images to the `boot` directory.

   NOTE For Elxflash, each operating system architecture directory includes an `fwmatrix.txt` file. You must use the `fwmatrix.txt` directory that matches the current operating system architecture.
4. Change directory (`cd`) to the `lx` directory.

   The following files and directories must be present:

   — `i386\`
   — `x86_64\`
   — `ppc64\`
   — `elxflash.sh`
   — `linlpcfg.sh`

   The `elxflash.sh` script configures the environment. Run Elxflash, and revert any changes before exiting. This script installs the MILI daemons if they are not already installed.

The `linlpcfg.sh` script configures the environment. Run linlpcfg, and revert any changes before exiting. This script installs the MILI daemons if they are not already installed.

NOTE If OneCommand Manager is installed, the OneCommand Manager libraries are used by the Elxflash and linlpcfg utilities.

The `linlpcfg.sh` and `elxflash.sh` scripts call the native versions of Elxflash and linlpcfg. For example, on Linux x86_64, the 64-bit utilities are called.

The adapter being managed by the utilities included in the Elxflash Stand-alone kit must not be managed simultaneously by other Emulex utilities, including OneCommand Manager.

5. To update firmware and boot code on an adapter, type:

```
# ./elxflash.sh /auto up
```

To display a list of HBAs, type:

```
# ./linlpcfg.sh  listhba
```

Each time Elxflash or LpCfg is run, a log file is created. On Linux these files are called:

```
/var/log/clu/elxflash.log
/var/log/clu/linlpcfg.log
```

4.2.2 Windows

1. Extract the kit contents.
2. Change directory (`cd`) to `Elxflash Standalone-windows-<version>`.

   The following directories must be present:

   — `boot\`
   — `firmware\`
   — `win\`
3. For Elxflash, copy the firmware images to the `firmware` directory. Copy the boot images to the `boot` directory.

   NOTE For Elxflash, each operating system architecture directory includes an `fwmatrix.txt` file. You must use the `fwmatrix.txt` directory that matches the current operating system architecture.
4. Change directory (`cd`) to the `win` directory.

   The following files and directories must be present:

   — `win32\`
   — `x64\`
   — `elxflash.bat`
   — `winlpcfg.bat`

   The `elxflash.bat` script configures the environment. Run Elxflash, and revert any changes before exiting. This script installs the MILI service if it is not already installed.

   The `winlpcfg.bat` script configures the environment. Run `winlpcfg`, and revert any changes before exiting. This script installs the MILI service if it is not already installed.

   NOTE The MILI service is temporarily installed. The service is removed after the script runs.

   If OneCommand® Manager is installed, the OneCommand Manager libraries and OneCommand Manager MILI service are used by the Elxflash and winlpcfg utilities.

The `winlpcfg.bat` and `elxflash.bat` scripts call the native versions of `elxflash.exe` and `winlpcfg.exe`. For example, on Windows x64, the 64-bit utilities are called.

The adapter being managed by the utilities included in the Elxflash Stand-alone kit must not be managed simultaneously by other Emulex utilities, including OneCommand Manager.

5. To update firmware and boot code on an adapter in Windows, type:

```
# elxflash.bat /auto /up
```

To display a list of HBAs in Windows, type from the `win` directory:

```
# winlpcfg.bat  listhba
```

Each time Elxflash or LpCfg is run, a log file is created. On Windows these files are called:

```
C:\clu\log\elxflash.log
C:\clu\log\winlpcfg.log
```

4.3 Supported Commands

Table 1 lists all of the LpCfg commands that are supported on Emulex adapters on various platforms.

- ✓ indicates commands that are supported on both winlpcfg and linlpcfg.
- L indicates commands that are supported only on linlpcfg.
- *FCoE only* indicates commands that are supported only by LPe16202/OCe15100 adapters in FCoE mode.
- *NIC only* indicates commands that are supported by LPe16202/OCe15100 adapters in NIC mode.

Table 1 Supported Commands for LpCfg1

Commands	FC Adapters			LPe16202/OCe15100 Adapters in NIC+FCoE Mode	
	x86	x64	PPC	x86	x64
Operating Systems	RHEL 6.7+	RHEL 6.7+ RHEL 7.2+	RHEL 6.7+ RHEL 7.2+	RHEL 6.7+	RHEL 6.7+ RHEL 7.2+
	SLES 11 SP3+	SLES 11 SP3+ SLES 12 SP2	SLES 11 SP3+	SLES 11 SP3+	SLES 11 SP3+ SLES 12 SP2
	WinPE 5+	WinPE 5+		N/A	N/A
`config`	✓	✓	L	FCoE only	FCoE only
`directDownload` Supported on LPe12000-series adapters only.	L	L	L	N/A	N/A
`disableboot`	✓	✓	L	N/A	N/A
`disablebootdevice`	✓	✓	L	FCoE	FCoE
`download` `a=<adapter name>` is supported on LPe12000-series adapters only.	✓	✓	L	✓	✓
`enableboot`	✓	✓	L	N/A	N/A
`enablebootdevice`	✓	✓	L	FCoE	FCoE
`extloopback`	✓	✓	L	✓	✓
`factorydefaults` Not supported on LPe16202/OCe15100 adapters in NIC+FCoE mode.	✓	✓	✓	N/A	N/A

Table 1 Supported Commands for LpCfg1 (Continued)

Commands	FC Adapters			LPe16202/OCe15100 Adapters in NIC+FCoE Mode	
	x86	x64	PPC	x86	x64
hbaattr	✓	✓	✓	✓	✓
h (Help)	✓	✓	L	✓	✓
intloopback	✓	✓	L	✓	✓
listboot	✓	✓	L	N/A	N/A
listhba	✓	✓	L	✓	✓
listrev	✓	✓	L	FCoE only	FCoE only
listmac	N/A	N/A	N/A	NIC only	NIC only
listwwn	✓	✓	L	FCoE only	FCoE only
logfile	✓	✓	L	✓	✓
networkboot	N/A	N/A	N/A	✓	✓
pciloopback	✓	✓	L	FCoE only	FCoE only
posttest Supported on LPe12000-series adapters only.	✓	✓	L	N/A	N/A
readaltboot	✓	✓	L	FCoE only	FCoE only
readbootdevice	✓	✓	L	FCoE only	FCoE only
readconfig	✓	✓	L	FCoE only	FCoE only
readmac	N/A	N/A	N/A	✓	✓
readvlanprops	N/A	N/A	N/A	✓	✓
reset	✓	✓	L	FCoE only	FCoE only
restoredefsvmac[1]	N/A	N/A	N/A	✓	✓
restoredefwwn	✓	✓	L	FCoE only	FCoE only
restorenwwn	✓	✓	L	FCoE only	FCoE only
restorevlanprops	N/A	N/A	N/A	✓	✓
restorewwn	✓	✓	L	FCoE only	FCoE only
savewwn	✓	✓	L	FCoE only	FCoE only
screendisplay	✓	✓	L	FCoE only	FCoE only
scriptvwwnn	✓	✓	L	FCoE only	FCoE only
scriptvwwpn	✓	✓	L	FCoE only	FCoE only
scriptwwnn	✓	✓	L	FCoE only	FCoE only
scriptwwpn	✓	✓	L	FCoE only	FCoE only
setaltboot	✓	✓	L	FCoE only	FCoE only
setbootdevice	✓	✓	L	FCoE only	FCoE only
version	✓	✓	L	✓	✓
vpd	✓	✓	L	✓	✓
writesvmac[a]	N/A	N/A	N/A	✓	✓
writevlanprops	N/A	N/A	N/A	✓	✓
writewwn	✓	✓	L	FCoE only	FCoE only

1. This command is supported only on certain OEM-specific adapters.

4.3.1 Help Command and General Adapter Management Commands

Use the following commands to view the online help, reset the adapter, or reset the adapter to factory defaults. You can also run a POST.

4.3.1.1 Viewing the Syntax for Commands (help)

To view the syntax for all available commands, type:

```
oslpcfg help
```

To view the syntax for a specific command, type:

```
oslpcfg help <command>
```

For example:

```
oslpcfg help download
```

returns a response similar to the following:

- For Linux and VMware:

  ```
  download n=<adapter> i=<path/image_filename>
  ```

  or

  ```
  download a=<adaptertype> i=<path/image_filename>
  ```

- For Windows:

  ```
  download n=<adapter> i=<path\image_filename>
  ```

  or

  ```
  download a=<adaptertype> i=<path\image_filename>
  ```

4.3.1.2 Resetting an Adapter (reset)

This command resets a specific adapter or all adapters in the system.

ATTENTION Executing a reset on an FC adapter that is being used to boot from SAN is not recommended. The reset might cause a loss of connectivity to the SAN and possible loss of data. To reset an FC adapter, make sure the adapter is not currently being used to boot from SAN. There are two ways to do this:

- Move the adapter you want to reset to a non-boot-from-SAN host and reset it from there.
- If the host with the reset target adapter is also hosting other boot-from-SAN adapters, carry out a boot from SAN using one of the other boot-from-SAN adapters. The intended target adapter can now be reset, because it is not being actively used for boot from SAN.

To reset one adapter, type:

```
oslpcfg reset n=<adapter number>
```

To reset all adapters in the system, type

```
oslpcfg reset n=all
```

4.3.1.3 Resetting an Adapter to the Factory Defaults (factorydefaults)

This command returns the adapter to the factory default profile and configuration. A reboot is required for the changes to take effect.

NOTE The `factorydefaults` command is not supported on LPe16202/OCe15100 adapters in NIC+FCoE mode.

After a successful factory reset on LPe32000-series, LPe31000-series, or LPe16000-series adapters, the adapter is offline. Perform an immediate reboot to complete the reset and return the adapter to full functionality. An LPe12000-series adapter is online after the reset is complete and does not require a reboot.

ATTENTION Restoring the factory defaults on an FC adapter that is being used to boot from SAN is not recommended. The `restore` command might cause a loss of connectivity to the SAN and possible loss of data. To restore the factory defaults on an FC adapter, make sure the adapter is not currently being used to boot from SAN. There are two ways to do this:

- Move the adapter on which you want to restore the defaults to a non-boot-from-SAN host and perform the restore defaults command from there.
- If the host with the restore defaults target adapter is also hosting other boot-from-SAN adapters, carry out a boot from SAN using one of the other boot-from-SAN adapters. The factory defaults can now be restored on the intended target adapter, because it is not being actively used for boot from SAN.

To return an adapter to its default settings, type:

```
oslpcfg factorydefaults n=<adapter#|all>
```

The following example returns adapter number 1 to its default settings.

```
oslpcfg factorydefaults n=1
```

4.3.1.4 Running a Power-on Self-Test (posttest)

This command runs a POST on the selected adapter.

ATTENTION Running a POST on an FC adapter that is being used to boot from SAN is not recommended. The `posttest` command might cause a loss of connectivity to the SAN and possible loss of data. To run a POST on an FC adapter, make sure the adapter is not currently being used to boot from SAN. There are two ways to do this:

- Move the adapter on which you want to run POST to a non-boot-from-SAN host and run POST from there.
- If the host with the `posttest` target adapter is also hosting other boot-from-SAN adapters, carry out a boot from SAN using one of the other boot from SAN adapters. The POST can now be run on the intended target adapter, because it is not being actively used for boot from SAN.

To run the adapter POST, type:

```
oslpcfg posttest n=<all|adapter number>
```

The following example runs a POST on adapter number 1.

```
oslpcfg posttest n=1
```

4.3.2 Viewing Adapter Information

Use the following commands to view different kinds of adapter information, such as VPD, boot device information, and adapter attributes.

4.3.2.1 Viewing Emulex Conventional Names Instead of VPD (/c)

Adding `/c` to any command that uses the `a=` parameter to return adapter data causes the command to return an Emulex conventional model for the adapter rather than the model name of the adapter. `/c` can be used for any command that uses the model name. For instance, the following commands use the `a=` parameter and show this behavior.

- `listhba`
- `config`
- `download`
- `directdownload`

NOTE The offline utility does not always display the model name of the adapter being tested.

In the following example, the VPD model name is reported for each adapter:

```
# ./linlpcfg listhba
Command: listhba
HBA 1: 10000090 FA942ECA Functional FW: 11.2.50.48 devID:E300 Bus:8 Dev:0 Func:0
LPe32000-M2-D
HBA 2: 10000000 C9A18000 Functional FW: US2.02A1  devID:F100 Bus:E Dev:0 Func:0
LPe12002-M8
HBA 3: 10000000 C9A18001 Functional FW: US2.02A1  devID:F100 Bus:E Dev:0 Func:1
LPe12002-M8
```

To list adapter information using the conventional name rather than the VPD, type

```
# ./linlpcfg listhba /c
```

The following information is returned:

```
Command: listhba /c
HBA 1: 10000090 FA942ECA Functional FW: 11.2.50.48 devID:E300 Bus:8 Dev:0 Func:0
LPe32000
HBA 2: 10000000 C9A18000 Functional FW: US2.02A1  devID:F100 Bus:E Dev:0 Func:0
LPe12002
HBA 3: 10000000 C9A18001 Functional FW: US2.02A1  devID:F100 Bus:E Dev:0 Func:1
LPe12002
```

The output from this command includes the Emulex conventional model name – LPe12002, rather than the VPD model name – LPe12002-M8.

4.3.2.2 Viewing the LpCfg Utility Version Information (version)

This command shows the LpCfg utility version information.

To view this information, type:

```
oslpcfg version
```

4.3.2.3 Viewing VPD (vpd)

This command shows the VPD of the adapter specified by its number.

To display VPD, type:

```
oslpcfg vpd n=<adapter number>
```

4.3.2.4 Viewing Boot Device Information (readbootdevice)

This command shows the WWN, the LUN (in decimal format), and the topology in use for the currently selected boot device.

To show this information, type:

```
oslpcfg readbootdevice n=<adapter number>
```

The following example reads WWN and LUN for adapter number 1:

```
oslpcfg readbootdevice n=1
```

4.3.2.5 Viewing Boot Code Versions (listboot)

This command lists all the boot code versions that are loaded in the flash of the adapter. If the selected adapter does not have boot code, it returns error code 39.

To list boot code versions, type:

```
oslpcfg listboot n=<adapter number>
```

The following example lists boot code versions that are loaded on adapter number 3:

```
oslpcfg listboot n=3
```

4.3.2.6 Viewing Adapter Attributes (hbaattr)

This command displays adapter information.

To list the adapter attributes for all installed adapters, type:

```
oslpcfg hbaattr
```

To list the adapter attributes for one adapter, type

```
oslpcfg hbaattr n=<adapter_number>
```

4.3.2.7 Viewing All Adapters in the System (listhba)

This command lists all installed adapters in the system. Information includes the adapter number, the IEEE address assigned by the manufacturer, the functional firmware, the adapter type, and possible mailbox errors.

To list all adapters in the system, type:

```
oslpcfg listhba
```

NOTE `listhba` with option `/c` displays the conventional model names instead of the model names contained in the VPD.

4.3.2.8 Viewing the WWN of All Adapters in the System (listwwn)

This command lists all adapters installed in the system and shows the factory-assigned WWN, the nonvolatile WWPN, and the WWNN used to identify an adapter in the SAN.

The factory-assigned WWN is an IEEE address that cannot be changed in the field. The nonvolatile WWN can be modified in the field and persists after a restart of the operating system. The full factory-assigned WWN and nonvolatile WWN are a concatenation of the two 8-character values (word 0 and word 1) that are shown for each. You can modify the nonvolatile WWPN and WWNN using either the `writewwn` command or the `scriptwwpn` and `scriptwwnn` commands. For more information on the `writewwn` command, see Section 4.3.4.1, Writing WWN and Updating NVPARMS (writewwn).

If the system does not have any Emulex adapters installed, it returns error code 45.

To show the WWN information, type:

```
oslpcfg listwwn
```

4.3.2.9 Viewing the MAC Address (listmac)

This command shows the MAC address of a NIC port on an LPe16202/OCe15100 adapter in NIC+FCoE mode.

To view the MAC address of a NIC port, type:

```
oslpcfg listmac n=<adapter number>
```

4.3.2.10 Reading the MAC Address (readmac)

This command shows the current, factory, and semivolatile MAC addresses for the adapter number specified.

The `readmac` command is supported only on LPe16202/OCe15100 adapters in NIC+FCoE mode. To view the MAC address of the adapter number specified, type:

```
oslpcfg.exe readmac n=<adapter number>
```

The following example lists information for adapter number 1:

```
>oslpcfg.exe readmac n=1
Command: readmac n=1
adapter 1:
Current MAC Address : 00-90-FA-30-43-AA
Factory MAC Address : 00-90-FA-30-43-30
Semi-Volatile MAC Address: 00-90-FA-30-43-AA
Command completed, NO Error
```

4.3.2.11 Viewing Firmware Program Revisions (listrev)

This command shows the firmware versions in the adapter's flash memory, specified by their numbers.

To show revisions, type:

```
oslpcfg listrev n=<adapter number>
```

The following example lists information for adapter number 3:

```
oslpcfg listrev n=3
```

4.3.2.12 Viewing Selected Configuration Regions (readconfig)

This command shows the contents of the selected configuration region up to the initialized length or the specified byte count (if the initialized length is less than the specified byte count). Valid region numbers are 0 to 32. You must initialize the configuration region first by writing data to it.

NOTE The `readconfig` command only supports reading configuration regions 0, 8, and 32 on the following adapters:

- LPe16000-series adapters
- LPe31000-series adapters
- LPe32000-series adapters

To read a configuration, type:

```
oslpcfg readconfig n=<adapter number> r=<region number> l=<byte count>
```

The following example reads the configuration for adapter number 1, region 0, byte count 20:

```
oslpcfg readconfig n=1 r=0 l=20
```

4.3.3 Firmware and Boot Code Download Commands

The following firmware and boot code download commands include a command to download a firmware or boot code file and a command to access the flash device directly.

4.3.3.1 Downloading a File (download)

This command downloads a firmware or boot code file to a specific adapter.

ATTENTION Downloading a firmware or boot code file to an FC adapter that is being used to boot from SAN is not recommended. The download command might cause a loss of connectivity to the SAN and possible loss of data. To download a firmware or boot code file to an FC adapter, make sure the adapter is not currently being used to boot from SAN. There are two ways to do this:

- Move the adapter on which you want to download the file onto a non-boot-from-SAN host and perform the download command from there.
- If the host with the download target adapter is also hosting other boot-from-SAN adapters, carry out a boot from SAN using one of the other boot-from-SAN adapters. The firmware or boot code file can now be downloaded onto the intended target adapter, because it is not being actively used for boot from SAN.

The adapter name is the name that appears when you run the `listhba` command. For more information on the `listhba` command, see Section 4.3.2.7, Viewing All Adapters in the System (listhba).

To download a firmware image file to an adapter specified by its number, type:

```
oslpcfg download n=<adapter number> i=<firmware image filename>
```

The following example downloads the `uu513a10.prg` boot code file to adapter number 6; in this example, the boot code file is for an LPe12000 adapter:

```
oslpcfg download n=6 i=uu513a10.prg
```

NOTE Boot code on LPe16000-series, LPe31000-series, and LPe32000-series adapters is updated by downloading firmware. LpCfg does not support updating boot code separately for these adapters.

4.3.3.2 Accessing the Flash Device Directly (directdownload)

This command accesses the flash device on the adapter directly without using the adapter firmware. This feature is useful in downloading a ROM file image if the firmware has been corrupted. The adapter name is the name that appears when you run the `listhba` command. You can also use `default` for the adapter name if only one single-port adapter or one dual-port adapter is in the system. For more information on the `listhba` command, see Section 4.3.2.7, Viewing All Adapters in the System (listhba).

NOTE You cannot use `directdownload` in a script file.

This command is supported only on LPe12000-series adapters.

ATTENTION Accessing the flash device directly on an FC adapter that is being used to boot from SAN is not recommended. The `directdownload` command might cause a loss of connectivity to the SAN and possible loss of data. To run the `directdownload` command on an FC adapter, make sure the adapter is not currently being used to boot from SAN. There are two ways to do this:

- Move the adapter on which you want to run the `directdownload` command to a non-boot-from-SAN host and run the command from there.
- If the host with the directdownload target adapter is also hosting other boot-from-SAN adapters, carry out a boot from SAN using one of the other boot-from-SAN adapters. The `directdownload` command can now be run on the intended target adapter, because it is not being actively used for boot from SAN.

To access the flash device on the adapter directly, type one of the following (all on one line):

In Windows:

```
oslpcfg directdownload a=<adapter name/default> i=<path\image_filename>
s=<selection 0 or 1>
```

In Linux:

```
oslpcfg directdownload a=<adapter name/default> i=<path/image_filename>
s=<selection 0 or 1>
```

where `s=1` saves the existing VPD.

The following example accesses the flash device on an LPe12000 HBA in Windows:

```
winlpcfg directdownload a=lpe12000 i=C:\image\ud100a8.rom s=1
```

The following example accesses the flash device if the offline utility cannot detect the adapter type and only one single-port adapter or one dual-port adapter is in the Windows system:

```
winlpcfg directdownload a=default i=C:\image\ud100a8.rom s=1
```

ATTENTION If `s=0`, the ROM images used with the `directdownload` command might not contain certain VPD information; for example, serial number, adapter model, or manufacturer. Direct download of a ROM image that has not been confirmed to contain the correct VPD image updates the board's firmware, but it also clears the VPD. The board will function. If you use calls for VPD in your applications, the information might be changed or missing.

4.3.4 World Wide Name Commands

The following commands allow you to use the WWN to update NVPARAMS, save WWN data to a file, and restore WWN data while updating NVPARAMS. You can also restore the NVPARAMS and the IEEE address.

4.3.4.1 Writing WWN and Updating NVPARMS (writewwn)

This command allows you to enter word 0 and word 1 of the WWPN or WWNN from the keyboard or from a barcode scanner to update a specified adapter's NVPARMS with a new WWPN or WWNN. The new WWPN and WWNN are used the next time the adapter is discovered. The adapter stores the original WWPN and WWNN in another region of the memory so it can be used to identify the adapter as it was manufactured. The WWN can also be read with a barcode scanner.

ATTENTION Running a `writewwn` command on an FC adapter that is being used to boot from SAN is not recommended. The `writewwn` command might cause a loss of connectivity to the SAN and possible loss of data. To write to the WWN and Update NVPARMS on an FC adapter, make sure the adapter is not currently being used to boot from SAN. There are two ways to do this:

- Move the adapter on which you want to use the `writewwn` command to a non-boot-from-SAN host and run `writewwn` from there.
- If the host with the `writewwn` target adapter is also hosting other boot-from-SAN adapters, carry out a boot from SAN using one of the other boot-from-SAN adapters. You can now write the WWN or update the NVPARMS on the intended target adapter, because it is not being actively used for boot from SAN.

The `writewwn` command prompts you for the WWPN and WWNN data words, so it cannot be used in a script file. The `scriptwwnn` and `scriptwwpn` commands use values entered with the command, so they can be used in a script file.

ATTENTION Use the `writewwn` command with caution. If you use the same WWPN or WWNN on more than one adapter in a fabric, unpredictable results might occur.

NOTE Word 0 of WWNN and WWPN names must follow one of the following formats:

```
1 0 0 0 0 x x x
2 x x x x x x x
3 x x x x x x x
5 x x x x x x x
```

To modify the WWPN and WWNN, type:

```
oslpcfg writewwn n=<adapter number>
```

The offline utility prompts you to enter new data:

- WWPN word 0
- WWPN word 1
- WWNN word 0
- WWNN word 1

The following example writes the WWPN and WWNN for adapter number 1.

```
oslpcfg writewwn n=1
Enter or Scan value for WWPN word 0 now
```

Enter a value.

```
10000000
```

The system echoes what you entered, followed by the next prompt:

```
10000000
Enter or Scan value for WWPN word 1 now
```

4.3.4.2 Saving WWN Data to a File (savewwn)

This command reads the original words 0 and 1 of the IEEE address, installed by manufacturing, from configuration regions 16 (or 32) of the adapter (specified by its number), and it saves the configuration region information in the selected WWN file.

To save the WWN data to a file, type:

```
oslpcfg savewwn n=<adapter number> c=<wwn filename>
```

The following example reads the configuration region information on adapter number 4 and saves it to the contents of the `ctwwn.sav` file:

```
oslpcfg savewwn n=4 c=ctwwn.sav
```

4.3.4.3 Restoring WWN and Updating NVPARMS (restorewwn)

This command restores words 0 and 1 of the IEEE address from a specified file created with the `savewwn` command and uses them to update the NVPARMS port name with this IEEE address.

ATTENTION Running a `restorewwn` command on an FC adapter that is being used to boot from SAN is not recommended. The `restorewwn` command might cause a loss of connectivity to the SAN and possible loss of data. To restore the WWN and Update NVPARMS on an FC adapter, make sure the adapter is not currently being used to boot from SAN. There are two ways to do this:

- Move the adapter on which you want to use the `restorewwn` command to a non-boot-from-SAN host and run `restorewwn` from there.
- If the host with the `restorewwn` target adapter is also hosting other boot-from-SAN adapters, carry out a boot from SAN using one of the other boot-from-SAN adapters. You can now restore the WWN and update the NVPARMS on the intended target adapter, because it is not being actively used for boot from SAN.

To restore the WWN, type:

```
oslpcfg restorewwn n=<adapter number> c=<wwn filename>
```

The following example updates the NVPARMS on adapter number 4 with the `ctwwn.sav` file.

```
oslpcfg restorewwn n=4 c=ctwwn.sav
```

4.3.4.4 Restoring NVPARMS (restorenvwwn)

ATTENTION Running a `restorenvwwn` command on an FC adapter that is being used to boot from SAN is not recommended. The `restorenvwwn`

command might cause a loss of connectivity to the SAN and possible loss of data. To restore the NVPARMS on an FC adapter, make sure the adapter is not currently being used to boot from SAN. There are two ways to do this:

- Move the adapter on which you want to use the `restorenvwwn` command to a non-boot-from-SAN host and run the command from there.
- If the host with the `restorenvwwn` target adapter is also hosting other boot-from-SAN adapters, carry out a boot from SAN using one of the other boot-from-SAN adapters. You can now restore the NVPARMS on the intended target adapter, because it is not being actively used for boot from SAN.

This command restores the nonvolatile WWPN and WWNN to the adapter, replacing any volatile WWPN and WWNN data, without powering off the adapter.

If the adapter does not have firmware that supports the volatile WWN, the following error message appears:

```
Write Volatile Parms Error. Reported Error 48
```

If this occurs, install firmware that supports the volatile WWN.

To restore the nonvolatile WWN, type:

```
oslpcfg restorenvwwn n=<adapter number>
```

The following example restores the nonvolatile WWN on adapter number 2.

```
oslpcfg restorenvwwn n=2
```

4.3.4.5 Restoring the IEEE Address (restoredefwwn)

This command reads the IEEE address (assigned by the manufacturer) and writes it to the nonvolatile WWPN and WWNN.

ATTENTION Running a `restoredefwwn` command on an FC adapter that is being used to boot from SAN is not recommended. The `restoredefwwn` command might cause a loss of connectivity to the SAN and possible loss of data. To restore the IEEE address on an FC adapter, make sure the adapter is not currently being used to boot from SAN. There are two ways to do this:

- Move the adapter on which you want to use the `restoredefwwn` command to a non-boot-from-SAN host and run `restoredefwwn` from there.
- If the host with the `restoredefwwn` target adapter is also hosting other boot-from-SAN adapters, carry out a boot from SAN using one of the other boot-from-SAN adapters. You can now restore the IEEE address on the intended target adapter, because it is not being actively used for boot from SAN.

To restore the IEEE address, type:

```
oslpcfg restoredefwwn n=<adapter number>
```

The following example restores the IEEE address on adapter number 2.

```
oslpcfg restoredefwwn n=2
```

4.3.5 Boot Code Commands

NOTE Offline utility boot commands apply only to x86 BootBIOs. They do not apply to UEFIBoot.

You must enable boot code before you can issue the `setBootDevice` and `setAltBoot` commands. If necessary, use the `enableboot` command to enable the boot code. See Section 4.3.5.1, Enabling or Disabling Boot Code (enableboot or disableboot), for more information.

To set the boot device with the offline utility, run the following commands in this order:

1. Use the `listboot` command to verify that the boot code is present. See Section 4.3.2.5, Viewing Boot Code Versions (listboot), for more information.

   Example:

   ```
   oslpcfg listboot n=1
   ```

   Sample return for an LPe12000-series adapter with x86 BootBIOS enabled:

   ```
   BootBIOS 1 (enabled): UU11.40A2
   Command completed, NO Error
   ```

   Sample return for an LPe12000-series adapter with x86 BootBIOS disabled:

   ```
   BootBIOS 1 (disabled): UU11.40A2
   Command completed, NO Error
   ```

   Sample return for an LPe16000-series, LPe31000-series, or LPe32000- series adapter with x86 BootBIOS enabled:

   ```
   Boot Code (enabled)          = 11.4.120.0
   fc_universal_bios_version    = 11.4.120.0
   fc_x86_bios_version          = 11.4.113.0
   fc_efi_bios_version          = 11.4.120.0
   fc_fcode_version             = 11.4.13.0
   Command completed, NO Error
   ```

   Sample return for an LPe16000-series, LPe31000-series, or LPe32000-series adapter with x86 BootBIOS disabled:

   ```
   Boot Code (disabled)         = 11.4.120.0
   fc_universal_bios_version    = 11.4.120.0
   fc_x86_bios_version          = 11.4.113.0
   fc_efi_bios_version          = 11.4.120.0
   fc_fcode_version             = 11.4.13.0
   Command completed, NO Error
   ```

2. Enable the boot code using the `enableboot` command. See Section 4.3.5.1, Enabling or Disabling Boot Code (enableboot or disableboot), for more information.
3. Use the `setbootdevice` command to configure the boot device. See Section 4.3.5.2, Selecting a Boot Device (setbootdevice), for more information.
4. Enable the boot device that was configured in step 3 using the `enablebootdevice` command. See Section 4.3.5.3, Enabling or Disabling Boot Devices (enablebootdevice or disablebootdevice), for more information.
5. Run the `readbootdevice` command to verify the configuration. See Section 4.3.2.4, Viewing Boot Device Information (readbootdevice), for more information.
6. Configure the system BIOS so the adapter boot device is the highest in the boot order.

4.3.5.1 Enabling or Disabling Boot Code (enableboot or disableboot)

This command enables or disables the boot code for the specified adapter.

ATTENTION Enabling or disabling the boot code on an FC adapter that is being used to boot from SAN is not recommended. The `enableboot` or

`disableboot` command might cause a loss of connectivity to the SAN and possible loss of data. To run either of these commands on an FC adapter, make sure the adapter is not currently being used to boot from SAN. There are two ways to do this:

- Move the adapter on which you want to enable or disable the boot code to a non-boot-from-SAN host and run the command from there.
- If the host with the target adapter is also hosting other boot-from-SAN adapters, carry out a boot from SAN using one of the other boot-from-SAN adapters. The `enableboot` or `disableboot` command can now be run on the intended target adapter, because it is not being actively used for boot from SAN.

To enable boot code, type:

```
oslpcfg enableboot n=<adapter number> i=1
```

NOTE `i` is always set to `1`.

The following example enables boot code on adapter number 6:

```
oslpcfg enableboot n=6 i=1
```

To disable boot code, type:

```
oslpcfg disableboot n=<adapter number>
```

The following example disables boot code on adapter number 6:

```
oslpcfg disableboot n=6
```

4.3.5.2 Selecting a Boot Device (setbootdevice)

This command sets the boot device specified by its WWN, LUN, and desired topology.

- Set `t=0` for arbitrated loop.
- Set `t=1` for point-to-point.

The selected device boots when the system reboots.

ATTENTION Selecting a boot device on an FC adapter that is being used to boot from SAN is not recommended. The `setbootdevice` command might cause a loss of connectivity to the SAN and possible loss of data. To run the `setbootdevice` command on an FC adapter, make sure the adapter is not currently being used to boot from SAN. There are two ways to do this:

- Move the adapter on which you want to run the `setbootdevice` command to a non-boot-from-SAN host and run the command from there.
- If the host with the `setbootdevice` target adapter is also hosting other boot-from-SAN adapters, carry out a boot from SAN using one of the other boot-from-SAN adapters. The `setbootdevice` command can now be run on the intended target adapter, because it is not being actively used for boot from SAN.

The boot code must be enabled before issuing the `setbootdevice` command.

To set the boot device, type (all on one line):

```
oslpcfg setbootdevice n=<adapter number> w0=<wwpn word 0> w1=<wwpn word 1>
l=<Decimal ID of LUN> t=<topology>
```

NOTE Enter the LUN in decimal format.

The following example sets the boot device on adapter number 1, LUN number 46, with a desired topology of arbitrated loop:

```
oslpcfg setbootdevice n=1 w0=a1b2c3d4 w1=b946a4e8 l=46 t=0
```

NOTE If port login fails after 50 ms, the command is retried once.

4.3.5.3 Enabling or Disabling Boot Devices (enablebootdevice or disablebootdevice)

After using the `setbootdevice` command, you can enable or disable the boot device by using the `enablebootdevice` or `disablebootdevice` command.

ATTENTION Enabling or disabling the boot device on an FC adapter that is being used to boot from SAN is not recommended. The `enablebootdevice` or `disablebootdevice` command might cause a loss of connectivity to the SAN and possible loss of data. To run either command on an FC adapter, make sure the adapter is not currently being used to boot from SAN. There are two ways to do this:

- Move the adapter on which you want to enable or disable the boot device to a non-boot-from-SAN host and run the command from there.
- If the host with the target adapter is also hosting other boot-from-SAN adapters, carry out a boot from SAN using one of the other boot-from-SAN adapters. Either command can now be run on the intended target adapter, because it is not being actively used for boot from SAN.

To enable the boot device, type:

```
oslpcfg enablebootdevice n=<adapter number>
```

To disable the boot device, type:

```
oslpcfg disablebootdevice n=<adapter number>
```

NOTE For the change to take effect, perform a system reboot.

4.3.5.4 Read All Alternative Boot Devices (readaltboot)

This command shows the WWN and LUN numbers in decimal format of all possible alternate boot devices. You can have up to seven alternate boot devices.

This command also allows you to identify available alternate boot devices that can be configured using the `setaltboot` command.

To read all alternate boot devices, type:

```
oslpcfg readaltboot n=<adapter number>
```

4.3.5.5 Setting One or More Alternate Boot Devices (setaltboot)

This command configures alternate boot devices. You can set up to seven alternate boot devices; that is, `index i` can be from 1 to 7.

NOTE The boot code must be enabled before you issue the `setaltboot` command.
`index` in this command does not refer to the same index as in the `enableboot` command.

To set up one or more alternate boot devices, type (all on one line):

```
oslpcfg setaltboot n=<adapter number> i=<index> w0=<wwpn word 0> w1=<wwpn word 1>
l=<Decimal ID of LUN>
```

The following example sets the alternate boot device on adapter number 1, LUN number 3:

```
oslpcfg setaltboot n=1 i=1 w0=12345678 w1=a842b6ed l=3
```

4.3.5.6 Enable or Disable PXE Boot on NIC Devices (networkboot)

NOTE This command is supported only on LPe16202/OCe15100 adapters in NIC mode.

This command allows you to enable or disable PXE boot on NIC devices.

To enable or disable PXE boot on a NIC device, type:

```
oslpcfg networkboot n=<adapter_number> t=<type> s=<status>
```

The following example enables PXE boot on adapter 1:

```
oslpcfg networkboot n=1 t=pxe s=1
```

The following example disables PXE boot on adapter 1

```
oslpcfg networkboot n=1 t=pxe s=0
```

4.3.6 Configuration Commands

Configuration commands allow you to update the configuration region of an adapter by name or by number.

4.3.6.1 Updating Configuration Regions (config)

Two forms of configuration are available:

- Configure all adapters of a given adapter name at once
- Configure a single adapter by its number

Valid region numbers range from 0 to 32.

4.3.6.1.1 Update by Name

To update a specified configuration region on all adapters of the same selected name, type:

```
oslpcfg config a=<adapter name> r=<region number> c=<configuration filename>
```

NOTE The adapter name is the name that appears when you run the `listHBA` command. For more information on the `listHBA` command, see Section 4.3.2.7, Viewing All Adapters in the System (listhba).

The following example updates region 6 of all LP12000 adapters with `ctplus1.cfl`:

```
oslpcfg config a=lpe12000 r=6 c=ctplus1.cfl
```

The following example updates region 17 of all ABC24-FC56 adapters with `d:\dfplus1.cfl`:

```
oslpcfg config a=ABC24-FC56 r=17 c=d:\dfplus1.cfl
```

NOTE The size of the `.cfl` file for configuration region update can be up to 2028 bytes.

4.3.6.1.2 Update by Number

To update a specified configuration region for one adapter, type:

```
oslpcfg config n=<adapter number> r=<region number> c=<configuration filename>
```

The following example updates region 17 of adapter number 4 with `heplus1.cf1`:

```
oslpcfg config n=4 r=17 c=heplus1.cfl
```

The following example updates region 6 of adapter number 2 with `d:\dfplus1.cf1`:

```
oslpcfg config n=2 r=6 c=d:\dfplus1.cfl
```

4.3.6.2 Writing a Semi-Volatile MAC Address on a NIC Device (writesvmac)

NOTE This command is supported only on OEM-specific LPe16202/OCe15100 adapters in NIC mode.

To write a semivolatile MAC address to an adapter, type:

```
oslpcfg writesvmac n=<adapter_number> <m=mac_address>
```

The following example writes the MAC address `0090FA112233` to adapter 1:

```
oslpcfg writesvmac n=1 m=0090FA112233
```

4.3.6.3 Restoring the Default Semi-Volatile MAC Address on a NIC Device (restoredefsvmac)

NOTE This command is supported only on OEM-specific LPe16202/OCe15100 adapters in NIC mode.

To restore the default semivolatile MAC address to an adapter, type:

```
oslpcfg restoredefsvmac n=<adapter_number>
```

The following example restores the default MAC address to adapter 1:

```
oslpcfg restoredefsvmac n=1
```

4.3.7 Diagnostic Tests

The following commands let you run diagnostic tests on your adapter.

4.3.7.1 Running the External Loopback Test (extloopback)

NOTE `extloopback` is supported only on LPe12000-series adapters.

ATTENTION Performing an `extloopback` test on an FC adapter that is being used to boot from SAN is not recommended. The `extloopback` command might cause a loss of connectivity to the SAN and possible loss of data. To perform an `extloopback` command on an FC adapter, make sure the adapter is not currently being used to boot from SAN. There are two ways to do this:

- Move the adapter on which you want to perform the external loopback test onto a non-boot-from-SAN host and perform the test from there.
- If the host with the test target adapter is also hosting other boot-from-SAN adapters, carry out a boot from SAN using one of the other boot-from-SAN adapters. The `extloopback` command can now be run on the intended target adapter, because it is not being actively used for boot from SAN.

This command runs the external loopback test. You must put a loopback plug in each adapter port to be tested. You can test a specific adapter in the system. Specify the number of times you want the test to repeat, and direct the test response if an error is found.

The option on error choices are:

- `o=1` stops the test on the first error
- `o=2` ignores three errors and either stops the test on the fourth error, or continues testing the next adapter on the fourth error
- `o=3` ignores errors and continues the test

NOTE The `extloopback` command does not support testing all installed adapters using `n=all`.

To run the external loopback test on an FC adapter, type:

os`lpcfg extloopback n=<`*adapter number*`> r=<`*repeat count*`> o=<`*option on error*`>`

where:

- `r` = 1 to 4096

To run the external loopback test on an LPe16202/OCe15101 adapter in NIC+FCoE mode, type (all on one line):

os`lpcfg extloopback n=<`*adapter number*`> p=<`*pattern*`> c=<`*byte count*`> r=<`*repeat count*`>`

- `p` = 3 to 8-HEX byte pattern
- `c` = 1500 to 8192
- `r` = 1 to 4096

The following example runs the external loopback test 50 times on FC adapter number 1 and stops the test if an error occurs:

os`lpcfg extloopback n=1 r=50 o=1`

4.3.7.2 Running the Internal Loopback Test (intloopback)

ATTENTION Performing an internal loopback test on an FC adapter that is being used to boot from SAN is not recommended. The `intloopback` command might cause a loss of connectivity to the SAN and possible loss of data. To perform an `intloopback` command on an FC adapter, make sure the adapter is not currently being used to boot from SAN. There are two ways to do this:

- Move the adapter on which you want to perform the internal loopback test onto a non-boot-from-SAN host and perform the test from there.
- If the host with the test target adapter is also hosting other boot-from-SAN adapters, carry out a boot from SAN using one of the other boot-from-SAN adapters. The `intloopback` command can now be run on the intended target adapter, because it is not being actively used for boot from SAN.

This command runs the internal loopback test. You can run the test on a specific adapter in the system. Specify the number of times you want the test to repeat, and direct the test response if an error occurs.

The option on error choices are:

- `o=1` stops the test on the first error
- `o=2` ignores three errors and either stops the test on the fourth error, or continues testing the next adapter on the fourth error
- `o=3` ignores errors and continues the test

NOTE `intloopback` does not support testing all installed adapters using `n=all`.
Do not use the internal loopback test on Emulex blade adapters.

To run the internal loopback test on an FC adapter, type:

`oslpcfg intloopback n=<adapter number> r=<repeat count> o=<option on error>`

where:

- `r` = 1 to 4096

To run the internal loopback test on an LPe16202/OCe15100 adapter in NIC+FCoE mode, type (all on one line):

`oslpcfg intloopback n=<adapter number> p=<pattern> c=<byte count> r=<repeat count> t=<type>`

where:

- `p` = 3 to 8-HEX byte pattern
- `c` = 1500 to 8192
- `r` = 1 to 4096
- `t` = 2

The following example runs the internal loopback test 100 times on FC adapter number 1 and stops the test if an error occurs:

`oslpcfg intloopback n=1 r=100 o=1`

4.3.7.3 Running the PCI Loopback Test (pciloopback)

ATTENTION Performing a PCI loopback test on an FC adapter that is being used to boot from SAN is not recommended. The `pciloopback` command might cause a loss of connectivity to the SAN and possible loss of data. To perform a `pciloopback` command on an FC adapter, make sure the adapter is not currently being used to boot from SAN. There are two ways to do this:

- Move the adapter on which you want to perform the PCI loopback test onto a non-boot-from-SAN host and perform the test from there.
- If the host with the test target adapter is also hosting other boot-from-SAN adapters, carry out a boot from SAN using one of the other boot-from-SAN adapters. The `pciloopback` command can now be run on the intended target adapter, because it is not being actively used for boot from SAN.

This command runs the PCI loopback test. You can run the test on a specific adapter or on all adapters in the system. Specify the number of times you want the test to repeat, and direct the test response if an error occurs.

The option on error choices are:

- `o=1` stops the test on the first error
- `o=2` ignores three errors and stops the test on the fourth error
- `o=3` ignores errors and continues the test

NOTE The `pciloopback` command supports testing all installed adapters using `n=all` on LPe12000-series adapters only.

To run the PCI loopback test, type:

`oslpcfg pciloopback n=<all|adapter number> r=<repeat count> o=<option on error>`

The following example runs the PCI loopback test 100 times on all LPe12000-series adapters in the system and stops the testing if any errors occur.

```
oslpcfg pciloopback n=all r=100 o=1
```

4.3.8 Using Script Files

Use script files to efficiently perform tasks. Script files are common sequences of commands you use when performing tasks.

You can group commands together and run them using a script file. You can also enter comment lines, which begin with a semicolon. Each line follows the same command syntax as those documented in this manual. Using the offline utility, you can:

- Run commands entered in a script file. Use the `@` command to run the script file.
- Run commands multiple times. Add the `repeat` command as the last line of the script file.
- Create a log of test results. Add the `logfile` command as the first line of the script file.

To run a script file type:

`oslpcfg @<scriptname.txt>`

The following example runs the script `script1.txt`, which resides in the current directory and runs all the commands in that script file.

```
oslpcfg @script1.txt
```

NOTE To interrupt and stop any script, press **s** on the keyboard.

The following is a sample script file. Each command follows the syntax covered previously. The comment lines begin with a semicolon (`;`):

```
version
screendisplay o=0
;download a=lpe12000 i=c:\temp\ud201a12.all
;reset n=1 s=0
;reset n=2 s=0
reset n=all s=0
listboot n=1
enableboot n=1 i=2
; pciloopback n=1 r=10 o=1
; pciloopback n=2 r=10 o=2
; pciloopback n=all r=50 o=3
intloopback n=1 r=10 o=1
intloopback n=2 r=10 o=1
extloopback n=1 r=40 o=3
extloopback n=2 r=40 o=3
;repeat r=10
```

4.3.8.1 Repeating a Series of Commands (repeat)

Enter this command at the end of a script file to repeat a series of commands from the beginning of the script file a specific number of times.

NOTE To interrupt and stop the `repeat` command, press **s** on the keyboard.

To repeat the series of commands in the script file, add the following as the last line of the file:

```
repeat r=<repeat count>
```

The following example repeats the series of commands in the script file ten times.

```
repeat r=10
```

4.3.8.2 Enabling or Disabling Test Messages on the Screen (screendisplay)

This command enables or disables test message displays on the screen.

- `o=0` – Prevents messages from appearing
- `o=1` – Enables messages

NOTE This command is supported only in script files.

To enable or disable test message displays on the screen, add the following line to the script:

```
screendisplay o=<display option>
```

In the following script file example, messages from the `version` and `listhba` commands appear in the log file and on the screen. After the `screendisplay` command is set to `0`, the result messages for all successive commands (`download`, `listboot`, and `enableboot`) appear only in the log file, not on the screen.

```
version
listhba
```

```
screendisplay o=0
download a=lpe12000 i=c:\temp\ud201a12.all
listboot n=2
enableboot n=2 i=1
```

4.3.8.3 Updating Nonvolatile WWNN (scriptwwnn)

ATTENTION Running a `scriptwwnn` command on an FC adapter that is being used to boot from SAN is not recommended. The `scriptwwnn` command might cause a loss of connectivity to the SAN and possible loss of data. To update the nonvolatile WWNN on an FC adapter, make sure the adapter is not currently being used to boot from SAN. There are two ways to do this:

- Move the adapter on which you want to use the `scriptwwnn` command to a non-boot-from-SAN host and run the script from there.
- If the host with the target adapter is also hosting other boot-from SAN-adapters, carry out a boot from SAN using one of the other boot-from-SAN adapters. You can now run the script to update the nonvolatile WWNN on the intended target adapter, because it is not being actively used for boot from SAN.

This command reads the WWNN words 0 and 1 from the command line to update the nonvolatile WWNN. You can also include this command in a script file. When the adapter is discovered, the new WWNN value is used. The adapter retains the original WWNN in another region of the firmware.

ATTENTION Use the `scriptwwnn` command with caution. If you use the same WWNN on more than one adapter in a fabric, unpredictable results might occur.

To change WWNN words 0 and 1 from the command line, type:

os`lpcfg scriptwwnn n=<`*`adapter number`*`> w0=<`*`wwnn word 0`*`> w1=<`*`wwnn word 1`*`>`

The following example updates nonvolatile WWNN word 0 and word 1 for adapter number 1:

os`lpcfg scriptwwnn n=1 w0=10000345 w1=B620A1B2`

NOTE Word 0 of WWNN and WWPN names must follow one of the following formats:

1 0 0 0 0 *x x x*
2 *x x x x x x x*
3 *x x x x x x x*
5 *x x x x x x x*

If the `scriptvwwnn` command has been used previously, the adapter continues to use that WWNN until you change the WWNN with the `restorenvwwn` command.

4.3.8.4 Updating Nonvolatile WWPN (scriptwwpn)

ATTENTION Running a `scriptwwpn` command on an FC adapter that is being used to boot from SAN is not recommended. The `scriptwwpn` command might cause a loss of connectivity to the SAN and possible loss of data. To update the nonvolatile WWPN on an FC adapter, make sure the adapter is not currently being used to boot from SAN. There are two ways to do this:

- Move the adapter on which you want to use the `scriptwwpn` command to a non-boot-from-SAN host and run the script from there.
- If the host with the target adapter is also hosting other boot-from-SAN adapters, carry out a boot from SAN using one of the other boot-from-SAN adapters. You can now run the script to update the nonvolatile WWPN on the intended target adapter, because it is not being actively used for boot from SAN.

This command reads WWPN words 0 and 1 from the command line to update the nonvolatile WWPN. You can also include this command in a script file. When it is discovered, the adapter uses the new WWNN value. It does not use the original IEEE address assigned by manufacturing (located in Configuration Region 16 or 32).

ATTENTION Use the `scriptwwpn` command with caution. If you use the same WWPN on more than one adapter in a fabric, unpredictable results might occur.

NOTE If the `scriptvwwnn` command has been used previously, the adapter continues to use that WWPN until you change the WWPN with the `restorenvwwn` command.

To change WWPN words 0 and 1 from the command line, type:

`oslpcfg scriptwwpn n=<adapter number> w0=<wwpn word 0> w1=<wwpn word 1>`

The following example updates the nonvolatile WWPN word 0 word 1 for adapter number 1:

`oslpcfg scriptwwpn n=1 w0=20A2D6B8 w1=C920A1B2`

4.3.8.5 Updating Volatile WWNN (scriptvwwnn)

ATTENTION Running a `scriptvwwnn` command on an FC adapter that is being used to boot from SAN is not recommended. The `scriptvwwnn` command might cause a loss of connectivity to the SAN and possible loss of data. To update the volatile WWNN on an FC adapter, make sure the adapter is not currently being used to boot from SAN. There are two ways to do this:

- Move the adapter on which you want to use the `scriptvwwnn` command to a non-boot-from-SAN host and run the script from there.
- If the host with the target adapter is also hosting other boot-from-SAN adapters, carry out a boot from SAN using one of the other boot-from-SAN adapters. You can now run the script to update the volatile WWNN on the intended target adapter, because it is not being actively used for boot from SAN.

This command reads the WWNN words 0 and 1 from the command line to update the volatile WWNN. The next time the adapter is discovered, it uses this new WWNN. It does not use the original IEEE address assigned by manufacturing (located in Configuration Region 16 or 32), nor does it use the value entered by the `writewwn` or `scriptwwnn` commands.

If the adapter does not have firmware that supports the volatile WWN, the following error message appears:

```
Write Volatile Parms Error. Reported Error 48
```

If this occurs, install firmware that supports the volatile WWN.

NOTE Word 0 of WWNN and WWPN names must follow one of the following formats:

```
1 0 0 0 0 x x x
2 x x x x x x x
3 x x x x x x x
5 x x x x x x x
```

After you issue this command, the volatile WWNN is used by the adapter until the `restorenvwwn` command is issued or the system is restarted.

ATTENTION Use the `scriptvwwnn` command with caution. If you use the same volatile WWNN on more than one adapter in a fabric, unpredictable results might occur.

To change volatile WWNN words 0 and 1 from the command line, type:

os`lpcfg scriptvwwnn n=<`*adapter number*`> w0=<`*wwnn word 0*`> w1=<`*wwnn word 1*`>`

The following example updates the volatile WWNN word 0 word 1 for adapter number 1:

os`lpcfg scriptvwwnn n=1 w0=20A2D6B8 w1=C920A1B2`

4.3.8.6 Updating Volatile WWPN (scriptvwwpn)

ATTENTION Running a `scriptvwwpn` command on an FC adapter that is being used to boot from SAN is not recommended. The `scriptvwwpn` command might cause a loss of connectivity to the SAN and possible loss of data. To update the volatile WWPN on an FC adapter, make sure the adapter is not currently being used to boot from SAN. There are two ways to do this:

- Move the adapter on which you want to use the `scriptvwwpn` command to a non-boot-from-SAN host and run the script from there.
- If the host with the target adapter is also hosting other boot-from-SAN adapters, carry out a boot from SAN using one of the other boot-from-SAN adapters. You can now run the script to update the volatile WWPN on the intended target adapter, because it is not being actively used for boot from SAN.

This command reads the WWPN words 0 and 1 from the command line to update the volatile WWPN. The next time the adapter is discovered, it uses this new WWPN. It does not use the original IEEE address assigned by manufacturing (located in Configuration Region 16 or 32), nor does it use the value entered by the `writewwn` or `scriptwwpn` commands.

If the adapter does not have firmware that supports the volatile WWN, the following error message appears:

```
Write Volatile Parms Error. Reported Error 48
```

If this error message appears, install firmware that supports the volatile WWN.

NOTE Word 0 of WWNN and WWPN names must follow one of the following formats:

```
1 0 0 0 0 x x x
2 x x x x x x x
3 x x x x x x x
5 x x x x x x x
```

After you issue this command, the volatile WWNN is used by the adapter until the `restorenvwwn` command is issued or the system is restarted.

ATTENTION Use the `scriptvwwnn` command with caution. If you use the same VWWNN on more than one adapter in a fabric, unpredictable results might occur.

To change VWWPN words 0 and 1 from the command line, type:

`oslpcfg scriptvwwpn n=<adapter number> w0=<wwpn word 0> w1=<wwpn word 1>`

The following example updates the volatile WWPN word 0 word 1 for adapter number 1:

```
oslpcfg scriptvwwpn n=1 w0=20A2D6B8 w1=C920A1B2
```

4.3.8.7 Creating a Log (logfile)

This command creates a log file with a specified directory and file name. The default log file is `lpcfglog.txt` and is created in the system's current directory.

NOTE This command is supported only in script files. Make it the first command in the script.

To create a log file, type the following command in the script file:

`logfile l=<filename>`

The following example creates a file called `lplog.txt` in the `d:\` directory:

```
logfile l=d:\log\lplog.txt
```

Results of all commands are recorded in a log file. Unless otherwise specified by the `logfile` command, the default log file is `lpcfglog.txt` in the current directory.

4.3.9 VLAN Commands

Using the following commands, you can manage the VLAN ID and VLAN priorities on specified adapters.

4.3.9.1 Reading VLAN Priorities (readvlanprops)

This command shows the NIC function on an adapter's VLAN ID and VLAN Priority for the adapter number specified.

NOTE `readvlanprops` is supported only on LPe16202/OCe15100 adapters in NIC+FCoE mode.

To read the VLAN priorities for the adapter specified, type:

`oslpcfg.exe readmac n=<adapter number>`

The following example lists information for adapter number 1:

```
oslpcfg.exe readvlanprops n=1
Command: readvlanprops n=1
adapter 1:
VLANId          : 1024
VLANPriority: 1
Command completed, NO Error
```

4.3.9.2 Restoring the Default Values for the VLAN ID and VLAN Priorities (restorevlanprops)

This command restores the adapter's default VLAN ID and VLAN Priority values, which are both set to 0, on the adapter number specified.

NOTE `restoredefvlanprops` is supported only on LPe16202/OCe15100 adapters in NIC+FCoE mode..

To restore the VLAN ID and VLAN priorities for the adapter specified, type:

```
>oslpcfg.exe restorevlanprops n=<adapter number>
```

4.3.9.3 Changing the VLAN ID and VLAN Priority on a Specified Adapter's NIC Function (writevlanprops)

This command lets you change the VLAN ID and VLAN Priority on the specified adapter's NIC function. You can specify the VLAN ID and the VLAN Priority with the following values:

- `vlan_id` – 0 to 4095. A value of 0 disables the VLAN ID.
- `vlan_priority` – 0 to 7.

NOTE `writevlanprops` is supported only on LPe16202/OCe15100 adapters in NIC+FCoE mode..

To change the VLAN ID and the VLAN Priority on a specified adapter's NIC function, type:

```
oslpcfg.exe writevlanprops n=<adapter number> i=<vlan_id> p=<vlan_priority>
```

The following example changes the VLAN ID and VLAN Priority for adapter number 1:

```
oslpcfg.exe writevlanprops n=1 i=1024 p=1
Command: writevlanprops n=1 i=1024 p=1
adapter 1:
Command completed, NO Error
```

Chapter 5: Status Messages and Error Codes

5.1 Elxflash Utility Status Messages

The Elxflash utility takes a less stringent approach to reporting status messages when used with auto-discovery. An error message is returned only when the Elxflash utility reports a failure during a download or if the firmware and boot directories cannot be located.

Table 2 lists the supported Elxflash utility status messages.

Table 2 Elxflash Status Messages

Message Number	Message Title	Message Details
0	GOOD_ALL_UPGRADES_OK	Returned if any of the discovered adapters had a successful download performed.
1	GOOD_NO_UPDATES_NEEDED	Returned if any of the discovered adapters did not need an update, a downgrade, or a rewrite operation. This message is never returned during a force firmware or force boot code download operation.
2	ERROR_ALL_UPGRADES_FAILED	Returned if all attempted downloads failed.
3	ERROR_SOME_UPGRADES_FAILED	Returned if some of the attempted downloads failed.
4	ERROR_EMULEX_APPS_COMMAND	Returned if Elxflash cannot run an external executable, such as *os*lpcfg (elxflashOffline)
5	ERROR_NO_SUPPORTED_HBA_FOUND	Returned if no supported adapters are found.
6	ERROR_DIRECTORY_NOT_FOUND	Returned if the firmware or boot directories are missing. Depending on the operation, auto-discovery expects the firmware or boot directories to exist as subdirectories in the root of the package directory.
7	GOOD_NO_UPGRADES_AVAILABLE	Returned if none of the discovered adapters had a matching image in the firmware or boot directories. In this case only, a matching image is an image that the adapter accepts regardless of the download operation being performed.
8	ERROR_NOT_ADMIN_USER	Returned when you attempt to run the executable file and do not have administrator (Windows) or root (Linux) privileges.
9	ERROR_UNSUPPORTED_OS	Used by the Exlflash Stand-alone execution scripts. Returned when the scripts are run on an unsupported operating system.
10	ERROR_UNSUPPORTED_ARCH	Used by the Exlflash Stand-alone execution scripts. Returned when the scripts are run on an unsupported architecture.
19	ERROR_IMAGE_VERSION	An error occurred when decoding the image version. The image file has an invalid extension, or Elxflash was unable to decode the image file's version. **NOTE** This is returned only when using the /iv=<*image_version*> switch.
20	ERROR_MISSING_DEP	Returned when a missing dependency is detected.

5.2 LpCfg Error Codes

Table 3 lists the supported LpCfg error codes.

Table 3 LpCfg Error Codes

Error Code	Description
0	No error
1	Invalid adapter number
2	Mailbox command error
3	No valid boot (BIOS) code found
4	Open file error
5	Invalid configuration region
6	Invalid adapter name
7	Download error
8	Invalid boot code index
9	Link not up for external loopback test
10	Link not up for internal loopback test
11	Invalid jumper selection (in jumper command)
12	Invalid alternate configuration region (in jumper command)
13	PCI loopback test fails
14	Adapter reset error
15	Read configuration region error
16	No VPD information available
17	No command in command line
18	Open log file error
19	Read wakeup parameters error
20	Update wakeup parameters error
21	Incorrect test parameters
22	Stopped by user
23	Internal loopback test fails
24	External loopback test fails
25	Error exists after four retries
26	Invalid command
27	Incorrect syntax
28	Command supported only in script files
29	Read_rev error
30	Dump configuration region error
31	Read file error
32	Short file error
33	Read NVPARMS error
34	Write NVPARMS error
35	Command does not support all adapters
36	Invalid LUN

Table 3 LpCfg Error Codes (Continued)

Error Code	Description
37	No boot (BIOS) code enabled
38	Update configuration region error
39	No boot (BIOS) found
40	Dump memory error
41	Update EROM error
42	Delete load entry error
43	Write WWN error
44	Not supported in script files
45	No Emulex adapter found
46	Invalid alternate boot device index
47	Cannot restart adapter
48	Write volatile parameters error
49	POST test error
50	Incorrect symbols
51	Invalid length
52	Invalid topology
53	No event log
54	Read event log
55	Invalid input value
56	No libdfc library
57	Non-numeric input
58	No valid WWN
59	Region cleanup
60	Region initialization
62	Unable to allocate memory
63	DFC_InitDiagEnv error
64	DFC_ReadPciCfg error
65	No driver installed
66	No valid driver
67	Not valid adapter type
68	Not valid image
69	Long file error
70	Incompatible image
71	Not supported
72	MILI service not started
73	Script not supported
74	MILI not started
75	No NIC adapter
76	Personality information
77	Personality change

Table 3 LpCfg Error Codes (Continued)

Error Code	Description
78	Administrator rights error
79	SLI-4 management error
80	Reboot required
83	Application is missing one or more dependencies.
87	The system must be rebooted to activate the new firmware image and the image must be downloaded again.
88	Duplicate MAC address
89	Download failed due to missing digital signature. Please contact customer support for additional help. Download failed on adapter <adapter number> Stat 89.
90	Download failed due to invalid firmware digital signature. Please contact customer support for additional help. Download failed on adapter <adapter number> Stat 90.
141	General error from MILI
200	General error

Chapter 6: Troubleshooting

When using Elxflash, some commands take precedence over others. For instance, the force firmware (`/ff`), force boot (`/fb`), and force firmware and boot code (`/f`) switches take precedence over the `/downgrade`, `/update`, and `/rewrite` switches. If `/ff`, `/fb`, or `/f` are used, the `/downgrade`, `/update`, and `/rewrite` switches are ignored.

The following examples illustrate this behavior:

- `./elxflash /ff /update`
  - `/update` is ignored.
  - Performs a force firmware operation on all installed and supported adapters.
- `./elxflash /fb /update`
  - `/update` is ignored.
  - Performs a force boot code operation on all installed and supported adapters.
- `./elxflash /ff /downgrade`
  - `/downgrade` is ignored
  - Performs a force firmware operation on all installed and supported adapters.
- `/elxflash /fb /downgrade`
  - `/downgrade` is ignored.
  - Performs a force boot code operation on all installed and supported adapters.
- `./elxflash /ff /rewrite`
  - `/rewrite` is ignored
  - Performs a force firmware operation on all installed and supported adapters.
- `./elxflash /fb /rewrite`
  - `/rewrite` is ignored.
  - Performs a force boot code operation on all installed and supported adapters.

6.1 Unsupported Driver

The Elxflash utility reports an error similar to the following if an unsupported driver is installed on the system:

```
elxflash: no supported Emulex HBA's found - Return Code=1
```

If an error similar to this example occurs, verify that the correct version of the driver is installed. For supported adapters, and supported versions of operating systems and platforms, go to http://www.broadcom.com.

6.2 Updating Adapters without Boot Code

The Elxflash utility reports an adapter's boot code version as `NONE` when the adapter does not have boot code installed. The utility does not update, downgrade, or rewrite the boot area if boot code is not present, but firmware will still be updated or downgraded without boot code. You can force a boot code download using auto-discovery or the `fwmatrix.txt` file.


BROADCOM®

Emulex® Drivers for Windows for LightPulse® Adapters

Quick Installation Guide

Version 11.4
September 6, 2017

DRVWin-LPE-IG114-100

Corporate Headquarters
San Jose, CA

Website
www.broadcom.com

Emulex Driver for Windows for LightPulse Adapters Quick Installation Guide

1 Introduction

This document explains how to install the Emulex® drivers for Windows on your system.

This product supports the following Emulex LightPulse® adapters:

- LPe12000-series adapters
- LPe16000-series adapters, including LPe16202/OCe15100 adapters
- LPe31000-series adapters
- LPe32000-series adapters

2 Driver Installation Options

Install the Windows drivers for LightPulse adapters in one of two ways:

- OneInstall Installer contains the Emulex Storport Miniport driver, Emulex PLUS (ElxPlus) driver, and the OneCommand® Manager application for LightPulse® Adapters in a single download package.

  NOTE OneInstaller Installer does not include NIC and FCoE drivers for LPe16202/OCe15100 adapters in NIC+FCoE mode. To install these drivers, you must use the individual driver kits.

- Driver kits and AutoPilot Installer provide installation options ranging from simple installations with a few mouse clicks to unattended installations that use predefined script files and text-only installations.

  NOTE For LPe16202 adapters, OneInstall Installer, or network interface card (NIC) driver installations, Microsoft patch KB2846340 must be installed on your system to successfully install the NIC driver. If the patch is not installed on your system, the installation stops and prompts you to install it. This patch, from Microsoft's Knowledge Base, is required for Windows Server 2012, Windows 8, and Windows 10 client operating systems.

  NOTE If you are installing the NIC driver kit as an update to the Windows Server 2012 driver, Windows Server 2012 driver R2, or Windows Server 2016 driver, some parameter defaults are different from the inbox driver. After you install the Emulex out-of-box driver, select **reset to default** on the **Advanced** tab of the **Device Manager Property** page. This action returns all adapter and driver settings to the default values listed in this guide.

  NOTE Low performance might occur if the Emulex NIC driver is installed on a system meeting the following conditions before installing Microsoft patch KB2846837:

  - A Windows 8, Windows 8.1, or Windows Server 2012 computer with multi-core processors is in use

- Three or more Ethernet ports are installed on the computer
- RSS is enabled and sets the RSS profile to use the `Closest` parameter for the Ethernet adapters

If any or all of these conditions exist, install patch KB2846837 before installing the Emulex NIC driver.

NOTE Windows 8 x64 and Windows 8.1 x64 drivers are Emulex signed. You must accept the Emulex certificate to install these kits. Support is provided by Broadcom, but not by Microsoft.

NOTE Check the Broadcom® website for required updates to the Windows operating system or the Emulex drivers.

2.1 OneInstall Installer

OneInstall Installer can be run in Interactive mode or Silent mode.

NOTE OneInstall Installer does not allow you to perform pre-installation tasks or text-only installations. For these tasks, use the driver kits.

NOTE OneInstaller Installer does not include NIC and Fibre Channel over Ethernet (FCoE) drivers for LPe16202/OCe15100 adapters in NIC+FCoE mode. To install these drivers, you must use the individual driver kits.

OneInstall Installer is a self-extracting executable file that installs the following software on your system:

- Fibre Channel (FC) drivers
- Emulex PLUS (ElxPlus) driver (supports the OneCommand Manager application, persistent binding, logical unit number (LUN) mapping and masking, and virtual port creation)
- OneCommand Manager application

NOTE The Enterprise kit for the OneCommand Manager application does not operate locally on Windows Server Core. You must install the OneCommand Manager Core Kit (command line interface only) to the Windows Server Core. Refer to the *OneCommand Manager Application for LightPulse Adapters User Guide* for installation instructions.

2.1.1 Loading OneInstall Installer Using Interactive Mode

To install the drivers using Interactive mode, perform these steps:

1. Download OneInstall Installer from the Broadcom website.
2. Navigate to OneInstall Installer in Windows Explorer.
3. Double-click **OneInstall Installer**.

   The **Welcome** screen is displayed.
4. Click **Next**.

   The **Installation Options** screen is displayed
5. Select the drivers and applications that you want to install and click **Next**.

   A progress screen is displayed while OneInstall installer loads the selected drivers and applications. After the drivers and application software are loaded, an **Installation completed** screen is displayed.
6. Click **Finish**.

2.1.2 Loading OneInstall Installer Using Silent Mode

Silent mode installation must be run from a batch file or from the command line.

If you run OneInstall Installer from a batch file or from a command line prompt, the default Windows behavior starts OneInstall Installer, then immediately continues with the next command. It does not wait until OneInstall Installer has finished.

As a result, the value of %ERRORLEVEL% will always be 0 because Windows successfully started OneInstall. It does *not* reflect an accurate OneInstall exit code.

To remedy this, run OneInstall Installer setup as follows:

```
START /wait OneInstall-Setup-<version>.exe
echo %ERRORLEVEL%
```

The `START /wait` ensures that the command does not return until setup has exited. The value of %ERRORLEVEL% now accurately reflects the OneInstall Installer exit code.

2.1.3 Command Format

The format of the command is:

`OneInstall-Setup-<version>.exe <install-mode> <options>`

Where:

<*version*> is the version number of OneInstall Installer

<*install-mode*> is one of the following:

- `/q0` – (Interactive, non-silent install) This is the default.
- `/q1` – (non-Interactive install) This option displays status pages.
- `/q2` – (Silent install) This option is completely silent, no pages are displayed.
- `/q` – This is the same as `/q1`.

<*options*> specifies the kit, or kits, to install:

- `ALL=1` – Install theFC driver and the OneCommand Manager application (default).

  NOTE On a Windows Server Core system, this installs all drivers and the OneCommand Manager Core Kit.

- `ALLCORE=1` – Install the FC driver and the OneCommand Manager Core Kit.
- `FC=1` – Install the FC driver only.
- `OCM=1` – Install the OneCommand Manager Enterprise Kit only.

  NOTE On a Windows Server Core system, this installs the OneCommand Manager Core Kit.

- `OMCORE=1` – Install the OneCommand Manager Core Kit only.

To install the drivers using Silent mode, perform these steps:

1. Download OneInstall Installer from the Broadcom website.
2. Open a DOS window.
3. Change the directory to the folder containing your OneInstall Installer.

The following are examples of Silent mode commands:

```
Start /wait OneInstall-Setup-<version>.exe  /q2  ALL=1
Start /wait OneInstall-Setup-<version>.exe  /q2  FC=1
Start /wait OneInstall-Setup-<version>.exe  /q2  OCM=1
```

```
Start /wait OneInstall-Setup-<version>.exe  /q2  ALLCORE=1
Start /wait OneInstall-Setup-<version>.exe  /q2  OCMCORE=1
```

2.2 Driver Kit Installer

Each driver kit contains and loads all the Windows drivers for a specific protocol.

- FC driver package includes ElxPlus (`elxdrvr-fc-<version>.exe`)
- FCoE driver package includes ElxPlus (`elxdrvr-fcoe-<version>.exe`)
- NIC driver package (`elxdrvr-nic-<version>.exe`)

NOTE FCoE and NIC driver packages are used on LPe16202/OCe15100 adapters only.

NOTE Updating the NIC protocol driver can temporarily disrupt operation of any NIC teams configured on the system.

2.2.1 Loading the Driver Kit

The driver kit copies the selected Emulex drivers and applications onto your computer.

NOTE This procedure does not install drivers, and no driver changes are made until you run AutoPilot Installer.

To load the driver kit, perform these steps:

1. Download the driver kit from the Broadcom website to your system.
2. Double-click the driver kit to run it.

   The **Emulex Kit Welcome** page is displayed.
3. Click **Next**.

   The **Installation Options** window is displayed.
4. Select one or both of the following options:
   - **Perform Installation of Software** – Copies the driver kit for your operating system to your computer.
   - **Unpack All Drivers** – Extracts all drivers to the current user's `Documents` folder. Select this option to perform boot from SAN installations.

   The **Operation in progress** window shows the kit file-loading progress. After the kit files are loaded, the **Installation completed** window is displayed.
5. To continue with the installation, ensure that **Start AutoPilot Installer** is selected.

2.3 AutoPilot Installer

AutoPilot Installer runs after the driver kit is loaded and the OneCommand Manager application is installed. AutoPilot Installer can be installed at these times:

- Immediately after the driver kit has been loaded
- At a later time using an interactive installation
- Through an unattended installation

AutoPilot Installer provides the following functions:

- Command line functionality – Initiates an installation from a command prompt or script. Configuration settings can be specified in the command line.

- Compatibility verification – Verifies that the driver to be installed is compatible with the operating system and platform.
- Driver installation and update – Installs and updates drivers.
- Multiple adapter installation capability – Installs drivers on multiple adapters, alleviating the requirement to manually install the same driver on all adapters in the system.

NOTE Refer to the *Emulex Software Kit Migration User Guide* for information about installing drivers on a system containing a mix of OneConnect® and LightPulse adapters.

- Driver diagnostics – Determines whether the driver is operating properly.
- Silent installation mode – Suppresses all screen output (necessary for unattended installation).

NOTE AutoPilot Installer does not allow you to install the driver if the minimum Windows service pack or Microsoft Storport driver update is not installed.

You can install the driver by using any of the following methods:

NOTE These installation methods are not mutually exclusive.

- **Hardware-first installation**. At least one Emulex adapter must be installed before you can install the Emulex drivers and utilities.
- **Software-first installation**. You can install drivers and utilities using AutoPilot Installer prior to the installation of any adapters. You do not need to specify the adapter models to be installed later. The appropriate drivers and utilities automatically load when you install the adapters.
- **Utility-only installation**. If the drivers in the driver kit share the same version with those already installed on the system, you can reinstall or update the previously installed utility without reinstalling the drivers.
- **Text-only installation**. Text-based installation mode is used automatically when AutoPilot Installer is run on a Windows Server Core system.
- **Network installation.** You can place the driver kit installers on a shared network drive and install them across your local area network (LAN). Network-based installation is often used with unattended installation and scripting, which allows you to configure and install the same driver and utility versions on all of the hosts in a storage area network (SAN).
- **Unattended installation**. You can run the driver kit installers and AutoPilot Installer with no user interaction from a command line or script. Unattended installation works for both hardware-first and software-first installations and all driver kits. An unattended installation operates in Silent mode (also referred to as Quiet mode) and creates an extensive report file with installation status.

NOTE Complete driver and utilities documentation can be downloaded from the Broadcom website.

2.4 Starting Installers from a Command Prompt or Script

If a driver kit or AutoPilot Installer is run from a command prompt or command script (batch file), the Windows command processor does not wait for the installer to run to completion. As a result, you cannot check the exit code of the installer before the next command is executed. For command line invocation, always use the `START` command with the `/wait` option, which causes the command processor to wait for the installer to finish before the command processor continues.

For additional information on command line installation and configuration parameters, refer to Section C, AutoPilot Installer Command Line and Configuration File Parameters, in the *Emulex Drivers for Windows for LightPulse Adapters User Guide*.

2.5 Running a Software Installation Interactively

Two options are available when performing an installation interactively. These options assume you have already downloaded the driver kit from the Broadcom website.

- Option 1 allows you to automatically run AutoPilot Installer, which completes the driver kit loading and installation with a few mouse clicks.
- Option 2 allows you to run AutoPilot Installer separately. This option is recommended when you are:
  - Changing installation settings for a limited number of systems
  - Familiarizing yourself with AutoPilot Installer configuration options

2.5.1 Option 1: Automatically Run AutoPilot Installer

Use this option unless you have specific configuration requirements.

1. Double-click the driver kit or run it from a command line. The command line parameter `APargs` allows you to specify arguments that are automatically passed to the AutoPilot Installer command.

   A **Welcome** window is displayed with driver kit version information and Emulex contact information (refer to Appendix C, AutoPilot Installer Command Line and Configuration File Parameters in the *Emulex Driver for Windows for LightPulse Adapters User Guide* for additional information on command line installations).
2. Click **Next** to proceed to the **Installation Options** window.

   For each installation option, the default installation location for that option is displayed. Browse to a different location, if desired.
3. Click **Install** to continue the installation.

   The **Progress** dialog is displayed.

   After all tasks complete, a **Finish** window is displayed. The **Start AutoPilot Installer** check box is automatically selected.
4. Click **Finish**.

   AutoPilot Installer runs automatically and completes one of the following installations:
   - Section 2.6, Hardware-First Installation or Driver Update
   - Section 2.7, Software-First Installation

2.5.2 Option 2: Run AutoPilot Installer Separately

To access these options, run AutoPilot Installer after the driver kit loading has been completed, which allows you to change the configuration options supplied to AutoPilot Installer.

1. Perform steps 1 through 3 in Option 1: Automatically Run AutoPilot Installer.
2. Clear the **Run AutoPilot Installer** check box on the **Finish** dialog.
3. Click **Finish**.

   The driver kit installer exits.

After the driver kit loading is complete, change the configuration in one of two ways:

- Change the configuration file.
- Supply parameters on the command line.

NOTE Refer to Appendix C, AutoPilot Installer Command Line and Configuration File Parameters, in the *Emulex Drivers for Windows for LightPulse Adapters User Guide* for additional information on either of these configuration methods.

After you have finished this step, you can run AutoPilot Installer at a later time.

4. Run AutoPilot Installer using the following command:

```
C:\Program Files\Emulex\AutoPilot Installer\APInstall.exe
```

NOTE The location of `APInstall.exe` might differ on your system, depending on your system's Program Files location. You can also specify a different location when you install the driver package.

2.6 Hardware-First Installation or Driver Update

The driver kit installer must be downloaded from the Broadcom website and installed before performing this installation.

NOTE Updating the NIC protocol driver can temporarily disrupt operation of any NIC teams configured on the system.

NOTE To update the Emulex protocol drivers, begin this procedure at step 2.

To perform a hardware-first installation, perform these steps:

1. Install a new Emulex adapter and power-on the system. **If the Windows Found New Hardware** wizard is displayed, click **Cancel** to exit; AutoPilot Installer performs this function.

   NOTE If there are multiple adapters in the system, the **Windows Found New Hardware** wizard is displayed multiple times. Click **Cancel** to exit the wizard each time it is displayed.

2. Run AutoPilot Installer using one of the two options listed in Section 2.5, Running a Software Installation Interactively.

   Consider the following:

   — If you are updating the driver, the existing port settings are used, unless otherwise specified in the configuration file. These settings are pre-selected but can be changed. Set or change the settings, then click **Next**.

   — If you are initially installing a vendor-specific version of the Emulex driver installation program, a **Driver Configuration** window may be displayed. This window includes one or more windows with questions that you must answer before continuing the installation process. In this case, answer each question and click **Next** on each window to continue.

3. Click **Next**. The installation is completed automatically.

   A dialog is displayed if Windows requires a reboot. After the installation is successful, a **Finish** window is displayed.

4. View or print a report, if desired.

   — **View Installation Report** – The installation report is a text file with current Emulex adapter inventory, configuration information, and task results.

   — **Print Installation Report** – The Windows **Print** dialog is displayed to select options for printing the installation report.

5. Click **Finish** to exit AutoPilot Installer.

6. If the system must be rebooted, you are prompted to do so as indicated in step 3; you must reboot before using the drivers or utilities.

2.7 Software-First Installation

The driver kit must be downloaded from the Broadcom website and loaded.

To perform a software-first installation:

1. Run AutoPilot Installer using one of the two options listed in Section 2.5, Running a Software Installation Interactively.

   Figure 1 is displayed.

Figure 1 AutoPilot Installer Warning (Software-First Installation)

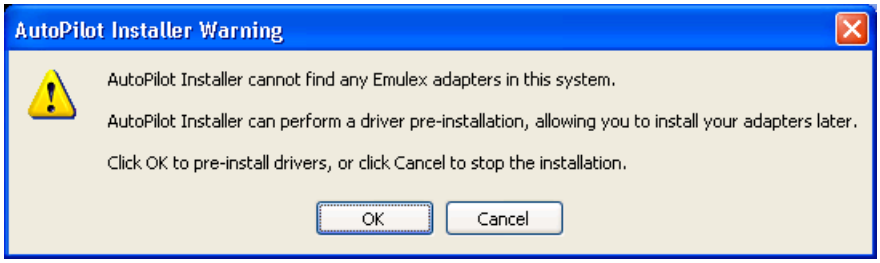

2. Click **OK**.

   A **Welcome** window is displayed.
3. Click **Next**. The installation automatically progresses.

   After the installation is successful, the **Finish** window is displayed.
4. View or print a report, if desired.
   - **View Installation Report** – The installation report is a text file with current Emulex adapter inventory, configuration information, and task results.
   - **Print Installation Report** – The Windows **Print** dialog is displayed to select options for printing the installation report.
5. Click **Finish** to exit AutoPilot Installer.

2.8 Text-Only Driver Installation

Text-based Installation mode is used automatically when the driver kit installer runs on a server with the Server Core installation option of Windows Server. During text-based installations, AutoPilot Installer uses a command prompt window. The driver kit installer notifies you when the driver is installed and also gives you an opportunity to stop the installation.

Whether AutoPilot Installer is launched from the command line or run as a program, Windows always starts AutoPilot Installer as a separate stand-alone task. This means that AutoPilot Installer has its own command prompt window and cannot access other windows.

3 Installation Failure

The two possible installation failures are described in this section.

3.1 AutoPilot Installer Failure

If AutoPilot Installer fails, the **Diagnostics** window shows that the adapter failed. If the adapter fails, perform these steps:

1. Select the adapter to view the reason why the adapter failed. The reason and suggested corrective action are displayed.
2. Perform the suggested corrective action, and run AutoPilot Installer again.

NOTE You can run `APInstall.exe` from a command prompt.

3.2 OneInstall Installer Failure

OneInstall Installer might fail for any of the following reasons:

- The operating system prerequisites have not been met.
- The individual kit installation failed. To check, run the installation interactively. If you encounter error messages when you run the installation interactively, those issues would also apply to an unattended installation.

If an individual package failed to install properly, run that package's installer directly. This method displays status and error messages that can be used to diagnose the issue. (OneInstall Installer does not provide this information because each package is installed silently.)


BROADCOM®

Emulex® Drivers for Windows for LightPulse® Adapters

User Guide

Version 11.4
September 6, 2017

DRVWin-LPE-UG114-100

Corporate Headquarters
San Jose, CA

Website
www.broadcom.com

Chapter 1: Introduction

This product supports the following Emulex® LightPulse® adapters:

- LPe12000-series adapters
- LPe16000-series adapters, including LPe16202/OCe15100 adapters
- LPe31000-series adapters
- LPe32000-series adapters

The Emulex drivers for Windows support the following protocols:

- Fibre Channel (FC)
- FC over Ethernet (FCoE) for LPe16202/OCe15100 adapters in NIC+FCoE mode
- Ethernet Network Interface Card (NIC) for LPe16202/OCe15100 adapters in NIC+FCoE mode

1.1 Driver Information

This document explains how to install the Emulex drivers for Windows on your system and configure the drivers' capabilities based on the supported networking protocols:

- FC and FCoE
  - Configuring the FC and FCoE driver parameters
  - Improving server performance with FC and FCoE drivers
- Ethernet
  - Configuring NIC driver options
  - Tuning network driver performance

1.2 Abbreviations

1GbE	1 Gigabit Ethernet
10GbE	10 Gigabit Ethernet
AL_PA	arbitrated loop physical address
ARM	Advanced RISC Machines
BIOS	basic input-output system
CPU	central processing unit
CRC	cyclic redundancy check
DCBX	Data Center Bridging Capabilities Exchange
DPC	deferred procedure call
DHCP	Dynamic Host Control Protocol
DID	device ID
DIMM	dual inline memory module
DMA	direct memory access
DNS	Domain Name Server
DOS	disk operating system

ETS	Enhanced Transmission Selection
FC	Fibre Channel
FC-AL	Fibre Channel Arbitrated Loop
FCoE	Fibre Channel over Ethernet
FCP	Fibre Channel Protocol
FDMI	Fabric-Device Management Interface
FLOGI	fabric login
FSB	front-side bus
GB	gigabyte
Gb/s	gigabits per second
GUI	graphical user interface
HBA	host bus adapter
IEEE	Institution of Electrical and Electronics Engineers
iSCSI	Internet Small Computer Systems Interface
I/O	input/output
IOMMU	input/output memory management unit
IP	Internet Protocol
KB	Knowledge Base
LACP	Link Aggregation Control Protocol
LAN	local area network
LSO	large send offload
LUN	logical unit number
MAC	Media Access Control
MDS	Multilayer Director Switch
MSI	message signaled interrupts
MTU	maximum transmission unit
N/A	not applicable
NDIS	Network Driver Interface Specification
NIC	network interface card
NPIV	N_Port ID virtualization
NTFS	New Technology File System
NUMA	non-uniform memory access
NVGRE	network virtualization using generic routing encapsulation
OS	operating system
PCI	Peripheral Component Interconnect
PCIe	PCI Express
PF	PCI function
PFC	process flow control or priority flow control
PLOGI	port login
POST	power-on self-test

PT-PT	point-to-point
PXE	Preboot Execution Environment
QFE	Quick Fix Engineering
RAID	redundant array of independent disks
RCMD	Remote Command Service
ROM	read-only memory
RSC	receive segment coalescing
RSCN	registered state change notification
RSS	receive-side scaling
RX	receive
SACK	selective acknowledgement
SAN	storage area network
SCSI	Small Computer System Interface
SFP	small form-factor pluggable
SLI	Service Level Interface
SR-IOV	Single Root I/O Virtualization
TCP	Transmission Control Protocol
TCP/IP	TCP over Internet Protocol
TX	transmit
UDP	User Datagram Protocol
ULP	unit of least precision
UEFI	Unified Extensible Firmware Interface
VF	virtual function
VHD	virtual hard disk
VLAN	virtual local area network
VLAN ID	VLAN identifier
VM	virtual machine
VMQ	virtual machine queue
WWN	World Wide Name
WWNN	World Wide Node Name
WWPN	World Wide Port Name
XRI	extensible resource indicator

Chapter 2: Installation

Install the Windows drivers for LightPulse adapters in one of two ways:

- OneInstall Installer contains the Emulex Storport Miniport driver, Emulex PLUS (ElxPlus) driver, and the OneCommand® Manager application for LightPulse® Adapters in a single download package.

  NOTE OneInstaller Installer does not include NIC and FCoE drivers for LPe16202/OCe15100 adapters in NIC+FCoE mode. To install these drivers, you must use the individual driver kits.

- Driver kits and AutoPilot Installer provide installation options ranging from simple installations with a few mouse clicks to unattended installations that use predefined script files and text-only installations.

  NOTE For LPe16202 adapters, OneInstall Installer, or NIC driver installations, Microsoft patch KB2846340 must be installed on your system to successfully install the NIC driver. If the patch is not installed on your system, the installation stops and prompts you to install it. This patch, from Microsoft's Knowledge Base, is required for Windows Server 2012, Windows 8, and Windows 10 client operating systems.

  NOTE If you are installing the NIC driver kit as an update to the Windows Server 2012 driver, Windows Server 2012 driver R2, or Windows Server 2016 driver, some parameter defaults are different from the inbox driver. After you install the Emulex out-of-box driver, select **reset to default** on the **Advanced** tab of the **Device Manager Property** page. This action returns all adapter and driver settings to the default values listed in this manual.

  NOTE Low performance might occur if the Emulex NIC driver is installed on a system meeting the following conditions before installing Microsoft patch KB2846837:
  - A Windows 8, Windows 8.1, or Windows Server 2012 computer with multi-core processors is in use
  - Three or more Ethernet ports are installed on the computer
  - RSS is enabled and sets the RSS profile to use the `Closest` parameter for the Ethernet adapters

  If any or all of these conditions exist, install patch KB2846837 before installing the Emulex NIC driver.

  NOTE Windows 8 x64 and Windows 8.1 x64 drivers are Emulex signed. You must accept the Emulex certificate to install these kits. Support is provided by Broadcom®, but not by Microsoft.

  NOTE Check the Broadcom website for required updates to the Windows operating system or the Emulex drivers.

2.1 OneInstall Installer

OneInstall Installer can be run in Interactive mode or Silent mode.

NOTE OneInstall Installer does not allow you to perform pre-installation tasks or text-only installations. For these tasks, use the driver kits.

NOTE OneInstaller Installer does not include NIC and FCoE drivers for LPe16202/OCe15100 adapters in NIC+FCoE mode. To install these drivers, you must use the individual driver kits.

OneInstall Installer is a self-extracting executable file that installs the following software on your system:

- FC drivers
- Emulex PLUS (ElxPlus) driver (supports the OneCommand Manager application, persistent binding, logical unit numbering (LUN) mapping and masking, and virtual port creation)
- OneCommand Manager application

NOTE The Enterprise kit for the OneCommand Manager application does not operate locally on Windows Server Core. You must install the OneCommand Manager Core Kit (command line interface only) to the Windows Server Core. Refer to the *Emulex OneCommand Manager Application for LightPulse Adapters User Guide* for installation instructions.

2.1.1 Loading OneInstall Installer Using Interactive Mode

To install the drivers using Interactive mode, perform these steps:

1. Download OneInstall Installer from the Broadcom website.
2. Navigate to OneInstall Installer in Windows Explorer.
3. Double-click **OneInstall Installer**.

   The **Welcome** screen is displayed.
4. Click **Next**.

   The **Installation Options** screen is displayed.
5. Select the drivers and applications that you want to install and click **Next**.

   A progress screen is displayed while OneInstall installer loads the selected drivers and applications. After the drivers and application software are loaded, an **Installation completed** screen is displayed.
6. Click **Finish**.

2.1.2 Loading OneInstall Installer Using Silent Mode

Silent mode installation must be run from a batch file or from the command line.

If you run OneInstall Installer from a batch file or from a command line prompt, the default Windows behavior starts OneInstall, then immediately continues with the next command. It does not wait until OneInstall has finished.

As a result, the value of %ERRORLEVEL% will always be 0 because Windows successfully started OneInstall. It does *not* reflect an accurate OneInstall exit code.

To remedy this, run OneInstall Installer setup as follows:

```
START /wait OneInstall-Setup-<version>.exe
```

```
echo %ERRORLEVEL%
```

The `START /wait` ensures that the command does not return until setup has exited. The value of %ERRORLEVEL% now accurately reflects the OneInstall Installer exit code.

2.1.2.1 Command Format

The format of the command is:

`OneInstall-Setup-<version>.exe <install-mode> <options>`

Where:

<*version*> is the version number of OneInstall Installer

<*install-mode*> is one of the following:

`/q0` – (Interactive, non-silent install) This is the default.

`/q1` – (non-Interactive install) This option displays status pages.

`/q2` – (Silent install) This option is completely silent, no pages are displayed.

`/q` – This is the same as `/q1`.

<*options*> specifies the kit, or kits, to install:

`ALL=1` – Install theFC driver and the OneCommand Manager application (default).

NOTE On a Windows Server Core system, this installs all drivers and the OneCommand Manager Core Kit.

`ALLCORE=1` – Install the FC driver and the OneCommand Manager Core Kit.

`FC=1` – Install the FC driver only.

`OCM=1` – Install the OneCommand Manager Enterprise Kit only.

NOTE On a Windows Server Core system, this installs the OneCommand Manager Core Kit.

`OMCORE=1` – Install the OneCommand Manager Core Kit only.

To install the drivers using Silent mode, perform these steps:

1. Download OneInstall Installer from the Broadcom website.
2. Open a DOS window.
3. Change the directory to the folder containing your OneInstall package.

The following are examples of Silent mode commands:

```
Start /wait OneInstall-Setup-<version>.exe  /q2  ALL=1
Start /wait OneInstall-Setup-<version>.exe  /q2  FC=1
Start /wait OneInstall-Setup-<version>.exe  /q2  OCM=1
Start /wait OneInstall-Setup-<version>.exe  /q2  ALLCORE=1
Start /wait OneInstall-Setup-<version>.exe  /q2  OCMCORE=1
```

2.2 Driver Kit Installer

Each driver kit contains and loads all the Windows drivers for a specific protocol.

- FC driver package includes ElxPlus (`elxdrvr-fc-<version>.exe`)
- FCoE driver package includes ElxPlus (`elxdrvr-fcoe-<version>.exe`)

- NIC driver package (`elxdrvr-nic-<version>.exe`)

NOTE FCoE and NIC driver packages are used on LPe16202/OCe15100 adapters only.

NOTE Updating the NIC protocol driver can temporarily disrupt operation of any NIC teams configured on the system.

2.2.1 Loading the Driver Kit

The driver kit copies the selected Emulex drivers and applications onto your computer.

NOTE This procedure does not install drivers, and no driver changes are made until you run AutoPilot Installer.

To load the driver kit, perform these steps:

1. Download the driver kit from the Broadcom website to your system.
2. Double-click the driver kit to run it.

   The **Emulex Kit Welcome** page is displayed.
3. Click **Next**.

   The **Installation Options** window is displayed.
4. Select one or both of the following options:
   - **Perform Installation of Software** – Copies the driver kit for your operating system to your computer.
   - **Unpack All Drivers** – Extracts all drivers to the current user's `Documents` folder. Select this option to perform boot from SAN installations.

   The **Operation in progress** window shows the kit file-loading progress. After the kit files are loaded, the **Installation completed** window is displayed.
5. To continue with the installation, ensure that **Start AutoPilot Installer** is selected.

2.3 AutoPilot Installer

AutoPilot Installer runs after the driver kit is loaded and the OneCommand Manager application is installed. AutoPilot Installer can be installed at these times:

- Immediately after the driver kit has been loaded
- At a later time using an interactive installation
- Through an unattended installation

AutoPilot Installer provides the following functions:

- Command line functionality – Initiates an installation from a command prompt or script. Configuration settings can be specified in the command line.
- Compatibility verification – Verifies that the driver to be installed is compatible with the operating system and platform.
- Driver installation and update – Installs and updates drivers.
- Multiple adapter installation capability – Installs drivers on multiple adapters, alleviating the requirement to manually install the same driver on all adapters in the system.

NOTE Refer to the *Emulex Software Kit Migration User Guide* for information about installing drivers on a system containing a mix of OneConnect® and LightPulse adapters.

- Driver diagnostics – Determines whether the driver is operating properly.
- Silent installation mode – Suppresses all screen output (necessary for unattended installation).

NOTE AutoPilot Installer does not allow you to install the driver if the minimum Windows service pack or Microsoft Storport driver update is not installed.

You can install the driver by using any of the following methods:

NOTE These installation methods are not mutually exclusive.

- **Hardware-first installation**. At least one Emulex adapter must be installed before you can install the Emulex drivers and utilities.
- **Software-first installation**. You can install drivers and utilities using AutoPilot Installer prior to the installation of any adapters. You do not need to specify the adapter models to be installed later. The appropriate drivers and utilities automatically load when you install the adapters.
- **Utility-Only installation**. If the drivers in the driver kit share the same version with those already installed on the system, you can reinstall or update the previously installed utility without reinstalling the drivers.
- **Text-Only installation**. Text-based installation mode is used automatically when AutoPilot Installer is run on a Windows Server Core system.
- **Network installation.** You can place the driver kit installers on a shared network drive and install them across your local area network (LAN). Network-based installation is often used with unattended installation and scripting, which allows you to configure and install the same driver and utility versions on all the hosts in a storage area network (SAN).
- **Unattended installation**. You can run the driver kit installers and AutoPilot Installer with no user interaction from a command line or script. Unattended installation works for both hardware-first and software-first installations and all driver kits. An unattended installation operates in Silent mode (also referred to as Quiet mode) and creates an extensive report file with installation status.

NOTE Complete driver and utilities documentation can be downloaded from the Broadcom website.

2.3.1 Starting Installers from a Command Prompt or Script

If a driver kit or an AutoPilot Installer is run from a command prompt or command script (batch file), the Windows command processor does not wait for the installer to run to completion. As a result, you cannot check the exit code of the installer before the next command is executed. For command line invocation, always use the `START` command with the `/wait` option, which causes the command processor to wait for the installer to finish before the command processor continues.

For additional information on command line installation and configuration parameters, see Appendix C.

2.3.2 Running a Software Installation Interactively

Two options are available when performing an installation interactively. These options assume you have already downloaded the driver kit from the Broadcom website.

- Option 1 allows you to automatically run AutoPilot Installer, which completes the driver kit loading and installation with a few mouse clicks.
- Option 2 allows you to run AutoPilot Installer separately. This option is recommended when you are:

— Changing installation settings for a limited number of systems
— Familiarizing yourself with AutoPilot Installer configuration options

2.3.2.1 Option 1: Automatically Run AutoPilot Installer

Use this option unless you have specific configuration requirements.

1. Double-click the driver kit or run it from a command line. The command line parameter `APargs` allows you to specify arguments that are automatically passed to the AutoPilot Installer command.

   A **Welcome** window is displayed with driver kit version information and Emulex contact information (see Appendix C for additional information on command line installations).
2. Click **Next** to proceed to the **Installation Options** window.

   For each installation option, the default installation location for that option is displayed. Browse to a different location, if desired.
3. Click **Install** to continue the installation.

   The **Progress** dialog is displayed.

   After all tasks complete, a **Finish** window is displayed. The **Start AutoPilot Installer** check box is automatically selected.
4. Click **Finish**.

   AutoPilot Installer runs automatically and completes one of the following installations:
   — Section 2.3.3, Hardware-First Installation or Driver Update
   — Section 2.3.4, Software-First Installation

2.3.2.2 Option 2: Run AutoPilot Installer Separately

To access these options, run AutoPilot Installer after the driver kit loading has been completed, which allows you to change the configuration options supplied to AutoPilot Installer.

1. Perform steps 1 through 3 in Option 1: Automatically Run AutoPilot Installer.
2. Clear the **Run AutoPilot Installer** check box on the **Finish** dialog.
3. Click **Finish**.

   The driver kit installer exits.

After the driver kit loading is complete, change the configuration in one of two ways:

- Change the configuration file.
- Supply parameters on the command line.

NOTE See Appendix C, AutoPilot Installer Command Line and Configuration File Parameters for additional information on either of these configuration methods.

After you have finished this step, you can run AutoPilot Installer at a later time.

4. Run AutoPilot Installer using the following command:

```
C:\Program Files\Emulex\AutoPilot Installer\APInstall.exe
```

NOTE The location of `APInstall.exe` might differ on your system, depending on your system's Program Files location. You can also specify a different location when you install the driver package.

2.3.3 Hardware-First Installation or Driver Update

The driver kit installer must be downloaded from the Broadcom website and installed before performing this installation.

NOTE Updating the NIC protocol driver can temporarily disrupt operation of any NIC teams configured on the system.

NOTE To update the Emulex protocol drivers, begin this procedure at step 2.

To perform a hardware-first installation, perform these steps:

1. Install a new Emulex adapter and power-on the system. If the **Windows Found New Hardware** wizard is displayed, click **Cancel** to exit; AutoPilot Installer performs this function.

   NOTE If there are multiple adapters in the system, the **Windows Found New Hardware** wizard is displayed multiple times. Click **Cancel** to exit the wizard each time it is displayed.

2. Run AutoPilot Installer using one of the two options listed in Section 2.3.2, Running a Software Installation Interactively.

   Consider the following:

   — If you are updating the driver, the existing port settings are used, unless otherwise specified in the configuration file. These settings are pre-selected but can be changed. Set or change the settings, then click **Next**.

   — If you are initially installing a vendor-specific version of the Emulex driver installation program, a **Driver Configuration** window may be displayed. This window includes one or more windows with questions that you must answer before continuing the installation process. In this case, answer each question and click **Next** on each window to continue.

3. Click **Next**. The installation is completed automatically.

   A dialog is displayed if Windows requires a reboot. After the installation is successful, a **Finish** window is displayed.

4. View or print a report, if desired.

   — **View Installation Report** – The installation report is a text file with current Emulex adapter inventory, configuration information, and task results.

   — **Print Installation Report** – The Windows **Print** dialog is displayed to select options for printing the installation report.

5. Click **Finish** to exit AutoPilot Installer.

6. If the system must be rebooted, you are prompted to do so as indicated in step 3; you must reboot before using the drivers or utilities.

2.3.4 Software-First Installation

The driver kit must be downloaded from the Broadcom website and loaded.

To do a software-first installation, perform these steps:

1. Run AutoPilot Installer using one of the two options listed in Section 2.3.2, Running a Software Installation Interactively.

   Figure 1 is displayed.

Figure 1 AutoPilot Installer Warning (Software-First Installation)

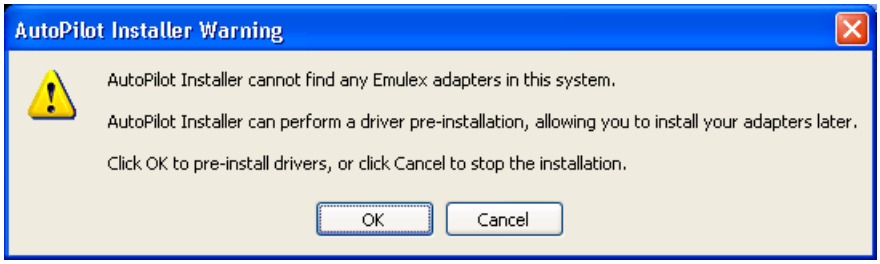

2. Click **OK**.

   A **Welcome** window is displayed.
3. Click **Next**. The installation automatically progresses.

   After the installation is successful, the **Finish** window is displayed.
4. View or print a report, if desired.
   - **View Installation Report** – The installation report is a text file with current Emulex adapter inventory, configuration information, and task results.
   - **Print Installation Report** – The Windows **Print** dialog is displayed to select options for printing the installation report.
5. Click **Finish** to exit AutoPilot Installer.

2.3.5 Text-Only Driver Installation

Text-based Installation mode is used automatically when the driver kit installer runs on a server with the Server Core installation option of Windows Server. During text-based installations, AutoPilot Installer uses a command prompt window. The driver kit installer notifies you when the driver is installed and also gives you an opportunity to stop the installation.

Whether AutoPilot Installer is launched from the command line or run as a program, Windows always starts AutoPilot Installer as a separate stand-alone task. This means that AutoPilot Installer has its own command prompt window and cannot access other windows.

2.3.6 Unattended Driver Installation

An unattended driver installation, sometimes referred to as a quiet or silent installation, requires no user input. This is useful for performing an installation remotely from a command script, or if you want to ensure that a custom configuration is not changed by a user during installation.

If in unattended installation mode, AutoPilot Installer does the following:

- Reads the configuration file
- Reads any options that might be specified on the command line, overriding the configuration file settings as appropriate
- Opens the installation report file
- Validates the operating system
- Discovers adapters and records the adapter inventory in the report file
- Verifies mandatory configuration file parameters
- Searches for drivers to install based on the LocalDriverLocation setting in the configuration file
- Verifies, if appropriate, that the selected driver is either a different type than the currently installed driver or a more recent version of the currently installed driver

- Copies the driver parameters from the configuration file into the registry for the driver's co-installer (FC and FCoE drivers only)
- Installs or updates the driver
- Rediscovers adapters and records the updated adapter inventory in the report file
- Records the final results and closes the report file

An unattended installation can be performed in two ways:

- Install the driver silently.
- Run the driver kit installer separately.

2.3.6.1 Option 1: Install the Driver Silently

Run the driver kit from a command prompt or script. Specify the `/q` (quiet) command line option. For example:

`elxdrvr-fc<version>.exe /q`

NOTE The name of the driver kit depends on the current version identifier. For other command line options, see Appendix C.

2.3.6.2 Option 2: Run the Driver Kit Installer Separately

1. Follow steps 1 to 3 in Section 2.3.2, Running a Software Installation Interactively.
2. Clear the **Run AutoPilot Installer** check box on the **Finish** dialog.
3. Choose one of the following options:
   — Run AutoPilot Installer from a command prompt or script with the silent option:
   `APInstall.exe /silent`
   — Edit the AutoPilot Installer configuration file before running AutoPilot Installer. The configuration file is typically located in:
   `C:\Program Files\Emulex\AutoPilot Installer\<driver type>\APInstall.cfg`
   Uncomment the line that sets `SilentInstallEnable` to `True`. You may also want to edit other settings in the same section of the configuration file related to unattended installations. See Appendix C.2.3 for additional information. After editing the file, run AutoPilot Installer from the **Start** menu, a command prompt, or a script.

2.3.7 Installation Failure

The two possible installation failures are described in this section.

2.3.7.1 AutoPilot Installer Failure

If AutoPilot Installer fails, the **Diagnostics** window shows that the adapter failed. If the adapter fails, perform these steps:

1. Select the adapter to view the reason why the adapter failed. The reason and suggested corrective action are displayed.
2. Perform the suggested corrective action, and run AutoPilot Installer again.

NOTE You can run `APInstall.exe` from a command prompt.

2.3.7.2 OneInstall Installer Failure

OneInstall Installer might fail for any of the following reasons:

- The operating system prerequisites have not been met.

- The individual kit installation failed. To check the installation, run the installation interactively. If you encounter error messages when you run the installation interactively, those issues would also apply to an unattended installation.
- If an individual package failed to install properly, run that package's installer directly. This method displays status and error messages that can be used to diagnose the issue. (OneInstall Installer does not provide this information because each package is installed silently.)

2.4 Manually Installing or Updating the Emulex Protocol Drivers

You can install or update the Emulex protocol drivers and utilities manually without using AutoPilot Installer.

The ElxPlus driver supports the OneCommand Manager application, persistent binding, and LUN mapping and masking.

NOTE The ElxPlus driver must be installed before you install the Emulex protocol drivers.

2.4.1 Installing the Emulex PLUS (ElxPlus) Driver for the First Time

NOTE Only one instance of the ElxPlus driver must be installed, even if you have multiple adapter ports installed in your system.

To install the ElxPlus driver from the desktop, perform these steps:

1. Run the driver kit installer, but do not run AutoPilot Installer. See Section 2.3.2, Running a Software Installation Interactively for instructions.
2. Select **Start > Settings > Control Panel > Add Hardware**. The **Add Hardware Wizard** window is displayed. Click **Next**.
3. Select **Yes, I have already connected the hardware** and click **Next**.
4. Select **Add a new hardware device** and click **Next**.
5. Select **Install the hardware that I manually select from a list (Advanced)** and click **Next**.
6. Select **Show All Devices** and click **Next**.
7. Click **Have Disk** and direct the **Device Wizard** to the location of `elxplus.inf`. If you have installed the driver installer kit in the default folder and `C:\` is your Windows system drive, the path is:

   `C:\Program Files\Emulex\AutoPilot Installer\FC\Drivers\Storport\x64\<OS>`
8. Click **OK**.
9. Select **elxplus**. Click **Next** and click **Next** again to install the driver.
10. Click **Finish**.

    The initial ElxPlus driver installation has completed. Continue with manual installation of the Storport Miniport Driver. See Section 2.4.3, Installing or Updating the FC and FCoE Storport Miniport Drivers for this procedure.

2.4.2 Updating the Emulex PLUS (ElxPlus) Driver

NOTE Only one instance of the ElxPlus driver must be installed, even if you have multiple adapter ports installed in your system.

To update an existing ElxPlus driver from the desktop, perform these steps:

1. Run the driver kit installer, but do not run AutoPilot Installer. See Section 2.3.2, Running a Software Installation Interactively for instructions.

2. Select **Start > Settings > Control Panel > Administrative Tools > Computer Management**.
3. Click **Device Manager** (left pane).
4. Click **Emulex PLUS class** (right pane) to show the ElxPlus driver entry.
5. Right-click the ElxPlus driver entry and select **Update Driver** from the menu.
6. Select **No, not this time**. Click **Next** on the **Welcome to the Hardware Update Wizard** window. Click **Next**.
7. Select **Install from a list or specific location (Advanced)** and click **Next**.
8. Select **Don't Search. I will choose the driver to install**.
9. Click **Have Disk** and direct the **Device Wizard** to the location of the driver's distribution kit. If you have installed the driver installer kit in the default folder, the path is:

   `C:\Program Files\Emulex\AutoPilot Installer\FC\Drivers\Storport\x64<OS>`
10. Click **OK**. Select Emulex PLUS.
11. Click **Next** to install the driver.
12. Click **Finish**.

    The ElxPlus driver update is finished. Continue with the manual installation of the Storport Miniport Driver.

2.4.3 Installing or Updating the FC and FCoE Storport Miniport Drivers

To update or install the FC and FCoE Storport Miniport driver from the desktop, perform these steps:

1. Select **Start > Settings > Control Panel > System**.
2. Select the **Hardware** tab.
3. Click **Device Manager**.
4. Open the **Storage Controllers** item.
5. Double-click the desired Emulex adapter.

   NOTE The driver affects only the selected adapter. If there are other adapters in the system, you must repeat this process for each adapter. All dual-channel adapter models are displayed in Device Manager as two adapters, and each adapter must be updated.
6. Select the **Driver** tab.
7. Click **Update Driver**. The **Update Driver** wizard starts.
8. Select **No, not this time**. Click **Next** on the **Welcome to the Hardware Update Wizard** window.
9. Select **Install from a list or specific location (Advanced)** and click **Next**.
10. Select **Don't search. I will choose the driver to install** and click **Next**.
11. Click **Have Disk** and direct the **Device Wizard** to the location of `oemsetup.inf`. If you have installed the driver installer kit in the default folder, the path is:

    `C:\Program Files\Emulex\AutoPilot Installer\FC(or FCoE)\Drivers\Storport\x64\<OS>`
12. Click **OK**. Select Emulex LightPulse LPX000, PCI Slot X, Storport Miniport Driver (your adapter model is displayed here).
13. Click **Next**.
14. Click **Finish**.

    The driver installation has completed. The driver will start automatically. If the adapter is connected to a SAN or data storage device, a blinking yellow light on the back of the adapter indicates a link up condition.

2.4.4 Installing or Updating the NIC Driver

NOTE The Microsoft patch KB2846340 must be installed on your system. This patch, from Microsoft's KB, is available for Windows Server 2012 on the Microsoft website.

2.4.4.1 Installing or Updating the NIC Driver on Windows Server 2012, Windows Server 2012 R2, and Windows Server 2016

1. Select **Server Manager > Dashboard > Tools > Computer Management > Device Manager**.

   NOTE Server Manager is set to open by default when booting Windows Server 2012. If it does not open automatically, you can open it with the **Server Manager** icon at the bottom left of the screen.

2. Open the **Network Adapters** item.
3. Double-click the desired Emulex adapter.
4. Select the **Driver** tab.
5. Click **Update Driver**.

   The **Update Driver** wizard starts.

6. Click **Browse my computer for driver software**.

   The driver affects only the selected adapter. If there are other adapters in the system, you must repeat this process for each adapter. All dual-channel adapter models are displayed in the Device Manager as two adapters, therefore, you must update each adapter.

7. Click **Let me pick from a list of device drivers on my computer**.
8. Select the network adapter that matches your hardware and click **Have Disk**.
9. Direct the **Device Wizard** to the location of `lpnic.inf`. Select the desired `oemsetup.inf` file and click **Open**.

   If you have installed the driver installer kit in the default folder, the path is:

   `C:\Program Files\Emulex\AutoPilot Installer\NIC\Drivers\NDIS\x64\<os>`

10. Click **Next**.
11. After the device driver finishes installing, click **Close**.

    The driver installation has completed. The driver will start automatically.

2.5 Removing the Emulex Driver Kits and Drivers

This section details procedures to uninstall the driver kits.

2.5.1 Uninstalling the Emulex Driver Kits

NOTE If you uninstall the Emulex driver kit, AutoPilot Installer is automatically uninstalled.

2.5.1.1 Uninstalling an Emulex Driver Kit on Windows Server 2012

To uninstall a driver kit on a Windows Server 2012 system, perform these steps:

1. Select **Start > Control Panel**.
2. From the Control Panel, select **Programs > Uninstall a Program**.

3. Select one of the following in the program list and click the **Uninstall** icon in the tool bar above the program list. If you have User Access Control enabled, click **Continue** when asked for permission.
   - `Emulex /FC kit-2.xx.xxx`
   - `Emulex/FCoE kit-2.xx.xxx`
   - `Emulex/NIC kit-4.xx.xxx`
4. Click **Yes** when prompted to remove the kit. After the kit is removed from the system, click **OK**.

2.5.1.2 Uninstalling an Emulex Driver Kit on a Server Core System

To uninstall a driver kit on a Server Core system, perform these steps:

1. From the system prompt, navigate to the Program Files folder.
2. Navigate to Emulex\AutoPilot Installer.
3. Run the following batch files:
   - `Uninstall_cna_kit.bat`
   - `Uninstall_fc_kit.bat`
   - `Uninstall_nic_kit.bat`

The driver files are removed from the system.

On all platforms, the reports folder in the `Emulex\AutoPilot` Installer folder is not removed, so you can still view installation history and the drivers that have been installed on the system. You can delete the reports folder at any time.

2.5.2 Uninstalling the Emulex Drivers

The Emulex Storport Miniport and ElxPlus drivers are uninstalled using the Device Manager.

2.5.2.0.1 Uninstalling an Emulex Storport Miniport Driver

To uninstall the Emulex Storport Miniport driver, perform these steps:

1. Select **Start > All Programs > Administrative Tools > Computer Management**.
2. Click **Device Manager**.
3. Double-click the adapter from which you want to remove the Storport Miniport driver. A device-specific console window is displayed. Select the **Driver** tab.
4. Click **Uninstall** and click **OK** to uninstall.

2.5.2.0.2 Uninstalling an ElxPlus Driver

NOTE Uninstall the ElxPlus driver only if all adapters and installations of Emulex miniport drivers are uninstalled.

To uninstall the ElxPlus driver, perform these steps:

1. Select **Start > All Programs > Administrative Tools > Computer Management**.
2. Click **Device Manager**.
3. Click the **Emulex PLUS driver** class.
4. Right-click the Emulex driver and click **Uninstall**.
5. Click **OK** in the **Confirm Device Removal** window.

2.5.2.1 Uninstalling the Emulex Driver on Windows Server 2012

The Emulex Storport Miniport and ElxPlus drivers are uninstalled using the Device Manager.

NOTE On Windows 2012 and Windows 2012 R2, after the message `Warning - you are about to uninstall this device from your system` is displayed, you must select the checkbox **Delete the software for this device** to uninstall the driver.

2.5.2.1.1 Uninstalling the Emulex Storport Miniport Driver

To uninstall the Emulex Storport Miniport driver in Windows Server 2012, perform these steps:

1. Select **Server Manager > Dashboard > Tools > Computer Management > Device Manager**.
2. Double-click the adapter from which you want to remove the Storport Miniport driver. A device-specific console window is displayed.
3. Select the **Driver** tab.
4. Click **Uninstall** and click **OK** to uninstall.

2.5.2.1.2 Uninstalling the ElxPlus Driver

NOTE Uninstall the ElxPlus driver only if all adapters and installations of Emulex miniport drivers are uninstalled.

To uninstall the ElxPlus driver, perform these steps.

1. Select **Server Manager > Dashboard > Tools > Computer Management > Device Manager**.
2. Click the **Emulex PLUS driver** class.
3. Right-click the Emulex driver and click **Uninstall**.
4. Click **OK** in the **Confirm Device Removal** window.

Chapter 3: Configuration

3.1 FC and FCoE Driver Configuration

The Emulex Storport Miniport driver has many options that you can modify to provide different behavior. You can set Storport Miniport driver parameters using the OneCommand Manager application. Refer to the *Emulex OneCommand Manager Application for LightPulse Adapters User Guide* for information on using this utility to configure the driver.

3.1.1 Configuring FC and FCoE Driver Parameters

Table 1, Storport Miniport Driver Parameters, provides information, such as the range of permissable values and the factory defaults. Parameters can be entered in decimal or hexadecimal format.

A parameter has one of the following activation requirements:

- Dynamic – The change takes effect while the system is running.
- Reset – An adapter reset from the utility is required before the change takes effect.
- Reboot – A reboot of the entire machine is required before the change takes effect. In this case, you are prompted to perform a reboot when you exit the utility.

NOTE If you are creating custom unattended installation scripts, any driver parameter can be modified and included in the script.

Most parameters default to a setting that optimizes adapter performance.

Table 1 Storport Miniport Driver Parameters

Parameter	Definitions	Activation Requirement	Notes
AutoMap	AutoMap controls the way targets are assigned SCSI IDs. Discovered targets are assigned persistent SCSI IDs according to the selected binding method. Persistent bindings do not take effect with the driver in stand-alone mode. ■ 0 = Automap is disabled. The OneCommand Manager application persistently sets the SCSI address of a discovered FCP-capable FC node (target). ■ 1 = Automap by WWNN. ■ 2 =Automap by WWPN. ■ 3 = Automap by DID. Value: 0–3 Default = 2	Reboot	
Class	Class selects the class of service on FCP commands. ■ If set to 2, class of service is 2. ■ If set to 3, class of service is 3. Value: 2–3 Default = 3	Dynamic	Supported on FC only.
CoalesceMsCnt	CoalesceMsCnt specifies wait time in milliseconds to generate an interrupt response if CoalesceRspCnt has not been satisfied. Zero specifies an immediate interrupt response notification. A nonzero value enables response coalescing at the specified interval in milliseconds. Value: 0–63 (decimal) or 0x0–0x3F (hexadecimal) Default = 0 (0x0)	Reset	Supported on FC only. Supported on LPe12000-series adapters only.
CoalesceRspCnt	CoalesceRspCnt specifies the number of response entries that trigger an interrupt response. Value: 0–255 (decimal) or 0x1–0xFF (hexadecimal) Default = 8 (0x8)	Reset	Supported on FC only. Supported on LPe12000-series adapters only.
ConfigScale	ConfigScale sets the memory footprint profile in accord with the anticipated use case on a per-port basis. ConfigScale is always set at 4. The maximum transfer size is set according to the value of the ExtTransferSize parameter. Default = 4	Reboot	Not supported on LPe12000-series adapters.
DriverTraceMask	The DriverTraceMask parameter is only available on operating systems that support extended system event logging. ■ If set to 0 = The parameter is disabled. ■ If set to 1 = Error events logging is enabled. ■ If set to 4 = Warning events logging is enabled. ■ If set to 8 = Informational events logging is enabled. The values can be masked to generate multiple levels of event logging. Values: 0, 1, 4, and 8. Default = 0.	Dynamic	

Table 1 Storport Miniport Driver Parameters (Continued)

Parameter	Definitions	Activation Requirement	Notes
EnableAck0	Set to 1 to force sequence rather than frame level acknowledgement for class 2 traffic over an exchange. This applies to FCP data exchanges on IREAD and IWRITE commands. Value: 0–1 (decimal) Default = 1	Reset	Supported on FC only.
EnableAUTH	EnableAUTH enables fabric authentication. This parameter requires the authentication to be supported by the fabric. Authentication is enabled if this value is set to 1. Value: 0–1 Default = 0	Reboot	Supported on LPe12000-series adapters only.
EnableFDMI	If set to 1, enables management server logon on fabric discovery, which allows FDMI to operate on switches that have FDMI-capable firmware. FDMI operates as FDMI-1. If set to 2, FDMI operates as FDMI-2. If set to 0, FDMI is disabled. Value: 0–2 (decimal) Default = 2	Reset	
EnableMDSD	If set to 1, Cisco Multilayer Director Switch (MDS) diagnostics are enabled. Value: 0–1 Default = 0 The parameter should be disabled (set to 0) after diagnostics are complete.	Reset	
EnableNPIV	If set to 1, enables NPIV. Requires NPIV-supported firmware for the adapter. Value: 0–1 Default = 1 (enabled) **NOTE** SLIMode=2 does not support NPIV mode on LPe12000 adapters.	Reboot	
EnableXLane	EnableXLane enables ExpressLane™ If set to 1, enables the driver to set the CS_CTL priority according to the value of XLanePriority driver parameter. Value: 0–1 Default = 0	Reboot	Not supported on LPe12000-series adapters.
ExtTransferSize	ExtTransferSize is an initialization-time parameter that affects the maximum SGL that the driver can handle, which determines the maximum I/O size that a port will support. ■ If set to 0 = The maximum default transfer size is 512 KB for all controller models. ■ If set to 1= The maximum transfer size is 1 MB. ■ If set to 2 = The maximum transfer size is 2 MB. ■ If set to 3 = The maximum transfer size is 4 MB. Value: 0–3 Default = 0 (disabled)		

Table 1 Storport Miniport Driver Parameters (Continued)

Parameter	Definitions	Activation Requirement	Notes
FrameSizeMSB	FrameSizeMSB controls the upper byte of receive FrameSize if issued in PLOGI. This allows the FrameSize to be constrained on 256-byte increments from 256 (1) to 2048 (8). Value: 0–8 Default = 0	Reset	
InitTimeout	Determines the number of timeout seconds during driver initialization for the link to come up. If the link fails to come up by the InitTimeout, driver initialization exits but is still successful. If the link comes up before the value specified by InitTimeout, the driver sets double the amount for discovery to complete. Value: 5–30 seconds or 0x5–0x1E (hexadecimal) Default = 15 seconds (0xF)	Reboot	
LimTransferSize	Limits the maximum transfer size to selectable values if this parameter is nonzero. Values: ■ 0 = Port (default) ■ 1 = 64 KB ■ 2 = 128 KB ■ 3 = 256 KB	Reboot	
LinkSpeed	LinkSpeed has significance only if the adapter supports speeds other than 1 Gb/s. Value: Auto-select, 2 Gb/s, 4 Gb/s, and 8 Gb/s Default = Auto-select **NOTE** Setting this option incorrectly can cause the adapter to fail to initialize. **NOTE** If you configure the link speed in a BIOS utility, the link speed might be overridden by the Emulex driver for Windows according to its LinkSpeed setting. To avoid this issue, configure the link speed in both the Emulex driver for Windows and the Boot BIOS or UEFI driver.	Reset	Supported on FC only. Supported on LPe12000-series adapters only.
LinkTimeOut	LinkTimeOut applies to a private loop only. A timer is started on all mapped targets using the LinkTimeOut value. If the timer expires before discovery is resolved, commands issued to timed-out devices return a SELECTION_TIMEOUT. The Storport driver is notified of a bus change event, which leads to the removal of all LUNs on the timed-out devices. Value: 1–500 seconds or 0x0–0xFE (hexadecimal) Default = 30 (0x1E)	Dynamic	

Table 1 Storport Miniport Driver Parameters (Continued)

Parameter	Definitions	Activation Requirement	Notes
LogErrors	LogErrors determine the minimum severity level required to enable entry of a logged error into the system event log. Errors are classified as severe, malfunction, or command level. A severe error requires user intervention to correct a firmware or adapter issue. An invalid link speed selection is an example of a severe error. A malfunction error indicates a system problem, but user intervention is not required. An invalid fabric command type is an example of a malfunction error. An object allocation failure is an example of a command error. ■ If set to 0 = All errors are logged. ■ If set to 1 = Command level errors are logged. ■ If set to 2 = Malfunction errors are logged. ■ If set to 3 = Severe errors are logged. Value: 0–3 Default = 3	Dynamic	
NodeTimeout	The node timer starts when a node (that is, a discovered target or adapter) becomes unavailable. If the node fails to become available before the NodeTimeout interval expires, the operating system is notified so that any associated devices (if the node is a target) can be removed. If the node becomes available before the NodeTimeout interval expires, the timer is canceled and no notification is made. Value: 1–255 seconds or 0x0–0xFF (hexadecimal) Default = 30 (0x1E)	Dynamic	
QueueDepth	QueueDepth requests per LUN/target (see the QueueTarget parameter). If you expect the number of outstanding I/Os per device to exceed 32, you must increase to a value greater than the number of expected I/Os per device (up to a value of 254). If the QueueDepth value is set too low, a performance degradation can occur due to driver throttling of its device queue. QueueDepth supports more than 1000 outstanding commands per port. Value: 1–254 or 0x1–0xFE (hexadecimal) Default = 32 (0x20)	Dynamic	
QueueTarget	QueueTarget controls I/O depth limiting on a per-target or per-LUN basis. ■ If set to 0 = Depth limitation is applied to individual LUNs. ■ If set to 1 = Depth limitation is applied across the entire target. Value: 0–1 or 0x0–0x1 (hexadecimal) Default = 0 (0x0)	Dynamic	

Table 1 Storport Miniport Driver Parameters (Continued)

Parameter	Definitions	Activation Requirement	Notes
RmaDepth	RmaDepth sets the remote management buffer queue depth. The greater the depth, the more concurrent management controls can be handled by the local node. Value: 8–64, or 0x8–0x40 (hexadecimal) Default = 16 (0x10) **NOTE** The RmaDepth driver parameter pertains to the functionality of the OneCommand Manager application.	Reboot	
ScanDown	■ If set to 0 (= lowest AL_PA) = Lowest physical disk (ascending AL_PA order). ■ If set to 1 (= highest AL_PA) = Lowest physical disk (ascending SEL_ID order). Value: 0–1 Default = 1 **NOTE** This option applies to private loop only in DID mode.	Reboot	Supported on FC only.
SLIMode	■ If set to 0 = Autoselect firmware, use the latest firmware installed. ■ If set to 2 = Implies running the adapter firmware in SLI-2 mode. ■ If set to 3 = Implies running the adapter firmware in SLI-3 mode. Value: 0, 2, and 3 Default = 0	Reboot	Supported on LPe12000-series adapters only.
SrbTimeout	SrbTimeout limits the SCSI timeout value to 60 seconds if set to 1 or enabled. This parameter is a non-displayed parameter that must be set manually in the registry. This parameter alters the I/O timeout behavior, where an I/O will be returned in a maximum timeout of 60 seconds on long I/O timeouts. ■ If set to 1 = Enabled ■ If set to 0 = Disabled Values: 0–1 Default = 0		

Table 1 Storport Miniport Driver Parameters (Continued)

Parameter	Definitions	Activation Requirement	Notes
`Topology`	■ If set to `0` (0x0) = FC-AL. ■ If set to `1` (0x1) = PT-PT fabric. ■ If set to `2` (0x2) = *FC-AL first, then attempt PT-PT. ■ If set to `3` (0x3) = *PT-PT fabric first, then attempt FC-AL. **NOTE** Topology fail-over requires firmware version v3.20 or higher. If the firmware does not support topology failover, options 0,2 and 1,3 are analogous. Value: 0–3 Default = 2 (0x2)	Reset	Supported on LPe12000-series and LPe16000-series adapters only. Not supported on LPe16202 adapters.
`TraceBufSiz`	`TraceBufSiz` sets the size in bytes for the internal driver trace buffer, which acts as an internal log of the driver's activity. Value: 250000–2000000 or 0x3D090–0x1E8480 (hexadecimal). Default = 250000 (0x3D090)	Reboot	
`XLanePriority`	`XLanePriority` sets the frame priority level for the LUN. Refer to the switch vendor administration guide to set the value. Value: 0–7F (hexadecimal) Default = 0	Dynamic	Not Supported on LPe12000-series adapters.

3.1.2 Server Performance with LPe12000-Series Adapters

3.1.2.1 I/O Coalescing

I/O Coalescing is enabled and controlled by two driver parameters: `CoalesceMsCnt` and `CoalesceRspCnt`. The effect of I/O Coalescing depends on the CPU resources available on the server. With I/O Coalescing turned on, interrupts are batched, which reduces the number of interrupts and maximizes the number of commands processed with each interrupt. For heavily loaded systems, this setting provides better throughput.

With I/O Coalescing turned off (the default setting), each I/O processes immediately, one CPU interrupt per I/O. For systems with light loads, the default setting provides better throughput. Table 2 shows recommendations based upon the number of I/Os per adapter.

Table 2 Recommended Settings for I/O Coalescing

I/Os per Second	Suggested `CoalesceMsCnt`	Suggested `CoalesceRspCnt`
I/Os < 10,000	0	8
10,000 < I/Os < 18,000	1	8
18,000 < I/Os < 26,000	1	16
I/Os > 26,000	1	24

3.1.2.1.1 CoalesceMsCnt

The `CoalesceMsCnt` parameter controls the maximum elapsed time in milliseconds that the adapter waits before it generates a CPU interrupt. The value range is 0–63 (decimal) or 0x0–0x3F (hexadecimal). The default is 0 and disables I/O Coalescing.

NOTE A port reset is required to make changes to `CoalesceMsCnt` and `CoalesceRspCnt`.

3.1.2.1.2 CoalesceRspCnt

The `CoalesceRspCnt` parameter controls the maximum number of responses to batch before an interrupt generates. If `CoalesceRspCnt` expires, an interrupt generates for all responses collected up to that point. With `CoalesceRspCnt` set to less than 2, response coalescing is disabled, and an interrupt triggers for each response. The value range for `CoalesceRspCnt` is 1–255 (decimal) or 0x1–0xFF (hexadecimal). The default value is 8.

NOTE A port reset is required to make changes to `CoalesceMsCnt` and `CoalesceRspCnt`.

3.1.2.2 Performance Testing

Three driver parameters must be considered (and perhaps changed from the default) for better performance testing: `QueueDepth`, `CoalesceMsCnt`, and `CoalesceRspCnt`.

3.1.2.2.1 QueueDepth

If the number of outstanding I/Os per device is expected to exceed 32, increase this parameter to a value greater than the number of expected I/Os per device, to a maximum of 254. The `QueueDepth` parameter defaults to 32. If the default setting is not a high enough value, performance degradation might occur due to Storport throttling its device queue.

3.1.2.2.2 CoalesceMsCnt

`CoalesceMsCnt` defaults to 0. If you are using a performance evaluation tool, such as IOMETER, and if you expect the I/O activity to be greater than 8000 I/Os per second, set `CoalesceMsCnt` to `1` and reset the adapter or reboot the system.

3.1.2.2.3 CoalesceRspCnt

`CoalesceRspCnt` defaults to 8. For all other values up to the maximum of 63, the adapter does not interrupt the host with a completion until either `CoalesceMsCnt` milliseconds has elapsed or `CoalesceRspCnt` responses are pending. The values of these two driver parameters reduces the number of interrupts per second, which improves overall CPU utilization. However, a point exists where the number of I/Os per second is small relative to `CoalesceMsCnt`, and this situation will slow down the completion process, causing performance degradation.

3.1.2.2.4 Examples

Test scenario 1:

- IOMETER runs with an I/O depth of 1 I/O per device in a small-scale configuration (16 devices). In this case, the test does not exceed the adapter's performance limits, and the number of I/Os per second are in the low thousands.
- Recommendation: Set `CoalesceMsCnt` to `0` (or use the default value).

Test scenario 2:

- IOMETER runs with an I/O depth of 48 I/Os per device in a small-scale configuration (16 devices).
- Recommendation: Set `QueueDepth` to be greater than 48 (for example, 64).

3.1.3 Server Performance with FC Drivers

3.1.3.1 Performance Testing

The `QueueDepth` parameter must be considered (and perhaps changed from the default) for better performance testing.

If the number of outstanding I/Os per device is expected to exceed 32, increase this parameter to a value greater than the number of expected I/Os per device, to a maximum of 254. The `QueueDepth` parameter defaults to 32. If the default setting is not a high enough value, performance degradation might occur due to Storport throttling its device queue.

Test scenario:

- IOMETER is running with an I/O depth of 48 I/Os per device in a small-scale configuration (16 devices).
- Recommendation: Set `QueueDepth` to be greater than 48 (for example, 64).

3.2 NIC Driver Configuration

NOTE This section applies to LPe16202/OCe15100 adapters in NIC+FCoE mode only.

3.2.1 Configuring NIC Driver Options

The Windows Server NIC driver supports configurable driver options through the **Advanced Property** page in the **Windows Device Manager**. For information on how to configure the options through the **Advanced Property** page, see Section 3.2.2.1, Modifying Advanced Properties.

For additional information on NIC driver options, see Section 3.2.5, Network Driver Performance Tuning.

You can also set configurable driver options using Microsoft PowerShell on Windows Server 2012 and Windows Server 2016.

Use the `Get-NetAdapter` PowerShell command to list all available adapters in the system. The `Get-Help <cmdl> -full` command returns descriptions and help for the cmdlets.

Use the following PowerShell commands to get and set driver parameter values.

To get the driver parameter value, perform this step:

```
Get-NetAdapterAdvancedProperty -Name <adapter name> -AllProperties
-RegistryKeyword <registry keyword>
```

Example:

```
Get-NetAdapterAdvancedProperty -Name ""SLOT 6 Port 1"" -AllProperties
-RegistryKeyword *RSS
```

To set the driver parameter value, perform this step:

```
Set-NetAdapterAdvancedProperty -Name <adapter name> -AllProperties
-RegistryKeyword <registry keyword> -RegistryValue <valid registry value>
```

Example:

```
Set-NetAdapterAdvancedProperty -Name ""SLOT 6 Port 1"" -AllProperties
-RegistryKeyword *RSS -RegistryValue 1
```

NOTE Refer to the documentation that accompanies the Windows Server 2012 and Windows Server 2016 operating systems for additional information on using PowerShell.

Table 3 lists the NIC driver options.

Table 3 Windows Server NIC Driver Options

Option Name	Acceptable Values	Supported Operating Systems	Definition
Class of Service (802.1p)	Automatic Priority (default) Filtered Priority User Priority Disable Priority	Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	The following modes are supported for selecting 802.1p priority tags: ■ Automatic Priority – The DCBX standard allows the network adapter to negotiate priority class usage with DCBX-aware endpoints, such as switches or network cards. If the peer indicates that priority pause is supported for a nonzero priority, the NIC automatically inserts the default priority in all transmitted packets. This mode is the default mode, which allows priority pause to operate for both storage and network traffic. ■ Filtered Priority – This mode coerces the user priorities in each packet to avoid sending packets on the network function that might disrupt the adapter's storage traffic. The network device uses the next lower priority if a conflict exists. This mode is useful if multiple network priorities are necessary. Only a limited number of classes are supported for priority pause, so typically, it does not function optimally in this mode. ■ User Priority – This mode allows any user-specified priority value and must be limited to cases where storage functions are not used. ■ Disable Priority – The adapter always transmits either untagged packets, or VLAN ID (802.1q) tagged packets with a priority value (802.1p) of zero.
Flow Control	Disabled RX and TX Enabled (default) RX Enable/TXDisable TX Enable/RX Disable	Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	Flow control is almost always advantageous to avoid packet drops on the network. The switch or network peer must also have flow control enabled. The IEEE 802.3x Ethernet specification defines a control frame between peers that can request a pause in packet transmissions. Flow control allows one system to request a temporary halt of all incoming traffic if receive buffer space is exhausted. **NOTE** The network device can be configured to respond to pause frames (RX Enable) and also to send pause frames (TX Enable).
Interrupt Moderation	0 (None) 8 (Static 90k Int/s) 9 (Static 70k Int/s) 10 (Static 50k Int/s) 11 (Static 40k Int/s) 12 (Static 30k Int/s) 2 (Static 25k Int/s) 13 (Static 20k Int/s) 14 (Static 15k Int/s) 15 (Static 10k Int/s) 16 (Static 5k Int/s) 4 (Adaptive) (default)	Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	The network device uses interrupt moderation algorithms to reduce the total amount of CPU cycles spent processing interrupts, which increases efficiency for the system. However, interrupt moderation increases the latency of each send and receive. It should only be disabled when short latencies are more important than efficient CPU utilization. The No Moderation setting disables all delays to minimize latency. The Static Moderation setting uses a constant interrupt delay to avoid any spikes in the interrupt rate. The Adaptive (default) setting causes the driver to dynamically maintain a target interrupt rate. The Adaptive setting value is controlled by a dynamic algorithm that scales well for various adapter link speeds.

Table 3 Windows Server NIC Driver Options (Continued)

Option Name	Acceptable Values	Supported Operating Systems	Definition
IP Checksum Offload (IPv4)	Disabled RX and TX Enabled (default) RX Enabled TX Enabled	Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	This option offloads the transmit and the receive IPv4 checksum computation. Offloading checksums increases system efficiency.
Large Send Offload v1 (IPv4)	Disabled Enabled (default)	Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	Large Send Offload allows the NIC hardware to segment large TCP packets (up to 64 KB) into smaller packets (≤ Packet Size) that can be transmitted. This segmentation increases transmit efficiency for TCP applications that send large buffers. During segmentation, the hardware computes the IPv4 and TCP checksums for each individual packet. The Windows Version 1 LSO supports only IPv4.
Large Send Offload v2 (IPv4)	Disabled Enabled (default)	Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	Large Send Offload allows the NIC hardware to segment large TCP packets (up to 64 KB) into smaller packets (≤ Packet Size) that can be transmitted. This segmentation increases transmit efficiency for TCP applications that send large buffers. During segmentation, the hardware computes the IPv4 and TCP checksums for each individual packet. The Windows Version 2 LSO supports larger offload sizes.
Large Send Offload v2 (IPv6)	Disabled Enabled (default)	Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	Large Send Offload allows the NIC hardware to segment large TCP packets (up to 64 KB) into smaller packets (less than the MTU) that can be transmitted. This segmentation increases transmit efficiency for TCP applications that send large buffers. During segmentation, the hardware computes the TCP checksums for each individual packet. IPv6 support requires Windows Version 2 LSO.
Maximum Number of RSS Processors	Min. 0 Max. The number of CPU cores installed on your system	Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	This property sets the maximum number of processors that can be used for RSS.
Maximum Number of RSS Queues	Min. 1 Max. 16 Default: 8	Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	If RSS is enabled, this parameter controls the number of receive queues. Typically, this option is left at the maximum value. Windows reduces the number of queues as necessary based on the number of installed CPU cores. This value can be reduced during performance tuning for a particular application. It is possible that system performance might improve by limiting the number of RSS queues.
Maximum RSS Processor Number	Min. 1 Max. The number of CPU cores installed on your system	Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	This parameter sets the maximum processor number for the RSS CPUs. This value is the highest processor number of any processors from the RSSMaxProcGroup parameter.

Table 3 Windows Server NIC Driver Options (Continued)

Option Name	Acceptable Values	Supported Operating Systems	Definition
`Network Address`	Valid MAC address The default setting is `None`	Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	This option overrides the permanent MAC address for the interface. The MAC address must follow this format *XX:XX:XX:XX:XX:XX*, where *X* is a hexadecimal digit (0–9 or A–F). ■ The address cannot be a multicast address, which has the lowest bit in the first byte set. ■ The address cannot be all zeros. For example, 01:00:00:00:00:00 is not valid, while 02:00:00:00:00:00 is valid.

`Packet Size`	`1514` (default)
	`9014`
	`8222`
	`4088`
		Windows Server 2012
		Windows Server 2012 R2
		Windows Server 2016
			This option configures the packet size for OneConnect NIC adapters only.
			This option determines the maximum packet size transmitted and received on the interface. A 1514-byte frame size is standard, while larger packets are called jumbo frames.
			Using a higher frame size is generally more efficient, but it uses more system memory. A larger frame size also requires support on the network switch.
			Jumbo frames are IPv4-only frames; IPv6 packets will be fragmented by LSO. Switches and the peer must be configured to accept the specified packet size or the size will be negotiated to the common smallest size.

`Performance Tuning`	`Maximum performance` (default)
	`Statically balanced`
	`Dynamically balanced`
		Windows Server 2012
		Windows Server 2016
			This parameter selects the driver algorithm for performance tuning, which allows you to balance raw networking throughput with overall system fairness among multiple devices and applications.
			■ Maximum performance – This mode maximizes the network performance for this adapter. This mode is the recommended mode. However, in systems with a large number of network or storage adapters, this mode can limit the performance of other devices.
			■ Statically balanced – This mode configures the network adapter to throttles CPU usage in all cases, allowing more balance among hardware devices and applications. If system responsiveness is poor, this mode can improve the overall system behavior.
			■ Dynamically balanced – Dynamic balancing adjusts the network adapter's performance based on system metrics, such as CPU usage. This mode can aggressively limit performance for the most stressful networking applications to ensure that all network adapters can share limited computer resources, yet this mode can maintain maximum performance if the system has resources available.

Table 3 Windows Server NIC Driver Options (Continued)

Option Name	Acceptable Values	Supported Operating Systems	Definition
Preferred NUMA Node	Not present or a value from 0–65535. Optional. No default setting is set.	Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	Most modern multi-socket servers have separate memory controllers for each CPU socket. These systems have NUMA latencies for a given CPU core to access the local versus remote memory node. By setting this property, the driver attempts to use both memory and CPU cores from the given NUMA node. If the Preferred NUMA node is not set, the driver uses the preferred NUMA node as specified by the computer's BIOS. For best performance, the network applications must use memory and CPU affinity from the same NUMA node. This level of tuning is primarily noticeable when multiple adapters are running.
Receive Buffers	64–32768, inclusive The default value is 896.	Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	This option determines the number of Ethernet receive buffers allocated per receive queue. This number can be adjusted by the driver as needed.
Receive CPU	Not Present or a value from 0 through (number of CPUs on the system – 1). Optional. A default setting is not available.	Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	This option sets the logical CPU used for processing the non-RSS receive packets. By default, the driver intelligently chooses a CPU in the system, so this parameter must only be used for advanced performance tuning. RSS packets are processed by the set of RSS CPUs provided by the Windows operating system.
Receive Side Scaling	Disabled Enabled (default)	Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	Support for multiple RSS queues if enabled. RSS scales receive processing over multiple CPUs in parallel. This scaling typically improves application performance; however, it tends to increase CPU usage on low-end machines. For additional PCI functions, RSS does not appear in the **Properties** list.
Recv Segment Coalescing (IPv4)	Disabled Enabled (default)	Windows Server 2012 Windows Server 2012 R2	RSC merges multiple TCP segments and identifies them as a single coalesced unit to the operating system's TCP/IP stack. This option reduces the per-packet receive processing overhead and CPU usage if standard 1514-byte-sized frames are in use. **NOTE** If checksum offloads are disabled, RSC must also be disabled. RSC depends on checksum offloads for better performance. **NOTE** Both RSC (IPv4) and RSC (IPv6) are coerced to 0 if TCP Connection Offload (IPv4) is enabled.

Table 3 Windows Server NIC Driver Options (Continued)

Option Name	Acceptable Values	Supported Operating Systems	Definition
Recv Segment Coalescing (IPv6)	Disabled Enabled (default)	Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	RSC merges multiple TCP segments and identifies them as a single coalesced unit to the operating system's TCP/IP stack. This option reduces the per-packet receive processing overhead and CPU usage if standard 1514 byte-sized-frames are in use. **NOTE** If checksum offloads are disabled, RSC must also be disabled. RSC depends on checksum offloads for better performance. **NOTE** Both RSC (IPv4) and RSC (IPv6) are coerced to 0 if TCP Connection Offload (IPv4) is enabled.
RSS Base Processor Group	Min.1 Max. 63	Windows Server 2012 Windows Server 2016 Windows Server 2016	This option defines the base processor group for the RSS queues on the network adapter. A processor group contains 64 logical processors. This value can be modified with RSS Base Processor Number to explicitly select the desired RSS processors for the adapter.
RSS Base Processor Number	Min. 1 Max. 63	Windows Server 2012 Windows Server 2016	This option defines the base processor number for the RSS queues on the network adapter within the given processor group. A processor group contains 64 logical processors, so this value ranges from 0 to 63. This value can be modified with RSS Base Processor Group to explicitly select the desired RSS processors for the adapter.
RSS Max Processor Group	Min. 0 Max. The number of processor groups present on your system	Windows Server 2012 Windows Server 2016	RSS Max Processor Group allows you to set the maximum number of processor groups for the RSS CPUs.
RSS Profile	Closest processor (default) Closest processor static NUMA scaling NUMA scaling static Conservative scaling	Windows Server 2012 Windows Server 2016	The RSS Profile setting determines the RSS load balancing profile implemented by Microsoft for this network adapter. The Closest Processor settings tend to localize the RSS CPUs to one NUMA node, allowing the device driver to allocate memory from the local node. The NUMA Scaling settings use all NUMA nodes on the system, and the memory allocation is not specific to a particular node. The driver ignores the Preferred NUMA Node setting.

Table 3 Windows Server NIC Driver Options (Continued)

Option Name	Acceptable Values	Supported Operating Systems	Definition
SR-IOV	Disabled (default) Enabled	Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	SR-IOV enables the adapter to allocate virtual PCI functions for each virtual machine in Hyper-V. **NOTE** The virtual switch and virtual network adapter must have SR-IOV enabled in the Hyper-V Manager. SR-IOV requires a platform with IOMMU virtualization (VT-d, AMD-Vi). If using SR-IOV, the Emulex NIC driver must be installed on each virtual function within the virtual machine. SR-IOV provides a direct hardware interface from the virtual machine to the networking adapter, which reduces latency and improves performance. The Windows Server 2012 and Windows Server 2012 R2 SR-IOV architecture establishes each Emulex virtual NIC with a corresponding emulated NIC. This architecture allows the virtual machine to seamlessly failover to the emulated NIC if SR-IOV is disabled. It also allows Live Migration to another system, regardless of the installed NIC hardware.
TCP Checksum Offload (IPv4)	0 (Disabled) 1 (TX Enabled) 2 (RX Enabled) 3 (RX and TX Enabled) (default)	Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	TCP Checksum Offload (IPv4) offloads the transmit or receive IPv4 TCP checksum computation. Offloading checksums increases system efficiency.
TCP Checksum Offload (IPv6)	0 (Disabled) 1 (TX Enabled) 2 (RX Enabled) 3 (RX and TX Enabled) (default)	Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	TCP Checksum Offload (IPv6) offloads the transmit or receive IPv6 TCP checksum computation. Offloading checksums increases system efficiency.
Transmit Buffers	64–256, inclusive The default setting is 256.	Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	Transmit Buffers sets the number of Ethernet transmits that might be posted to the hardware at any given time. The default value is sufficient to achieve maximum performance. Reducing this value conserves system memory.
Transmit CPU	Not Present or a value from 0 through (number of CPUs – 1). Optional. A default setting is not available.	Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	This option sets the CPU to be used to process transmit completions. By default, the driver intelligently chooses a CPU in the system, so this parameter must only be set for advanced performance tuning.
Transmit Side Scaling (TSS)	Enabled Disabled	Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	Transmit Side Scaling distributes transmit completions to be processed on multiple CPUs in parallel. It uses the RSS CPU table for distribution, and therefore, requires RSS to be enabled.
UDP Checksum Offload (IPv4)	Disabled RX and TX Enabled (default) RX Enabled TX Enabled	Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	UDP Checksum Offload settings offload the transmit or receive IPv4 UDP checksum computation. Offloading checksums increases system efficiency.

Table 3 Windows Server NIC Driver Options (Continued)

Option Name	Acceptable Values	Supported Operating Systems	Definition

`UDP Checksum Offload (IPv6)`	`Disabled`
	`RX and TX Enabled` (default)
	`RX Enabled`
	`TX Enabled`
		Windows Server 2012
		Windows Server 2012 R2
		Windows Server 2016
			`UDP Checksum Offload` settings offload the transmit or receive IPv6 UDP checksum computation.
			Offloading checksums increases system efficiency.

`Virtual Machine Queues`	`Enabled` (default) `Disabled`	VMQs require: Windows Server 2012 Windows Server 2012 R2 Windows Server 2016 with Hyper-V	VMQs are dedicated hardware receive queues for virtual machines that filter receive packets based on the destination MAC address or VLAN. Receive buffers can be allocated for each queue from VM memory. This option improves network throughput by distributing processing of network traffic for multiple VMs among multiple processors. It reduces CPU utilization by offloading receive packet filtering to NIC hardware. VMQs prove beneficial when four or more VMs are in use.
`Virtual Machine Queues Transmit`	`Enabled` (default) `Disabled`	Windows Server 2012 R2 Windows Server 2016	If this option is enabled with VMQs, separate transmit queues are created for each VM network interface. Send and receive interrupts for a VM network interface are processed on the same CPUs. Separate transmit queues increase system overall CPU utilization but offer greater system scalability.
`VLAN Identifier (802.1q)`	`Not Present` (default) `1-4094`	Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	If selected, the adapter adds a VLAN tag to all transmitted packets, and only receives packets with the matching VLAN tag. **NOTE** Do not use this property if the Emulex Teaming Driver is enabled. In that case, perform VLAN configuration in the Emulex Teaming Driver application. **NOTE** Do not use this property with Hyper-V. Use the Microsoft Hyper-V Manager to configure VLANs on each virtual machine.

3.2.1.1 Advisory: PowerShell Behavior

3.2.1.1.0.1 Issues with Capabilities Reported by Standard PowerShell Commands (Get-NetAdapter)

The default registry values for driver parameters are initially populated from the driver installation `.inf` file. Thereafter, the registry is written to only if the default settings are explicitly overridden. PowerShell uses these registry values to report capabilities with the result that the registry values might not always reflect what is supported in the current configuration.

The default settings can be modified through the **Driver Properties** page, standard PowerShell commands, and utilities, such as elxoccfg (for additional information on elxoccfg, see Section 3.2.3, Using ELXOCCFG for Windows NIC Driver Options).

Standard PowerShell (Get-NetAdapter) commands function in the following manner:

- If the feature is currently enabled, the driver reports its current capabilities. PowerShell reports all of the feature capabilities based on what the driver indicates. The feature capabilities are guaranteed to be what the NIC driver supports in the current configuration.
- If the feature is not enabled, the driver does not report any current capabilities. At that point, PowerShell searches the registry for keys related to the feature and reports their values. These values are either the default values

(`.inf`) or the last configured user values (if overwritten by the user). Default values are only intended as maximum upper bounds; they are not guaranteed resources supported in every configuration.

As a result, the driver can only report a feature's current capabilities (accurate for the present configuration) if the feature is currently enabled. However, standard PowerShell commands report whatever is present in the registry if the feature is not enabled. This information can conflict with what the driver actually supports in the current configuration.

3.2.1.1.0.2 Determining What PowerShell Is Reporting (Registry and Driver-Reported Capabilities)

You can usually tell whether PowerShell is using capabilities reported by the driver or is picking up registry values.

For SR-IOV, check the output of `Get-NetAdapterSRIOV` and `CurrentCapabilities`.

If `CurrentCapabilities` is empty, the driver is not currently enabled for SR-IOV. Any reported fields in `Get-NetAdapterSriov | fl *` are based on registry values. If `CurrentCapabilities` is not null, the driver is enabled for SR-IOV. `Get-NetAdapterSriov` fields are based on what the driver reports.

3.2.2 Configuring Windows Server NIC Driver Parameters

The Windows Server NIC drivers support driver options through the **Advanced Property** page in the **Windows Device Manager**.

NOTE Ensure that the OneCommand Manager application GUI is closed before opening the Windows Device Manager.

3.2.2.1 Modifying Advanced Properties

Modify the advanced properties for the driver for Windows with the Windows Device Manager. For additional information on advanced properties, see Section 3.2.5, Network Driver Performance Tuning.

To modify the advanced properties, perform these steps:

1. Enter the Windows Device Manager using one of the following options:
   - Click **Start > Control Panel > System** and click the **Device Manager** hyperlink.
   - Click **Start > Run**, and type
     `devmgmt.msc`
2. Click **OK**.
   The **Windows Device Manager** is displayed (Figure 2).

Figure 2 Partial View of Windows Device Manager

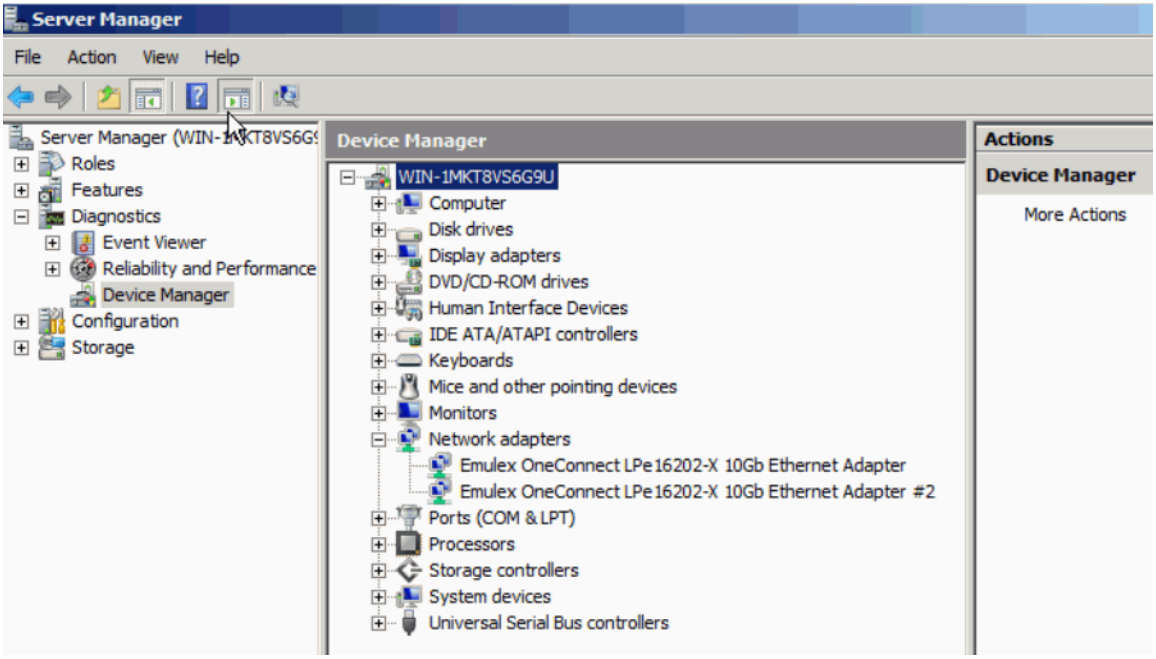

3. Right-click the network adapter for which you want to modify advanced properties.
4. Click **Properties**, and click the **Advanced** tab (Figure 3).
5. From the list of properties, click the property (parameter) you want to modify, then select the new value of the property by selecting it from the list under the property.
6. Click **OK**.

NOTE Modifying properties causes the network driver to reload, and some TCP connections might be temporarily dropped.

Figure 3 NIC Advanced Properties

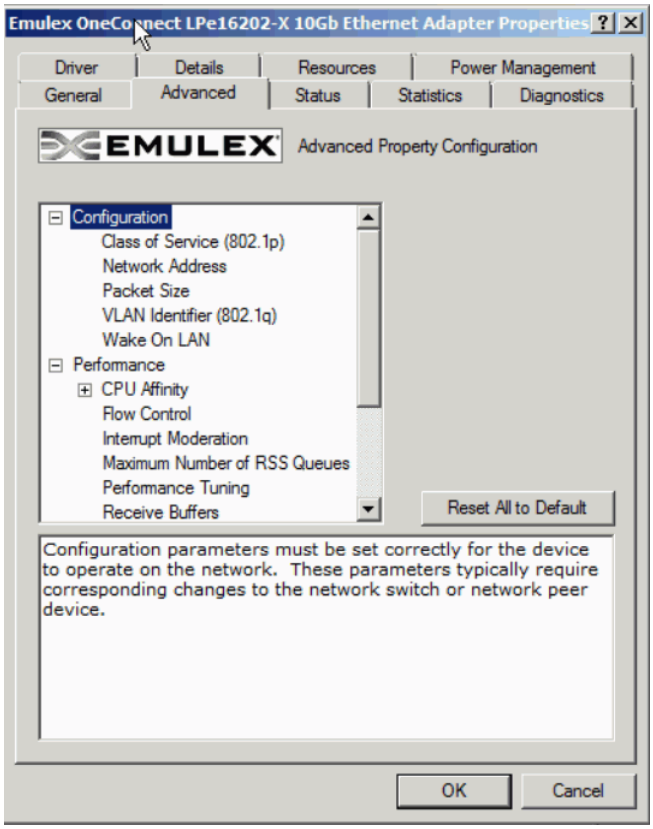

3.2.2.2 Statistics Property Page

Use the **Statistics** tab to view the performance of the device and network. By viewing the statistics properties, you can troubleshoot issues and performance-tune the system; for example, you can assess how different device properties change the system performance.

To view the statistics properties, perform these steps:

1. Enter the Windows Device Manager using one of the following options:
   - Click **Start > Control Panel > System** and click the Device Manager hyperlink.
   - Click **Start > Run**, then type:

     `devmgmt.msc`
2. Click **OK**.

   The **Windows Device Manager** is displayed (Figure 2).
3. Right-click the network adapter for which you want to view the statistics properties.
4. Click **Properties,** then click the **Statistics** tab (Figure 4).

Figure 4 NIC Statistics Properties in Windows Server 2016

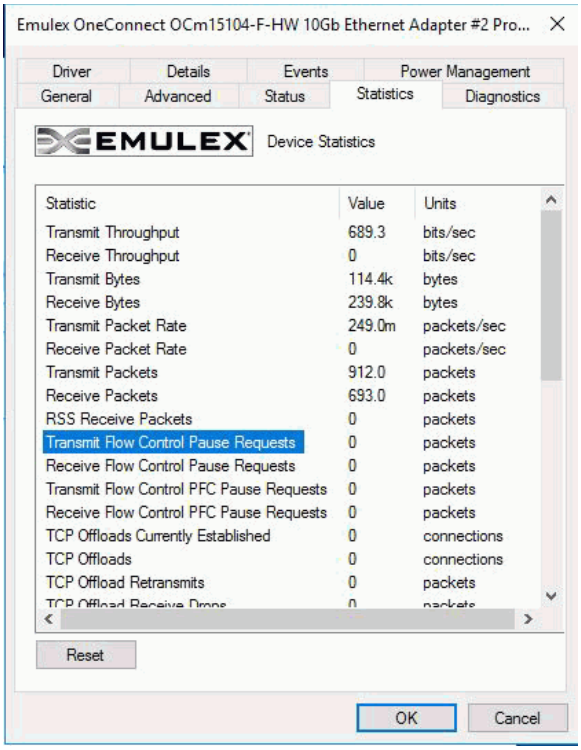

5. From the list of properties, select the property (parameter) you want to view.

Table 4 lists the NIC driver properties statistics.

Table 4 NIC Driver Properties Statistics

Statistic Name	Description
Transmit Throughput	The data rate for this adapter on the network, including all packet headers. It is expressed in terms of bits per second, where 1 byte = 8 bits. This range is computed as the average over approximately 3 seconds.
Receive Throughput	The receive rate for this adapter.
Transmit Bytes	The total number of bytes transmitted by this adapter, since the last statistics reset or the last driver reload.
Receive Bytes	The total number of bytes received by this adapter.
Transmit Packet Rate	The rate of transmit packets for the adapter.
Receive Packet Rate	The rate of receive packets for the adapter.
Transmit Packets	The total number of packets transmitted by the adapter since the last statistics reset, or the driver was reloaded.
Receive Packets	The total number of packets received. This number includes both RSS and non-RSS packets.
RSS Receive Packets	The number of receive packets that were suitable for RSS.
Transmit Flow Control Pause Requests	The number of times the network adapter sent a PAUSE frame to request that the peer stop sending data temporarily. This number indicates a potential bottleneck in the system. Typically, this bottleneck is the result of the DMA of packets from the adapter to host memory.
Receive Flow Control Pause Requests	The number of times the network adapter received a PAUSE frame from the peer. This number indicates a potential bottleneck in the attached switch or network peer device. This statistic increments only if the switch is correctly configured for flow control.

Table 4 NIC Driver Properties Statistics (Continued)

Statistic Name	Description
Transmit Flow Control PFC Pause Requests	IEEE 802.1Qbb PFC extends an Ethernet PAUSE frame to each of the eight traffic classes. The PFC PAUSE frame is a link-level mechanism used by overwhelmed receiver nodes to halt the transmission from the peer node for a specified period of time.
Receive Flow Control PFC Pause Requests	IEEE 802.1Qbb PFC extends an Ethernet PAUSE frame to each of the eight traffic classes. The PFC PAUSE frame is a link-level mechanism used by overwhelmed receiver nodes to halt the transmission from the peer node for a specified period of time.
VMQs Allocated	The current number of virtual machine queues allocated.
Interrupt Rate	The number of interrupts per second generated by the adapter. The interrupt rate can be tuned by modifying the **Interrupt Moderation** parameter.
Receive Drops No Memory (DMA Limited)	The number of packets dropped as a result of insufficient buffers posted by the driver. This value is generally the result of the CPU core used for any receive queue reaching 100%. The system might lack sufficient CPU cycles to post receive buffers at the necessary rate. Many small packets lead to this behavior on almost any CPU, because the processing time for small packets is very high in the networking stack. Using a teaming driver might also lead to this behavior, because it increases the CPU load during receive. Increasing the number of **Receive Buffers** in the **Advanced Property** page might alleviate some of these drops, in particular if the drops are the result of bursts of small receive packets on the network. However, if the CPU is the limit, increasing the buffer resources does not help because the driver cannot post them fast enough. Enabling RSS is another strategy to reduce drops because it allows the NIC driver to use additional CPU cores. The number of RSS queues can be increased to increase the total number of posted buffers available to the adapter. Enabling RSC can also reduce CPU consumption in the networking stack by combining multiple TCP packets into one larger packet. For best performance, the system BIOS must be set to **Maximum Performance** or manually disable C-states. The transitions to low power C-states might cause a steady trickle of drops due to increased latencies from packet reception until the driver's interrupt processing code is invoked.
Receive Drops No Fragments (CPU Limited)	The number of receive packets dropped because of a DMA bottleneck from the network adapter to host memory. This situation might be caused by bottlenecks in either the PCIe bus or main memory. In the **Status** tab of the **Custom property** page, the Emulex NIC reports the PCIe link parameters and the maximum supported parameters. For example, installing a 8x device in a 4x PCIe slot cuts the available PCIe bandwidth in half. The PCIe MTU and Read Request size are also reported, and these can be configured in the system BIOS. The performance of the main memory is the other major concern for networking throughput. The ideal situation uses high-speed memory with all memory channels populated per CPU; typically, three or four DIMMs per CPU socket. For the ideal performance, use the same DIMM size in each memory channel to allow perfect memory channel interleaving. Features, such as memory sparing or memory mirroring, dramatically decrease the memory bandwidth of the system and cause drops. TCP connection offload might lead to increased drops as a result of “no memory.” If TCP connection offload is used, enabling flow control might reduce the drops. Alternatively, disabling TCP connection offload might improve performance.
Receive CRC Errors	The number of packets dropped as the result of CRC errors on the layer 2 Ethernet packet. In products that expose multiple PCIe functions per Ethernet port, this statistic is incremented only for the lowest PCI function per port because the packet cannot be further classified because of the error.
Receive IP Checksum Errors	The number of receive packets with an incorrect IPv4 checksum. These packets are provided to the TCP/IP stack for disposal in the operating system.
Receive UDP Checksum Errors	The number of receive packets with an incorrect UDP checksum. These packets are provided to the TCP/IP stack for disposal in the operating system.
Receive TCP Errors	The number of receive packets with an incorrect TCP checksum. These packets are provided to the TCP/IP stack for disposal in the operating system.

Table 4 NIC Driver Properties Statistics (Continued)

Statistic Name	Description
Tunnels allocated	The number of interfaces converted to tunnel interfaces. Used with NVGRE offload enabled and on.
Tenants allocated	The number of interfaces converted into tenant interfaces. Used with NVGRE offload enabled and on and VMQ.
Virtual Functions allocated	The number of PCIe virtual functions created by the SR-IOV supporting adapter.

3.2.3 Using ELXOCCFG for Windows NIC Driver Options

The `elxoccfg.exe` program supports configuring parameters for the network functions on Emulex Ethernet adapters either through Interactive mode with a set of menus, or Command line mode that is scriptable.

If you performed a standard driver installation, the `elxoccfg.exe` file is located in the following directory:

```
Directory of C:\Program Files\Emulex\AutoPilot
Installer\NIC\Drivers\NDIS\<platform>\<OS>
```

The following section describes how to use the `elxoccfg.exe` program to configure the Windows device driver from the command line.

3.2.3.1 Using elxoccfg.exe Options

To display help, use the `-?` option by typing:

```
elxoccfg.exe -?
```

The following text is displayed:

```
OneConnect Network Config (0.0.9999.0)
Copyright 2011 Emulex
Usage:   elxoccfg.exe [-options]
```

Running the command with no arguments will display a menu to select the adapter and parameters to modify. Use the command line arguments to script this process.

Table 5 describes the available options.

Table 5 elxoccfg.exe Options

Option	Description
`-a str[,str]`	Selects all adapters with any of the given strings in the connection or device name. If omitted, `occfg` prompts for an adapter from a list.
`-s name=v,[name=v]`	Sets the parameter's value and reloads the devices.
`-g name[,name]`	Gets parameter value.
`-r`	Skips reloading the driver when setting a parameter.
`-f`	Forces reloading the driver.
`--`	Forces disabling the driver.
`-+`	Forces enabling the driver.
`-l`	Lists available adapters and exits.
`-T filename`	Saves the tinylog to a binary file.
`-L filename`	Loads a binary file and replays the tinylog.
`-x`	Resets all parameters to the default values.
`-p`	Shows all registry parameter values.

Table 5 elxoccfg.exe Options (Continued)

Option	Description
-q	Shows all driver parameter values.
-h	Shows help text for all parameters.
-?	Shows this help.
-M module=*trace level* [,module=*trace level*]	Continously downloads the ARM logon to a file. Arguments set a specific trace level on listed modules. Default argument is all=error. Refer to the ARM firmware for list of modules and debug trace levels. This is a special command argument.

3.2.3.1.1 Examples:

```
Run interactively with menus:                elxoccfg.exe
Set a parameter on all Emulex adapters:      elxoccfg.exe -a Emulex -s
                                             rss=1
Set multiple parameters on one adapter adapter:
       elxoccfg.exe -a "Local Area Connection 23" -s "Flow=3,rss=0"
```

3.2.3.2 Selecting an Adapter

In batch mode, the -a parameter must be followed by a substring that is contained within the adapter name. The name is a combination of the device manager name (for example, Emulex LightPulse LPe16202) and the network connection name (for example, Local Area Connection). The latter can be modified by using the Windows Network Connections applet (ncpa.cpl).

The most typical scenario involves setting parameters to be the same for all ports of a network adapter by specifying -a emulex.

It is often convenient to rename the connections to have a common name to easily operate on a group. For example, naming the network connections "dot1, dot2, dot3" allows operating on all adapters using the substring "dot", or on any individual adapter by specifying the exact name such as, "dot1".

3.2.3.3 Configuring Device Parameters

The elxoccfg program queries and modifies registry parameters for Emulex network devices. The registry keys are stored at:

HKLM/System/CurrentControlSet/Control/Class/{4D36E972-E325-11CE-BFC108002bE10318}/####

where #### is the device instance number.

The elxoccfg program allows you to modify registry keys on a set of network devices. After the driver is modified, it must be restarted to apply these parameters. In Batch mode, elxoccfg automatically restarts the driver when changing a parameter, and, in Interactive mode, you use a menu item to select to restart the driver.

In Batch mode, the commands to modify parameters look like the following examples:

```
elxoccfg -a emulex -s rss=0
elxoccfg -a emulex -s "Interrupt Moderation=4,Flow Control=3"
```

The parameter name must uniquely specify one parameter to modify, but it might be only a substring on the full parameter name. For example, the following examples are all equivalent:

```
elxoccfg -a emulex -s "Flow Control=3"
elxoccfg -a emulex -s flow=3
elxoccfg -a emulex -s control=3
```

Note that the parameter name is generally the text readable parameter description name, but you can specify the exact registry key name as well. Microsoft has defined many documented standard registry key names that start with a * character. The * is not a wildcard, it is part of the registry key name. The following examples are equivalent:

```
elxoccfg -a emulex -s "Flow Control=3"
elxoccfg -a emulex -s "*FlowControl=3"
```

NOTE Quotation marks are required if the parameter name contains a space character.

To modify a parameter without a driver reload, use `-r`. This setting is used to modify several parameters in sequence, without forcing a driver reload. To force a driver reload, use the `-f` parameter.

The following is an example of such a sequence:

```
elxoccfg -a emulex -r -s rss=0
elxoccfg -a emulex -r -s "interrupt moderation=0"
elxoccfg -a emulex -f
```

Registry keys can be set to two special values:

- The `delete` value causes the key to be entirely deleted and the driver uses the default value. This value is appropriate for keys that are optional, such as the "`Network Address`".
- The `default` value sets the key to the driver's default value. If the key is optional, the default value might be equivalent to deleting the key.

For example:

```
elxoccfg -a emulex -s vlan=delete
elxoccfg -a emulex -s rss=default
```

3.2.3.4 Viewing Device Parameters

The `elxoccfg.exe` program can query device parameters from either the registry or the device driver (if running driver version greater than or equal to 2.103.*x.x*).

The registry and driver values might differ until the driver is reloaded. If the driver reload fails for any reason (such as another application has an open handle to the device driver), it might be necessary to reboot the system to apply the registry changes.

NOTE If the driver has been disabled or if the driver failed to load due to any error, the driver query returns the error, `Failed to query driver for the parameter`.

The following are Batch mode examples:

```
elxoccfg -a emulex -g "Interrupt Moderation"
elxoccfg -a "(Local Area Connection)" -g interrupt,rss
Emulex LightPulse LPe16202, NIC (Local Area Connection):
   [Registry] Interrupt Moderation = 4 (Adaptive [Default])
   [Driver] Interrupt Moderation = 4 (Adaptive [Default])
Emulex LightPulse LPe16202, NIC (Local Area Connection):
   [Registry] RSS = 0 (Disable)
   [Driver] RSS = 0 (Disable)
```

3.2.3.5 Resetting All Parameters

Resetting all parameters restores the default values for each adapter. To reset all parameters, use the following command:

```
elxoccfg -a emulex -x
```

3.2.3.6 Displaying All Parameters

To display the current value of all parameters, use either `-p` or `-q` command line options. These options display the registry value or driver value of the parameter, or both when using `-pq` together.

For example:

```
elxoccfg.exe -a "SLOT 4 2 Port 1" -pq
OneConnect Network Config (11.2.107.0)
Copyright 2009-2016 Broadcom. All rights reserved.

Emulex OneConnect LPe16202-X 10Gb Ethernet Adapter #4 (SLOT 4 2 Port 1)
   Display all properties.

   [Registry] Class of Service (802.1p) = 1 (Auto Priority Pause)
   [Driver]   Class of Service (802.1p) = 1 (Auto Priority Pause)

   [Registry] Enhanced Transmission Selection = 0 (Disabled)
   [Driver]   Enhanced Transmission Selection = 0 (Disabled)

   [Registry] Flow Control = 3 (Rx & Tx Enabled)
   [Driver]   Flow Control = 3 (Rx & Tx Enabled)

   [Registry] IPv4 Checksum Offload = 3 (Rx & Tx Enabled)
   [Driver]   IPv4 Checksum Offload = 3 (Rx & Tx Enabled)

   [Registry] Interrupt Moderation = 4 (Adaptive 30k Int/sec (default))
   [Driver]   Interrupt Moderation = 4 (Adaptive 30k Int/sec (default))

   [Registry] Large Send Offload V1 (IPv4) = 1 (Enabled)
   [Driver]   Large Send Offload V1 (IPv4) = 1 (Enabled)

   [Registry] Large Send Offload V2 (IPv4) = 1 (Enabled)
   [Driver]   Large Send Offload V2 (IPv4) = 1 (Enabled)

   [Registry] Large Send Offload V2 (IPv6) = 1 (Enabled)
   [Driver]   Large Send Offload V2 (IPv6) = 1 (Enabled)

   [Registry] Maximum Number of RSS Processors = 12
   [Driver]   Maximum Number of RSS Processors = 16 (0x10)

   [Registry] Maximum Number of RSS Queues = 8
   [Driver]   Maximum Number of RSS Queues = 8

   [Registry] Maximum RSS Processor Number = <not set>
   [Driver]   Maximum RSS Processor Number = <not set>

   [Registry] Network Address = <not set>
   [Driver]   Network Address = <not set>

   [Registry] Packet Size = 1514 (1514)
   [Driver]   Packet Size = 1514 (0x5ea) (1514)

   [Registry] Performance Tuning = 0 (Maximum Performance)
   [Driver]   Performance Tuning = 0 (Maximum Performance)
```

```
[Registry] Preferred NUMA Node = <not set>
[Driver]   Preferred NUMA Node = <not set>

[Registry] RSS Base Processor Group = <not set>
[Driver]   RSS Base Processor Group = <not set>

[Registry] RSS Base Processor Number = <not set>
[Driver]   RSS Base Processor Number = <not set>

[Registry] RSS Max Processor Group = <not set>
[Driver]   RSS Max Processor Group = <not set>

[Registry] RSS Profile = 1 (Closest Processor)
[Driver]   RSS Profile = 1 (Closest Processor)

[Registry] Receive Buffers = 896
[Driver]   Receive Buffers = 1664 (0x680)

[Registry] Receive CPU = <not set>
[Driver]   Receive CPU = <not set>

[Registry] Receive Side Scaling = 1 (Enabled)
[Driver]   Receive Side Scaling = 1 (Enabled)

[Registry] Recv Segment Coalescing (IPv4) = 1 (Enabled)
[Driver]   Recv Segment Coalescing (IPv4) = 1 (Enabled)

[Registry] Recv Segment Coalescing (IPv6) = 1 (Enabled)
[Driver]   Recv Segment Coalescing (IPv6) = 1 (Enabled)

[Registry] SR-IOV = 1 (Enabled)
[Driver]   SR-IOV = 1 (Enabled)

[Registry] TCP Checksum Offload (IPv4) = 3 (Rx & Tx Enabled)
[Driver]   TCP Checksum Offload (IPv4) = 3 (Rx & Tx Enabled)

[Registry] TCP Checksum Offload (IPv6) = 3 (Rx & Tx Enabled)
[Driver]   TCP Checksum Offload (IPv6) = 3 (Rx & Tx Enabled)

[Registry] Transmit = 1 (Enabled)
[Driver]   Transmit = 1 (Enabled)

[Registry] Transmit Buffers = 2048 (2048)
[Driver]   Transmit Buffers = 2048 (0x800) (2048)

[Registry] Transmit CPU = <not set>
[Driver]   Transmit CPU = <not set>

[Registry] UDP Checksum Offload (IPv4) = 3 (Rx & Tx Enabled)
[Driver]   UDP Checksum Offload (IPv4) = 3 (Rx & Tx Enabled)

[Registry] UDP Checksum Offload (IPv6) = 3 (Rx & Tx Enabled)
[Driver]   UDP Checksum Offload (IPv6) = 3 (Rx & Tx Enabled)
```

```
[Registry] VLAN Identifier (802.1q) = 0
[Driver]   VLAN Identifier (802.1q) = 0

[Registry] Virtual Machine Queues = 1 (Enabled)
[Driver]   Virtual Machine Queues = 1 (Enabled)
```

3.2.3.7 Using Interactive Mode

The `elxoccfg.exe` program also supports Interactive mode with a set of menus.

To start this utility in Interactive mode, perform these steps:

1. Run `elxoccfg.exe` from a command console.

   A list of adapters is displayed on which to operate.
2. Type either a number of the list or a substring from any part of the name (for additional information, see Section 3.2.3.2, Selecting an Adapter).

   The program prompts for an operation, such as modifying or querying a parameter value.
3. Follow the prompt.

   The program provides a list of available registry parameters to modify or query.
4. Type either the number of the corresponding option or a substring in the parameter name. The substring must uniquely identify the parameter or typing `occfg` displays all potential options.
5. To apply the parameters, select the menu item to exit and reload the drivers. Pressing **Ctrl+ C** at any point might leave modifications in the registry, but the driver does not use the new parameters until it is reloaded.

3.2.3.8 Parameter Help

In interactive mode, setting a parameter displays help text and information regarding the valid values for each parameter. This information can be dumped for all parameters by specifying the `-h` option.

The following is an example help text for the RSS parameter.

```
RSS:
Receive Side Scaling (RSS) scales receive processing over multiple CPUs in
parallel. This scaling typically improves application performance; however, it
tends to increase CPU usage on low end machines.
RSS is only supported on two primary adapters per device. It will appear disabled
for additional PCI functions in blade server configurations.
RSS requires Windows Server 2012 and later.
Registry Key: *RSS
Default Value : 1 (Enable)
Valid Values :
   0 = Disable
   1 = Enable
```

3.2.4 Using SR-IOV with Emulex Devices

This section describes how to use SR-IOV with Emulex devices.

NOTES

- The operating system comes with an Emulex inbox driver. Use the Emulex out-of-box driver.
- For a list of supported drivers and adapters, refer to the latest Windows Drivers release notes, which are available for download

from the Documents and Downloads area of the Broadcom website.

- LPe16202 and OCe15100 adapters in NIC+FCoE mode support 63 virtual functions per physical function.

3.2.4.1 Server BIOS Configuration

SR-IOV requires support in the server chipset beyond standard virtualization technologies, including operating system control of PCIe and interrupt remapping. The server might have BIOS options to control SR-IOV, and typically these are disabled by default. The following items might require modification in your system BIOS during boot:

- Enable "Virtualization", such as Intel VT-x or AMD-V. This setting is required for any virtual machine.
- Explicitly enable SR-IOV in the system BIOS. The specific name for this option varies among vendors. For instance, it might be called `Intel VT-d` (Virtualization Technology for Direct I/O), `AMD-Vi` (AMD I/O Virtualization Technology), or `IOMMU`.

3.2.4.2 SR-IOV Server Validation

Use the following Microsoft PowerShell commands to determine if your server is capable of SR-IOV:

- `Get-NetAdapterSriov`
- `Get-VmHost`
- `Get-VmNetworkAdapter`
- `Get-VmSwitch`

Refer to the Microsoft documentation for additional information.

NOTE Early SR-IOV-capable chipsets had errors that might prevent SR-IOV from operating in Windows Server 2012 and Windows Server 2012 R2. The PowerShell command `Get-VmHost | fl *` includes `IovSupportReasons` that indicates if the chipset has this issue.

3.2.4.2.1 Enabling SR-IOV on Unqualified Servers

If Windows Server 2012 or Windows Server 2012 R2 detects a problem with the system I/O remapping hardware, you might still be able to use SR-IOV by explicitly enabling SR-IOV in the registry using `IovEnableOverride`.

NOTE Only use this procedure for trusted virtual machines.

CAUTION Use the registry editor at your own risk. Using the registry editor can cause serious issues that might require you to reinstall the computer's operating system. Broadcom cannot guarantee that issues resulting from changes you make to the registry can be repaired. Make a backup of your registry before making any changes.

3.2.4.2.2 Backing Up and Editing the Registry

To back up and edit the registry, perform these steps:

1. Create a system restore point.
2. Open the registry editor by running `regedit.exe` at the command prompt.
3. Select the hive (the top-level key) and export it to a `.reg` file.
4. Save the `.reg` file to a location off of the server as a precaution.
5. Navigate to:

   `HKLM\Software\Microsoft\Windows NT\CurrentVersion\Virtualization`
6. Create a Dword-type entry named `IovEnableOverride`.

7. Set the value of `IovEnableOverride` to `1`.
8. Reboot the system.

   If the system does not boot, press **F8** and select **Previous Known Good**, or use the system restore function while booting from an operating system installation disc or recovery disk.
9. If the system boots but does not work properly, restore from a previous restore point, or import the saved `.reg` file and reboot.

3.2.4.3 Verifying the Driver Version

To verify that the device driver meets the minimum requirements, perform these steps:

1. Select **Server Manager > Dashboard > Tools > Computer Management**.
2. Click **Device Manager**.

   The **Device Manager** opens (Figure 5).

Figure 5 Device Manager for Windows Server 2012

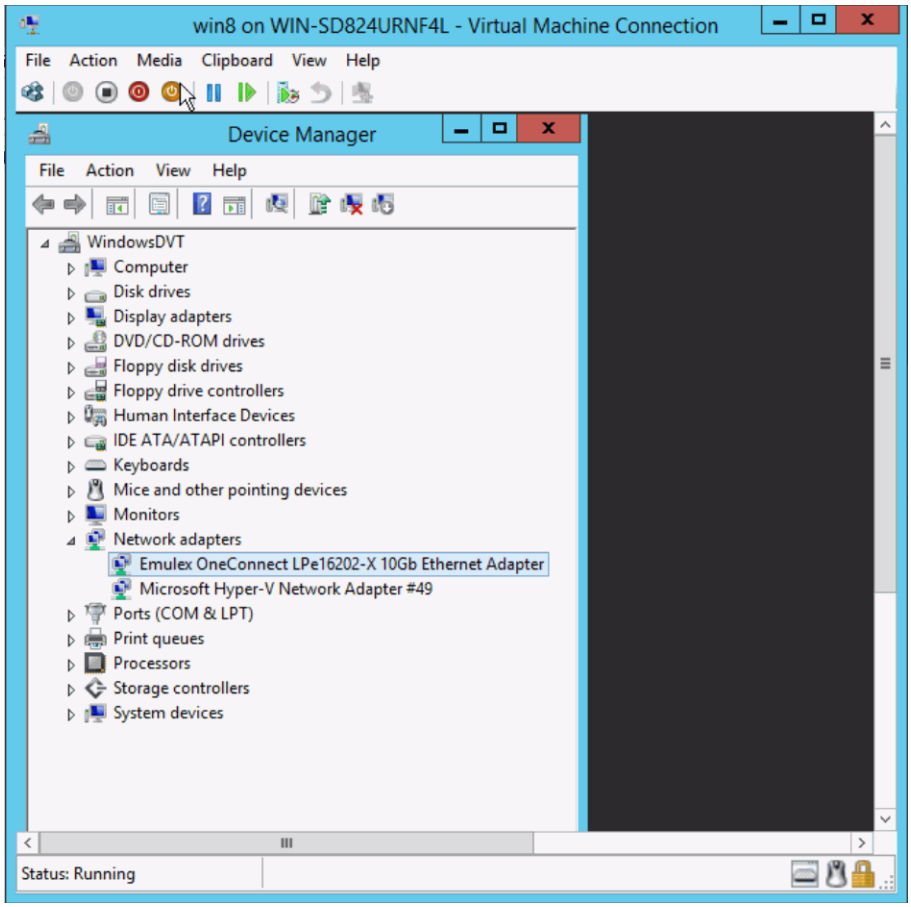

3. Open the **Network adapters** item, find the Emulex device and right-click it.
4. Select **Properties** from the context menu.

   The **Properties** dialog opens and shows the **Driver** page (Figure 6). The **Driver** page contains the driver version number.

Figure 6 Emulex NIC Driver Page

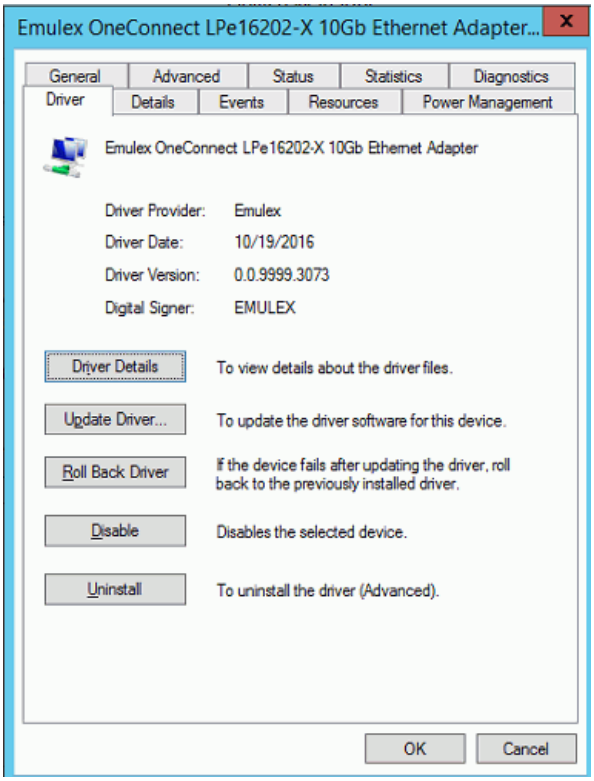

5. Click **Driver Details**.
   A window opens that displays the driver name.

3.2.4.4 Enabling SR-IOV in the Emulex Device

To enable SR-IOV in the Emulex device, perform these steps:

1. Select **Server Manager > Dashboard > Tools > Computer Management**.
2. Click **Device Manager**.
   The **Device Manager** opens (Figure 5).
3. Open the **Network adapters** item, find the Emulex device and right-click it.
4. Select **Properties** from the context menu.
   The **Properties** dialog opens (Figure 6).
5. Click the **Advanced** tab.
   The **Advanced Property Configuration** page opens (Figure 7).

Figure 7 Emulex NIC Advanced Configuration Page

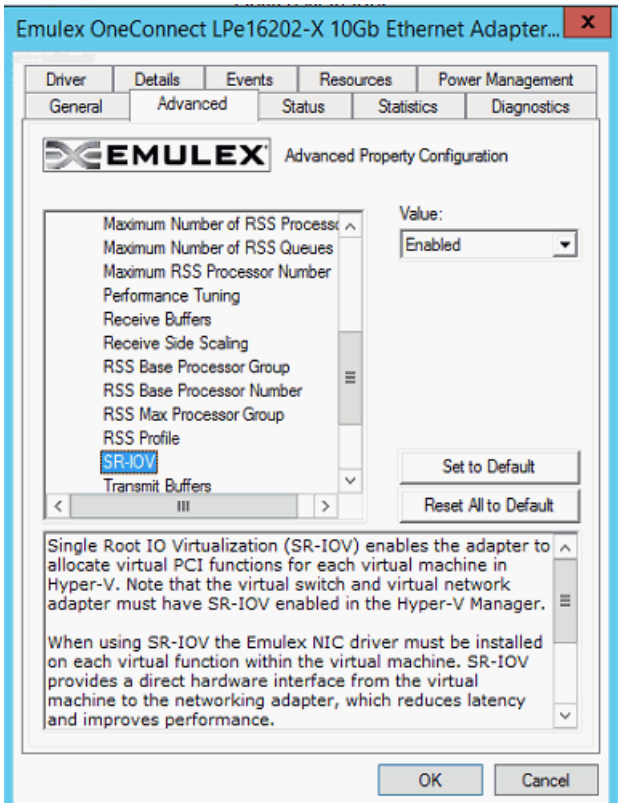

6. Select **SR-IOV** from the list and select **Enabled** from the **Value** drop-down list.

NOTE You must configure Hyper-V to create an SR-IOV-enabled virtual machine. Refer to the Microsoft Hyper-V documentation for additional information.

3.2.4.5 SR-IOV Link Behavior

Link status information for the virtual function (VF) is not available from the Network Connection Manager of the virtual machine. The link status of the VF adapter is displayed in the **Status** tab of the **Device Manager Advanced Property** page for the corresponding adapter.

If the link status of the VF is up, then it is shown as Connected along with the Link Speed. When PF link is disconnected, the VF adapter will be shown as Disconnected.

Figure 8 PF Link Disconnected, VF Adapter Disconnected

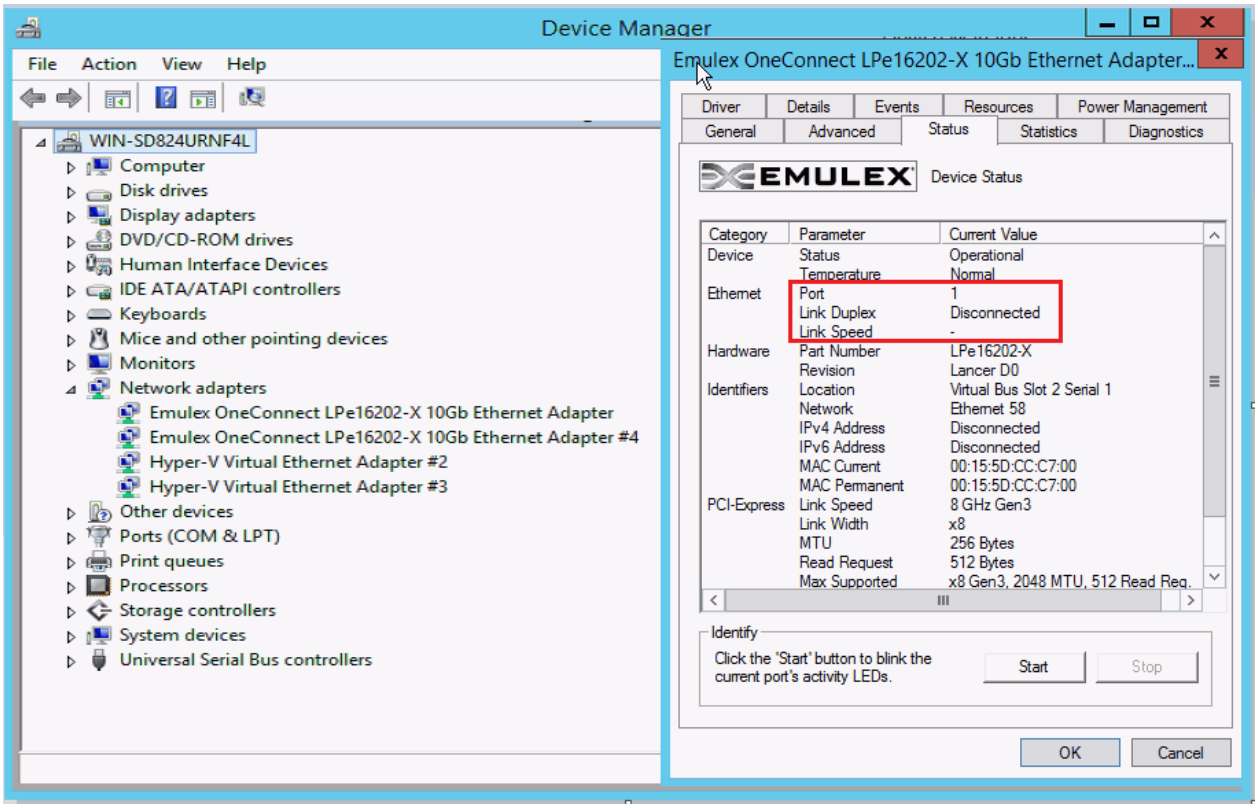

When a VF adapter link is down, the communication between the VF and its corresponding physical function (PF) adapter is still possible. The default synthetic data path provided by Hypervisor will be used instead of the VF data path of the adapter. In this case, statistics counters of the VF adapter are not incremented.

The virtual switch connected to the PF is always shown as link up regardless of the PF link status. The link status of the switch is controlled by the operating system and not the PF driver. Figure 9 depicts the **Network Connections** panel when the PF link is disconnected, but the SR-IOV enabled virtual switch link is shown as connected.

Figure 9 Network Connections Panel – PF Link Disconnected

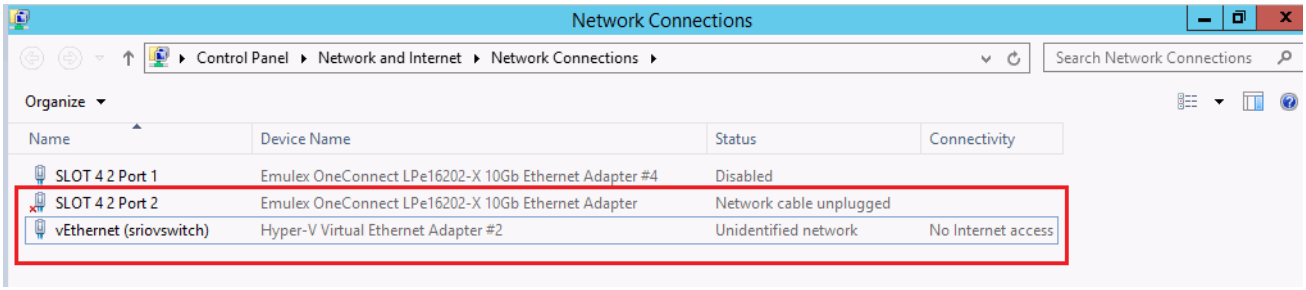

3.2.4.6 Hyper-V

The Hyper-V role must be added using the Server Manager. After the Hyper-V role is added, enable SR-IOV in the Hyper-V Manager by doing one of the following:

- Creating the virtual switch
- Creating each virtual NIC

Refer to the Microsoft documentation for additional information.

NOTE Ensure that SR-IOV is enabled on the server and on the Emulex adapter prior to configuring the Hyper-V virtual switch.

The Windows Server 2012, Windows Server 2012 R2, and Windows Server 2016 servers treat SR-IOV as an offload. Each VM using SR-IOV gets its own VF with access to a subset of HW resources on the NIC, allowing for superior performance by bypassing the Hyper-V vSwitch reducing overhead.

After the Emulex driver is loaded, the Emulex SR-IOV virtual function is used for all unicast receive and transmit traffic. Live migration and multicast are supported while using SR-IOV. If SR-IOV is disabled, the Emulex adapter is removed from the virtual machine, and all traffic automatically uses the Hyper-V synthetic NIC.

NOTE If multiple adapters are added to the virtual machine, use MAC addresses to map the Emulex network adapter to the corresponding Microsoft virtual network adapters.

3.2.4.7 Verifying SR-IOV

If SR-IOV is enabled, it can be verified by opening the Device Manager within the virtual machine and examining the information about the transmit and receive packets that are using the SR-IOV virtual function. This final verification shows that SR-IOV is working correctly. SR-IOV also can be verified from the host Hyper-V server.

NOTE Because current versions of Windows Server 2012 require that SR-IOV be enabled in different locations prior to creating the virtual switch, if SR-IOV is not working, delete the virtual switch and create it again. The SR-IOV option is always available during switch creation.

3.2.4.7.1 Verifying SR-IOV from the Virtual Machine

To verify SR-IOV from within the virtual machine, perform these steps:

1. From within the virtual machine, select **Server Manager > Dashboard > Tools > Computer Management**.
2. Click **Device Manager**.

   The **Device Manager** opens (Figure 5).
3. Open the **Network adapters** item, select the Emulex device and right-click it.
4. Select **Properties** from the context-menu.

   The **Properties** dialog opens showing the **Driver** page (Figure 6).

   NOTE The Emulex adapter might initially appear as a “Network Adapter” before the driver is loaded.
5. Select the **Statistics** tab (Figure 10).

   Information about the transmit and receive packets that are using the SR-IOV virtual function are displayed; specifically, the number of Transmit Bytes and Receive Bytes that are transmitted directly to hardware from the virtual function. If this number is greater than zero, the device is successfully using the SR-IOV direct hardware access.

Figure 10 Emulex NIC Statistics Page

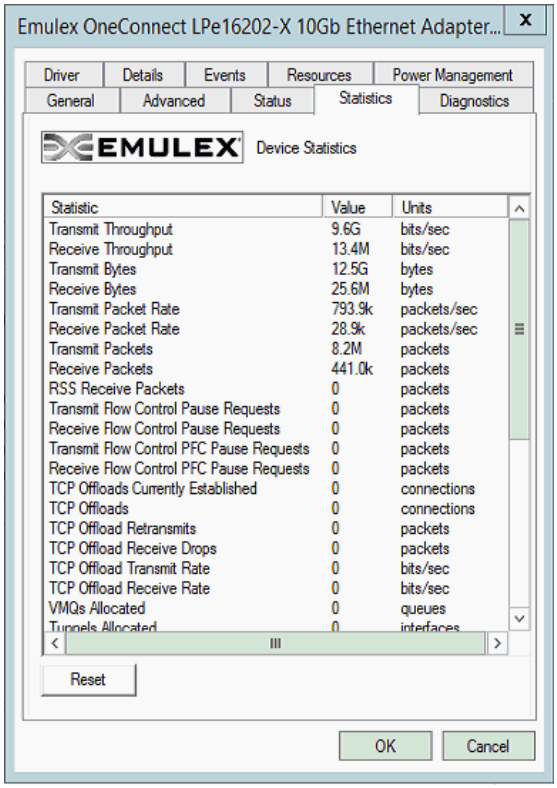

3.2.4.7.2 Verifying SR-IOV from the Host Hyper-V Server

To verify SR-IOV from the host Hyper-V server, perform these steps:

1. From the **Device Manager**, open the **Network adapters** item, select the Microsoft Hyper-V Network adapter and right-click.
2. Select **Properties** from the context-menu.

   The **Hyper-V Network adapter Properties** dialog opens and shows the **Driver** page.
3. Select the **Statistics** tab (Figure 10).
4. From the **Statistics** tab, locate the Virtual Functions Allocated item.

   "Virtual Functions Allocated" shows the count of currently enabled virtual functions.

NOTE The Microsoft Powershell command **Get-NetAdapterSriovVf** lists each SR-IOV virtual function. Refer to the Microsoft documentation for additional information.

3.2.5 Network Driver Performance Tuning

This section describes the tuning and configuration of the network drivers.

3.2.5.1 Optimizing Server Hardware and BIOS Configuration

Adapter performance can be improved by selecting a more efficient PCIe packet payload size. If the system BIOS allows selection of a larger PCIe packet size, selecting at least a 512-byte PCIe packet payload size provides the best efficiency for PCIe data transfers. This setting might be an option in the server's system BIOS. The current value is displayed in the **Device Manager** on the **Status property** page for the adapter.

Most computers offer multiple distinct memory channels, which must be configured for channel interleaving for optimal performance. Optimal interleaving is achieved by using the exact same DIMM configuration for each memory channel. Check the manufacturer's documentation and BIOS parameters for details about optimizing memory bandwidth. Typically, all of the DIMM slots must be populated to make use of all memory channels. As a general rule, more DIMMs provide better performance by allowing a higher degree of memory-access interleaving to occur. However, some servers decrease the memory speed if using more than two DIMMs per memory channel, so it is important to consider the trade-off for a particular server platform.

Some servers might allow memory mirroring or memory sparing, where the total memory is divided in half and each location is stored twice. Memory mirroring and memory sparing provide fault recovery if one memory location detects an error, but they greatly reduce the perceived memory bandwidth of the system.

Nearly any desktop or low-end server has enough memory bandwidth for the adapter to support DMA at 20 Gb/s of data (10 Gb/s read, 10 Gb/s write). However, most of the memory demands come from the processor accessing the data for either packet copies in the non-offloaded networking stack or application. Increasing the clock speed of the memory interface to the processor can be critical for achieving the best networking performance. This interface might be the FSB, Intel QPI, or AMD HyperTransport.

3.2.5.2 Windows Server Network Driver

Table 6 describes ways to use various NIC driver properties and Microsoft Windows properties to performance-tune a system.

Table 6 Windows Server Performance Tuning Situations

Situation	Answer/Solution
A large number of short-lived TCP connections, such as web server or email server, exist.	Enable RSS and increase the number of RSS queues.
Large data transfers, such as to a backup server, exist.	Enable jumbo packets, and use TCP offload.
A small server is struggling to keep up with larger servers on the network.	Disable RSS, enable jumbo packets, and increase the interrupt moderation to allow fewer interrupts per second.
A general-purpose server, such as Active Directory server, DHCP server, or a DNS server, exists.	Enable RSS.

3.2.5.2.0.1 Analyzing Performance Issues

Use the Windows Performance Monitor (perfmon) to view statistics for each network device.

1. Click **Start > Run** and type *perfmon* to launch the Windows Performance Monitor.
2. Right-click and select **Add Counters** to add additional statistics.

Table 7 is a partial list of the statistics to use to troubleshoot performance issues. For network performance, all of the counters from the table are useful: Network Interface, TCPv4, IPv4, and Processor.

Table 7 Statistics and Fine Tuning

Situation	Answer/Solution
Network Interface > Packets Received Errors.	If this value is incrementing even a small amount, a physical problem might exist on the network, such as a loose connection or bad cable, which causes CRC errors in Ethernet packets. Find and eliminate the physical problem.
Network Interface > Packets Received Discarded.	If this value is incrementing dramatically, the computer system might be receiving a lot of unsolicited traffic using network resources.

Table 7 Statistics and Fine Tuning (Continued)

Situation	Answer/Solution
IPv4 > Fragmented Datagrams / sec.	If this value is greater than 0, the computer system is sending or receiving IP fragments. This problem impacts performance. See Section 3.2.5.2.2, Jumbo Packet.
TCPv4 > Segments Retransmitted / sec.	TCP retransmits indicate that packets are being dropped by the receiving system (or in a network switch). Ideally, reduce retransmits to 0.
Processor > % Processor Time.	If CPU usage is high, try to enable all available offloads, such as checksum offloads, and use jumbo packets.

3.2.5.2.0.2 Jumbo Packet

The jumbo packet setting in the registry determines the maximum Ethernet packet size. It includes the Ethernet frame header (typically 14 bytes) but excludes the trailing CRC. The standard packet size is 1514 bytes plus a 4-byte trailing CRC.

Vendors use many terms that refer to this same quantity, such as packet size, frame size, or MTU. The MTU is the Ethernet packet payload size. The MTU does not include the Ethernet frame header or the trailing CRC. The standard MTU is 1500 bytes, which corresponds to a 1514-byte packet size plus a 4-byte trailing CRC. Historically, any 1514-byte frame is a standard packet, while any frame larger than 1514 bytes is called a jumbo packet. Windows Server attempts to standardize the terminology across vendors so that the jumbo packet parameter refers to the byte size of the packet.

The Windows Server driver supports several jumbo packet values. The larger packet size provides better throughput and CPU usage. Typically, all devices on the network, including switches, must be configured for the larger size. The drawbacks of using jumbo packets are interoperability and increased memory usage on the server.

To set a jumbo packet value, go to the **Advanced Properties** page in Windows **Device Manager**. For information on how to configure the options through the **Advanced** page, see Section 3.2.2.1, Modifying Advanced Properties.

The path MTU is the maximum MTU that can be used before IP fragmentation occurs, taking into account the MTU for the endpoints and all routers between the endpoints. To verify the path MTU, send a ping to a remote target with an increasing payload size. Eventually, the IP packet length exceeds the path MTU, and the packet fragments. This situation can be seen by using a packet sniffing application, such as Ethereal, Wireshark, or Microsoft Network Monitor.

IP fragmentation degrades performance dramatically, because all fragments must be received and reassembled before delivering the network packet to the upper layer protocol. In many cases, IP fragmentation can lead to a 10x performance degradation. The MTU parameter must be modified on all systems to avoid IP fragmentation for optimal network throughput.

Typical cases for using the MTU include the following:

- Server interconnects are typically deployed using jumbo frames. This configuration is the most efficient for high bandwidth server-to-server communication, such as Network Attached Storage, iSCSI, and database transactions.
- Servers connected to client systems that run desktop operating systems typically use standard 1500-byte frames. Most desktop systems do not support jumbo packets.
- Servers that require both high performance server-to-server communication and client access can be configured with jumbo frames with Path MTU Discovery enabled. Path MTU Discovery is enabled by default in the Windows Server, and it allows TCP connections to negotiate the optimal packet size that avoids IP fragmentation.

3.2.5.2.1 Flow Control

The adapter supports IEEE 802.3x standard flow control, which uses control packets to temporarily pause the transmission of packets between two endpoints. These control messages are point-to-point; they are not forwarded by switches or routers. You must configure both endpoints for flow control. The adapter can either respond to flow control packets (by temporarily pausing transmits) or send flow control PAUSE packets if the transmitter is

overwhelming the system's receive bandwidth. For best performance, flow control must be enabled on the switches as well as on adapters. Receive and transmit flow control are enabled by default. Flow control is not available if using FCoE on a converged network adapter. In this situation, priority pause is negotiated with the network switch and used only for the FCoE protocol packets.

The NIC function can also use priority pause if it is supported by the switch. This process requires tagging packets in the operating system with the correct priority value, and enabling ETS in the driver properties.

Configurations that support multiple PCI functions per port generally configure flow control from the switch or blade configuration application. Because flow control is an Ethernet port property, it must be the same for all PCI functions using the same port.

If multiple PCI functions are exposed for a single 10GbE port, such as in a blade configuration, the flow control parameter must be set the same on all adapters for the port. The results are unpredictable if the setting differs among PCI functions, because this is a shared property of the 10GbE port.

3.2.5.2.1.1 Examples

Flow control greatly improves the following situations:

- The adapter is installed in a 4x PCIe slot or an underpowered server system.

  If the PCIe bus does not provide 10 Gb/s of throughput due to chipset limitations or the bus width, the adapter cannot maintain 10 Gb/s of incoming receive data. The adapter starts dropping packets quickly. In this situation, it might be beneficial to enable receive flow control in the adapter, and enable flow control in the attached switch for all devices, which helps to slow down the transmitters.
- The adapter transmits to 1GbE devices, especially non-TCP protocol.

  If the adapter transmits to a 10GbE switch with attached 1GbE clients, the adapter can overwhelm the switch. The switch is then forced to start dropping packets because, although it might receive a 10-Gb/s stream, the client can only handle a 1-Gb/s stream. In this situation, it might be beneficial to enable transmit flow control in the adapter, and enable flow control for the 10GbE switch port.

NOTE If multiple PCI functions are exposed for a single 10GbE port, such as in a blade configuration, the flow control parameter must be set the same on all adapters for the port. The results are unpredictable if the setting differs among PCI functions, because this is a shared property of the 10GbE port.

For information on modifying the `Flow Control` parameter, see Section 3.2.1, Configuring NIC Driver Options.

3.2.5.2.2 NUMA Considerations for Windows Server 2012 R2 and Windows Server 2016

NUMA assignments can affect network performance and CPU efficiency. If your application is not NUMA-aware and network traffic is moderate to heavy, the CPU and memory access are managed by the operating system. As a result, the operating system can cross NUMA nodes or your application might be on the same NUMA node as other applications, decreasing your network efficiency. Regardless of whether your application is multi-threaded, and if data is not in parallel, consider the NUMA CPU defaults.

To improve network and CPU performance for heavy network loads under these conditions, you might need to make an appropriate NUMA CPU selection. For example, in Windows Server 2012 R2, and Windows Server 2016, you can use the Task Manager to adjust the `Set Affinity` property to bind the application to a specific NUMA node for maximum network performance and CPU efficiency.

3.2.5.2.3 Checksum Offloading and Large Send Offloading (LSO)

The adapter supports IP and UDP checksum offloading. These protocols are enabled by default. You can disable offloading through the Windows Device Manager Advanced Properties. Disabling checksum offloading is useful only for packet-sniffing applications, such as Ethereal or Microsoft Network Monitor, on the local system where the adapter

is installed and monitored. When packets are sniffed, transmit packets might appear to have incorrect checksums because the hardware has not yet calculated them.

The adapter supports transmit LSO, which allows the TCP stack to send one large block of data, and the hardware segments it into multiple TCP packets. Transmit LSO is recommended for performance, but it can be disabled for packet sniffing applications. LSO data appear as giant packets in the packet sniffer, because the hardware has not yet segmented them.

NOTE On Windows Server 2012 and later, `Recv Segment Coalescing` is enabled by default. You must disable `Recv Segment Coalescing` if you want to set the `Checksum Offload` setting to anything other than `Enabled`.

For information on modifying the `CheckSum Offload` or `Large Send Offload` parameter, see Section 3.2.1, Configuring NIC Driver Options.

3.2.5.2.4 Receive Side Scaling (RSS) for Non-Offloaded IP/TCP Network Traffic

The adapter can process TCP receive packets on multiple processors in parallel. This situation is ideal for applications that are CPU limited. Typically, these applications have numerous client TCP connections that might be short-lived. Web servers and database servers are prime examples. RSS typically increases the number of transactions per second for these applications.

3.2.5.2.5 Understanding RSS

To better understand RSS, it helps to understand the interrupt mechanism used in the network driver. Without RSS, a network driver receives an interrupt when a network packet arrives. This interrupt can occur on any CPU, or it might be limited to a set of CPUs for a given device, depending on the server architecture. The network driver launches one DPC that runs on the same CPU as the interrupt. Only one DPC ever runs at a time. In contrast, with RSS enabled, the network driver launches multiple parallel DPCs on different CPUs.

For example, on a four-processor server that interrupts all processors, without RSS the DPC jumps from CPU to CPU, but it only runs on one CPU at a time. Each processor is busy only 25 percent of the time. The total reported CPU usage of the system is about 25 percent (more if other applications are also using the CPU). This scenario is a sign that RSS might help performance. If the same four-processor server uses RSS, four parallel DPCs can run, one on each processor. The total CPU usage that is available for networking processing is increased from 25 percent to 100 percent.

Some server machines and some network traffic profiles do not benefit from RSS. Because the non-offloaded TCP stack includes a data copy during receive processing, it is possible that memory bandwidth will limit performance before the CPU. In this situation, the CPU usage is very high while all processors wait for memory accesses. To overcome this issue, you can reduce the number of RSS CPUs, or disable RSS entirely.

Poor RSS behavior is typical only in network performance testing applications that receive data, but perform no other processing. For other applications, RSS allows the application to scale other processing tasks across all CPUs, thereby improving overall performance. RSS offers the most benefit for applications that create numerous, short-lived connections. These applications are typically CPU-limited instead of network-bandwidth-limited.

For information on modifying the `RSS Queues` parameter, see Section 3.2.1, Configuring NIC Driver Options.

NOTE Microsoft currently does not schedule RSS processing on all hyper-threaded CPUs. For example, only CPU 1 and CPU 3 have RSS queues on a dual-core, hyper-threaded CPU.

Chapter 4: Troubleshooting

Your system may operate in an unexpected manner in certain circumstances. This section contains reference tables on event codes and error messages and provides information regarding unusual situations.

4.1 General Troubleshooting

The following table describes issues you may encounter and their solutions.

Table 8 General Troubleshooting

Issue	Answer/Solution
The operating system fails to install or does not successfully install the driver.	Verify that the operating system is supported by the driver.
The AutoPilot Installer fails.	If the AutoPilot Installer fails, the **Diagnostics** window shows that the adapter failed. If the adapter fails, perform these steps: 1. Select the adapter to view the reason why the adapter failed. The reason and suggested corrective action are displayed. 2. Perform the suggested corrective action and run AutoPilot Installer again. **NOTE** You can run AutoPilot Installer again from the **Start** menu (**Programs > Emulex > AutoPilot Installer**), or you can run `APInstall.exe` from a command prompt.
The OneInstall Installer fails.	If OneInstall Installer fails, it may be because of one of the following reasons: ■ The operating system prerequisites have not been met. ■ The individual kit installation failed. To check, run the installation interactively. If you encounter error messages when you run the installation interactively, those issues would also apply to an unattended installation. ■ If an individual package failed to install properly, run that package's installer directly. This method displays status and error messages that can be used to diagnose the issue. (The OneInstall Installer does not provide these displays because each package is installed silently.)
Windows Device Manager shows a code 10 or code 39 with a yellow or red exclamation point on the device.	The firmware image does not match the installed device drivers, or the firmware is corrupt. Using the OneCommand Manager application or one of the Windows PE offline or online utilities, install a version of firmware that is compatible with the driver.
The firmware is corrupt or non-responsive.	Using the OneCommand Manager application or one of the Windows PE offline or online utilities, install a version of firmware that is compatible with the driver.
Port names might differ for adapter ports, although they are running the same driver binary.	This is a display issue that does not affect functionality. Run the AutoPilot Installer to correct this issue. Example: Run `elxdrvr-fc-11.x.xxx.xx.exe` and reinstall the driver kit. **NOTE** A reboot may be required after the installation to see the correct names in Device Manager.

4.2 Troubleshooting the NIC Drivers

The followng table provides troubleshooting information for the NIC drivers.

Table 9 Troubleshooting the NIC Drivers

Issue	Answer/Solution
Performance is not as expected.	The adapter may be installed in the wrong type of PCIe slot. Verify that the adapter has been properly installed.
Frequent event log entries for link changes, or statistics that show more than expected CRC errors, occur.	Unload and reload the driver to reset available target IDs. Ensure that the SAN configuration is correct prior to reloading the driver. This action clears the driver's consistent binding table and frees target IDs for new target nodes.
The driver fails to load, and an event log entry states that the driver failed to load due to memory constraints.	There may not be enough memory installed in the system to provide sufficient memory for all devices installed in the system. Try installing more memory if possible.
Unpredictable results occur if the flow control setting differs among PCI functions.	If multiple PCI functions are exposed for a single 10GbE port, the flow control parameter must be set the same on all adapters for the port. Results are unpredictable if the setting differs among PCI functions because this is a shared property of the 10GbE port.
On servers that support PCIe hot unplug, the system may hang or produce a bug check if a PCIe hot unplug or replace is attempted.	Hot unplug is not supported in this release.

The system crashes or appears to hang. In the case of a hang, there could be a message might indicate that the driver experienced a hardware malfunction.	Several possible causes for this issue exist.
	■ Certain systems require an updated BIOS to properly manage the power states of newer Intel and AMD processors. Check with your vendor for information regarding BIOS and firmware updates that may be required to run well with the latest releases of the Windows operating systems. Also, certain BIOS settings may be required. For example, disable any low power processor states and low power settings for PCIe.
	■ On certain AMD systems, it is possible the `intelppm.sys` driver is enabled, and should not be. To query this system driver's run state, log on as administrator and at the command line, type:
	`sc query intelppm`
	If the results indicate that the intelppm driver is running, you must disable it. At the command line, type:
	`sc config intelppm start= disabled`
	■ On all systems, it may be necessary to set the power options to High Performance. Refer to the operating system documentation for details.

On Windows Server 2012 R2, the NIC driver will not load on a VM that is using passthrough in the host.	There is no workaround for this issue. PCI passthrough is not supported.

Appendix A: Error and Event Log Information

A.1 FC /FCoE Error and Event Logs

A.1.1 Viewing the FC /FCoE Error Log

The system event log is a standard feature of Windows Server software. All events logged by the Emulex Storport Miniport will be Event ID 11 with source `ELXFC/LPFCOE`.

To view the error log:

1. Open the **Event Viewer** window by doing one of the following:
   - Click **Start > Programs > Administrative Tools > Event Viewer**.
   - Right-click **My Computer > Manage and Event Viewer** in Computer Management.

   The **Event Viewer** window is displayed.
2. Double-click any event with the source name `ELXFC/LPFCOE`.
3. Examine the entry at offset 0x10 and Event ID 11. The Emulex event code is found in byte 0x10, and supplementary data is in the byte offsets 0x11 through 0x13.

For example, in Figure 11:

byte 0x10 = 9b, byte 0x11 = 00, byte 0x12 = 29, and byte 0x13 = 00

Figure 11 Event Properties

A.1.1.1 Severity Scheme

When the Event Viewer is launched, there are three branches: Application, Security, and System. All ELXFC/LPFCOE error log entries are found under the System branch, and all ELXFC/LPFCOE error log entries have the Event Viewer severity level of "error".

- A severe error code indicates that the driver, firmware, or adapter is behaving abnormally, and your intervention is required to correct the issue.
- A malfunction error code indicates that there is an issue with the system, but your intervention is not required.
- A command error code indicates that an event has transpired, but does not require your intervention. An event may be issue-oriented, such as an invalid fabric command sub-type. An event may not be issue-oriented, such as exhausted retries on PLOGI or PDISC.

A.1.1.2 Related Driver Parameter: LogError

The `LogError` driver parameter determines the minimum severity level to enable entry of a logged error into the system. See Chapter 3: Configuration for instructions on how to set driver parameters.

- If set to `0` = All errors regardless of severity are logged.
- If set to `1` = Severe, malfunction, and command level errors are logged.
- If set to `2` = Both severe and malfunction errors are logged.
- If set to `3` = Only severe errors are logged.

NOTE Set `LogError` to `1` if you are troubleshooting SAN connectivity or device discovery issues.

A.1.1.3 Format of an Error Log Entry

An error log entry takes the form of an event. This event is described by the following items:

- Date (date entry was logged)
- Source (elxfc/elxcna)
- Time (time entry was logged)
- Category (none)
- Type (error)
- Event ID (0)
- User (N/A)
- Computer (name of computer)

A.1.1.4 Error Codes Tables

This section provides tables listing error codes and their descriptions.

A.1.1.4.1 Severe Errors

Table 10 lists severe errors and their codes.

Table 10 Severe Errors

Byte 0x10	Interpretation
0x00	Invalid link speed selection (SLI2-3 mode)
0x01	READ_REV failed (SLI2-3 mode)
0x02	Invalid adapter type (SLI2-3 mode)
0x03	Invalid adapter type (SLI2-3 mode)
0x04	CONFIG_PORT failed

Table 10 Severe Errors (Continued)

Byte 0x10	Interpretation
0x06	READ_CONFIG_failed
0x07	CONFIG_RING 0 failed
0x08	CONFIG_RING 2 failed
0x09	CONFIG_RING 1 failed
0x0A	CONFIG_RING 3 failed
0x0B	INIT_LINK failed (SLI2-3 mode)
0x0C	INIT_LINK failed (SLI2-3 mode)
0x0D	READ_REV failed (SLI2-3 mode)
0x0E	Invalid adapter type (SLI2-3 mode)
0x0F	Invalid adapter type (SLI2-3 mode)
0x10	CONFIG_PORT failed (reinitialization)
0x12	READ_CONFIG command failed (reinitialization)
0x13	CONFIG_RING 0 failed (reinitialization)
0x14	CONFIG_RING 1 failed (reinitialization)
0x15	CONFIG_RING 2 failed (reinitialization)
0x16	CONFIG_RING 3 failed (reinitialization)
0x17	Unresponsive adapter port (SLI2-3 mode)
0x1C	Firmware trap: info1 (SLI2-3 mode)
0x1D	Firmware trap: info2 (SLI2-3 mode)
0x1E	Over-temperature error condition (SLI2-3 mode)
0x1F	Firmware-initiated adapter port reset (SLI2-3 mode)
0x20	Adapter port error attention (SLI2-3 mode)
0x22	Over-temperature warning (SLI2-3 mode)
0x23	Returned to safe temperature (SLI2-3 mode)
0x24	Invalid response tag (SLI2-3 mode)
0x25	Invalid response tag (SLI2-3 mode)
0x26	Invalid response tag (SLI2-3 mode)
0x27	Invalid response sequence (SLI2-3 mode)
0x28	Failure on REG_LOGIN mailbox command
0x29	Unable to initiate fabric binding operation
0x42	Re-simulate FCF after exhausted retries on FLOGI
0x51	ABTS timeout on path and target (0x11: path id; 0x12: target id)
0x2A	Attempted ADISC to non-existent node
0x2B	Failure on iocb context allocation
0x2C	Unable to initiate nport unbinding operation
0x2D	Unable to initiate nport binding operation
0x2E	Failed to allocate resources for Express Lane
0x30	Failure on mailbox context allocation
0x7C	Menlo initialization error
0x7D	Menlo initialization error

Table 10 Severe Errors (Continued)

Byte 0x10	Interpretation
0x7E	Menlo initialization error
0xA0	Failed to initialize adapter port (OneConnect)
0xA1	Failed to initialize adapter port (SLI2-3 mode)
0xCA	Invalid scatter gather list size
0xCB	Unsupported IFType (SLI4 mode)
0xC1	Failed to allocate miniport un-cached extension
0xC2	Insufficient un-cached extension space
0xC3	Port initialization failure (OneConnect)
0xC4	Port initialization failure (SLI2-3 mode)
0xC5	Utility mailbox command error
0xC6	SLI4 Pre-initialization failure
0xC7	UNREG_VPI failure requiring reset
0xC8	Invalid FLOGI response failure requiring reset
0xC9	REG_FCFI failure requiring resolicitation (SLI4 mode)
0xD3	NPIV memory allocation failure
0xE0	Unable to allocate exchange for unsolicited ELS command
0xE1	Misconfigured port event on indicated port. For LPE16000 and LPe32000 link effect and link state (SLI4 mode) 0x13: Port Name; 0x12: Link effect; 0x11: Link state. Link State Values - 0x11 0x0 “Physical Link is functional” 0x1 “Optics faulted/incorrectly installed/not installed – Reseat optics. If issue not resolved, replace.” 0x2 “Optics of two types installed – Remove one optic or install matching pair of optics.” 0x3 “Incompatible optics – Replace with compatible optics for card to function.” 0x4 “Unqualified optics – Replace with Avago optics for Warranty and Technical Support.“ See "Link Effect" 0x5 “Uncertified optics – Replace with Avago-certified optics to enable link operation.“ See "Link Effect" Link Effect 0x12 bit 0 set “Link is non-operational.” bit 0 clear “Link is operational.”
0xF0	Unresponsive adapter port (SLI4 mode)
0xF4	ULP Unrecoverable Error: low part (SLI4 mode)
0xF5	ULP Unrecoverable Error: high part (SLI4 mode)
0xF6	ARM Unrecoverable Error (SLI4 mode)
0xF7	READ_NV failed (SLI4 mode)
0xF8	READ_NV failed (SLI4 mode)
0xF9	READ_REV failed (SLI4 mode)
0xFA	READ_CONFIG failed (SLI4 mode)
0xFB	Failed to post header templates (SLI4 mode)
0xFC	Invalid Completion Queue Entry (SLI4 mode)
0xFD	Invalid Completion Queue Entry (SLI4 mode)
0xFE	Invalid Completion Queue Entry (SLI4 mode)

A.1.1.4.2 Malfunction Errors

Table 11 lists malfunction errors and their codes.

Table 11 Malfunction Errors

Byte 0x10	Interpretation
0x05	SET_VAR command failed
0x11	SET_VAR command failed (reinitialization)
0x21	Spurious mailbox command interrupt
0x31	Unrecognized mailbox command completion
0x32	Duplicate link attention: event tag unchanged
0x33	Invalid link attention: no link state indicated
0x34	Duplicate link attention: link state unchanged
0x35	Error reading common service parameters for port
0x36	Error reading common service parameters for fabric
0x37	Error reading common service parameters for nport
0xB1	Write check error
0x3B	Failed to create node object
0x3C	PRLI initiation failure
0x3D	Recoverable UNREG base VPI error (0x11: mailbox status)
0x3E	Recoverable UNREG VPI error (0x11: mailbox status)
0x42	Exhausted retries on FLOGI
0x45	ELS command rejected
0x49	Exhausted retries on PLOGI
0x4E	World Wide Port Name mismatch on ADISC
0x4F	World Wide Node Name mismatch on ADISC
0x50	ADISC response failure
0x55	LOGO response failure
0x57	PRLI to non-existent node
0x5A	PRLI response error
0x5F	CT command error
0x62	Name server response error
0x66	State Change Notification registration failure
0x6A	Unrecognized ELS command received
0x6F	Received PRLI from un-typed source
0x73	Failed to pend PRLI for authentication
0x77	Failed to allocate Node object
0x7A	REG_VPI failed
0xA3	Command context allocation failure
0xAB	SCSI command error
0xAC	Read check error

Table 11 Malfunction Errors (Continued)

Byte 0x10	Interpretation
0xB0	Node timeout: device removal signaled to Storport
0xB2	FCP_RSP short frame received
0xE1	Misconfigured port event on indicated port. For LPE16000 and LPe32000 link effect and link state (SLI4 mode) 0x13: Port Name; 0x12: Link effect; 0x11: Link state. Link State Values - 0x11 0x0 "Physical Link is functional" 0x1 "Optics faulted/incorrectly installed/not installed – Reseat optics. If issue not resolved, replace." 0x2 "Optics of two types installed – Remove one optic or install matching pair of optics." 0x3 "Incompatible optics – Replace with compatible optics for card to function." 0x4 "Unqualified optics – Replace with Avago optics for Warranty and Technical Support." See "Link Effect" 0x5 "Uncertified optics – Replace with Avago-certified optics to enable link operation." See "Link Effect" Link Effect 0x12 bit 0 set "Link is non-operational." bit 0 clear "Link is operational."

A.1.1.4.3 Command Errors

Table 12 lists command errors and their codes.

Table 12 Command Errors

Byte 0x10	Interpretation
0x43	Fabric logon succeeded
0x46	ELS command failed
0x47	Exhausted retries on ELS command
0x4A	PLOGI accepted
0x56	LOGO accepted
0x59	PRLI accepted
0x63	Fabric name server response
0x6B	ELS RSCN processed
0x71	LOGO received from fabric
0x79	FDISC accepted
0xA2	SCSI address assigned to discovered target
0xA4	Report LUNs error (initial I/O to discovered target)
0xA5	Local error indication on FCP command
0xA6	FCP Command error
0xA8	Data overrun
0xA9	FCP command error
0xAA	SCSI check condition
0xAD	Local reject indication on FCP command

Table 12 Command Errors (Continued)

Byte 0x10	Interpretation
0x43	Fabric logon succeeded
0xAE	Error on SCSI pass-through command
0xAF	Error on Menlo CT command
0xE1	Misconfigured port event on indicated port. For LPE16000 and LPe32000 link effect and link state (SLI4 mode) 0x13: Port Name; 0x12: Link effect; 0x11: Link state. Link State Values - 0x11 0x0 "Physical Link is functional" 0x1 "Optics faulted/incorrectly installed/not installed – Reseat optics. If issue not resolved, replace." 0x2 "Optics of two types installed – Remove one optic or install matching pair of optics." 0x3 "Incompatible optics – Replace with compatible optics for card to function." 0x4 "Unqualified optics – Replace with Avago optics for Warranty and Technical Support." See "Link Effect" 0x5 "Uncertified optics – Replace with Avago-certified optics to enable link operation." See "Link Effect" Link Effect 0x12 bit 0 set "Link is non-operational." bit 0 clear "Link is operational."

A.1.1.4.4 Event Indicators

Table 13 lists event indications and their codes.

Table 13 Event Indications

Byte 0x10	Interpretation
0x18	Port shutdown event (SLI2–3 mode)
0x19	Port in off-line state (SLI2–3 mode)
0x1A	Port in on-line state (SLI2–3 mode)
0x1B	Port in off-line state (SLI2–3 mode)
0xA7	Data underrun
0xD0	NPIV Virtual Port creation success (Virtual Port Did 0x11–0x13)
0xD1	NPIV Virtual Port creation failed (Virtual Port index 0x11–0x13)
0xD2	NPIV Virtual Port FDISC failed (Virtual Port index 0x11–0x13)
0xD4	Exceeded max Virtual Port supported (Virtual Port index 0x11–0x13)
0xD5	NPIV Virtual Port removal (Virtual Port Did 0x11–0x13)
0xEE	Authenticated successfully (remote Did 0x11–0x13)
0xEF	Failed to authenticate (remote Did 0x11–0x13)
0xE2	Authentication not support (remote Did 0x11–0x13)
0xE3	Authentication ELS command timeout (remote Did 0x11–0x13)
0xE4	Authentication transaction timeout (remote Did 0x11–0x13)
0xE5	LS_RJT other than Logical Busy received for Authentication transaction (remote Did 0x11–0x13)
0xE6	LS_RJT Logical Busy received for Authentication Transaction (remote Did 0x11–0x13)
0xE7	Received Authentication Reject other than Restart (remote Did 0x11–0x13)
0xE8	Received Authentication Reject Restart (remote Did 0x11–0x13)
0xE9	Received Authentication Negotiate (remote Did 0x11–0x13)

Table 13 Event Indications (Continued)

Byte 0x10	Interpretation
0xEA	Authentication spurious traffic (remote Did 0x11–0x13)
0xEB	Authentication policy has been changed (remote Did 0x11–0x13)
0xED	Same passed were set for both local and remote entities (remote Did 0x11–0x13)
0xF1	Port shutdown event (SLI4 mode)
0xF2	Port in off-line state (SLI4 mode)
0xF3	Port in on-line state (SLI4 mode)

A.1.2 Viewing the FC and FCoE Event Log

This section provides information on the FC and FCoE event logs.

A.1.2.1 Event Log Interpretation

- All events logged by Emulex Storport Miniport are in Event ID 11 with source "elxfc/lpfcoe".
- The Storport Miniport driver parameter `LogErrors` determines what type of events are logged by the driver; the default setting is `3`, which logs only events of a SEVERE nature; the optional setting of `2` logs events of both SEVERE and MALFUNCTION type; and the optional setting of `1` logs events of SEVERE, MALFUNCTION, and COMMAND type.

  NOTE For troubleshooting SAN connectivity or device discovery issues, set the `LogErrors` to `1`.
- The Emulex event code is found in byte 0010 and supplementary data is in byte offsets 0011 through 0013.

A.1.2.2 Additional Event Log Information

The following tables are not comprehensive but do include the codes that are most likely to appear in SAN environments where issues occur.

A.1.2.2.1 ELS/FCP Command Error Status Codes

Table 14 lists the internal firmware codes posted by the adapter firmware that explain why a particular ELS or FCP command failed at the FC level.

Table 14 ELS/FCP Command Error Status Codes

Code	Explanation
0x2	Remote Stop – Remote port sent an ABTS
0x3	Local Reject – Local Reject error detail
0x9	LS_RJT Received – Remote port sent LS_RJT
0xA	A_RJT Received – Remote port sent BA_RJT

A.1.2.2.2 CT Command Response Codes

Table 15 lists the codes that indicate the response to a FC Common Transport protocol command.

Table 15 CT Command Response Codes

Code	Explanation
0x8001	FC Common Transport Reject
0x8002	FC Common Transport Accept

A.1.2.2.3 FC-CT Reject Reason Codes

Table 16 lists the codes that indicate the reason a CT command was rejected.

Table 16 FC-CT Reject Reason Codes

Code	Explanation
0x01	Invalid command code
0x02	Invalid version level
0x05	Logical busy
0x07	Protocol error

A.1.2.2.4 ELS Command Codes

Table 17 lists the FC protocol codes that describe the Extended Link Services commands that were sent.

Table 17 ELS Command Codes

Code	Explanation
0x01	Link Service Reject (LS_RJT)
0x02	Accept (ACC)
0x03	N_Port Login (PLOGI)
0x04	Fabric Login (FLOGI
0x05	N_Port Logout (LOGO)
0x20	Process Login (PRLI)
0x21	Process Logout (PRLO)
0x51	Discover F_Port Service Params (FDISC)
0x52	Discover Address (ADISC)
0x61	Register State Change Notify (RSCN)

A.1.2.2.5 SCSI Status Codes

Table 18 lists the SCSI status codes returned from a SCSI device that receives a SCSI command.

Table 18 SCSI Status Codes

Code	Explanation
0x00	GOOD
0x02	CHECK CONDITION
0x08	BUSY
0x18	RESERVATION CONFLICT
0x28	QUEUE FULL

A.1.2.2.6 Local Reject Status Codes

Table 19 list the codes supplied by the Emulex adapter firmware that indicate why a command was failed by the adapter.

Table 19 Local Reject Status Codes

Code	Explanation
0x02	SEQUENCE TIMEOUT – Possible bad cable/link noise
0x04	INVALID RPI – Occurs if the link goes down
0x05	NO XRI – Possible host or SAN problem
0x0D	TX_DMA FAILED – Possible host system issue
0x0E	RX_DMA FAILED – Possible host system issue
0x0F	ILLEGAL FRAME – Possible bad cable/link noise
0x11	NO RESOURCES – Port out of exchanges or logons
0x18	LOOP OPEN FAILURE – FC_AL port not responding
0x1A	LINK DOWN – Queued cmds returned at link down
0x1D	OUT OF ORDER DATA – Possible bad cable or noise

A.1.2.2.7 SRB Status Codes

Table 20 lists the SCSI Request Block status codes provided by the driver to the operating system based upon the response from a SCSI device in the SAN.

Table 20 SRB Status Codes

Code	Explanation
0x04	ERROR
0x05	BUSY
0x09	TIMEOUT
0x0A	SELECTION TIMEOUT
0x0B	COMMAND TIMEOUT
0x0E	BUS RESET
0x12	DATA OVERUN

A.1.2.3 ASC/ASCQ

Additional Sense Code/Additional Sense Code Qualifier information can be found in any SCSI specification document; these codes contain detailed information about the status or condition of the SCSI device in question.

A.1.2.4 Additional Notes on Selected Error Codes

These error codes might be seen more frequently than others, or that indicate conditions that you might be able to solve by investigation and correction of issues in the SAN configuration.

NOTE The nomenclature of "0x" is used as the prefix for the byte code fields because those byte codes are actually hexadecimal values.

A.1.2.4.1 Node Timeout (Code 0xAA)

This event code indicates that a particular device has not been found (if the message is logged during device discovery) or that a particular device has been removed from the fabric. If this message appears, determine if there is

something wrong with the connection of that device to the SAN (cables, switches or switch ports, or status of the target device itself).

A.1.2.4.2 SCSI Command Error (Code 0x9A) and SCSI Check Condition (Code 0x9B)

Code 0x9A indicates that the SCSI command to a particular device was responded to with an error condition (the target and LUN information, along with the SCSI status, are provided).

In the specific case of code 0x9B, this code indicates that the device responded with the specific status of Check Condition – the ASC/ASCQ information provided in bytes 0x12 and 0x13 allows you to find out the status being reported by the target and to determine if there is an action that can be performed to return the device to functional status.

A.1.2.4.3 Name Server Response (Code 0x98)

This code is useful in determining if the expected number of targets in a SAN configuration are being presented by the name server to the requesting adapter. The number in byte 0x11 is the number of targets returned to the name server query made by the adapter. If the number of targets does not match expectations, examine the SAN configuration found in the switch tables and if that information shows targets or devices still missing, check connections between the switch ports and those devices.

A.1.2.4.4 Context Allocation Failures

A number of event codes for which the interpretation contains the phrase "context allocation failure" exist. These types of events refer to the internal memory constructs of the Emulex Storport Miniport driver and, as such, are intended for Broadcom design engineers' information. If you encounter this type of code, contact Broadcom Technical Support for assistance.

NOTE Context allocation failures are rare.

A.2 NIC Error and Event Logs

This section provides information on NIC error and event logs.

A.2.1 Viewing the NIC Error Log

For Windows Server operating systems, the network driver generates error codes in the system event log. These error codes can be viewed by using the Event Viewer application.

To view the error codes, perform these steps:

1. Click the **Start** tab on the bottom of the screen.
2. Click **Run**.
3. Type `eventvwr` and click **OK**.
4. Click **Windows Log**.
5. Click **System**.
6. Click the `be2net` error under System Events to show the event details.

A.2.2 NIC Event Log

The Windows Device Manager generates error log codes if any errors occur during the installation of the NIC driver. Each log contains a Message ID, Severity, and Symbolic Link. The Message ID is unique and tracks the error message (if

not displayed). Table 21 shows the list of error codes, the severity of the error, the message displayed, the meaning of the error, and recommended resolutions. When reporting an issue with the adapter to Broadcom, check the event log and report if any of these entries that are displayed.

Table 21 NIC Event Log Entries

Message ID Hexadecimal/ Decimal	Severity	Message	Recommended Resolution
0x4000003AL	Informational	N/A	N/A
0x0000038L/56	Warning	The device firmware does not support ETS functionality in SR-IOV or multichannel mode.	Revert to default mode for ETS support.
0x0049/73	Informational	%2 : Correct optics installed. (%2 is a place holder for the NIC controller name.)	This message is informational.
0x0046/70	Warning	Unqualified SFP+ module detected on %2, Port %3 from %4 part number %5. (%2 is a place holder for the NIC controller name. %3, %4, %5 are place holders for other arguments in the log message.)	Replace the SFP+ module.
0x0045/69	Informational	SFP+ module detected on %2, Port %3 from %4 part number %5. (%2 is a place holder for the NIC controller name. %3, %4, %5 are place holders for other arguments in the log message.)	This message is informational.
0x00037/55	Warning	This adapter may have an issue recovering from corrupted use of SR-IOV. Assigning an SR-IOV device to a Virtual Machine could leave the system vulnerable, and lead to instability. Assign SR-IOV devices only to Virtual Machines that run trusted workloads, or consider disabling the use of SR-IOV.	This adapter exposes a vulnerability to the VM that may allow the VM to crash the entire physical computer. This is no different than running a physical adapter. SR-IOV should be used only if the VM has a trusted server administrator.
0x00036/54	Warning	Incompatible optics. Replace with compatible optics for card to function.	Replace the incompatible SFP transceivers with compatible ones for the adapter to function correctly.
0x00035/53	Warning	Optics of two types installed-Remove one optic or install matching pair of optics.	Remove one SFP transceiver or install a matching pair of SFP transceivers.
0x00034/52	Warning	Optics faulted/incorrectly installed/not installed. Reseat optics, if issue not resolved, replace.	Reseat the SFP transceiver. If the issue is not resolved, replace the transceiver.
0x00033/51	Warning	SR-IOV virtualization failed initialization. Check system BIOS settings, or disable SR-IOV for the adapter.	Check system BIOS settings, or disable SR-IOV for the adapter.
0x00032/50	Warning	The Ethernet link is down due to PHY over-temperature condition. Improve cooling for the device.	Improve the cooling conditions for the device.
0x00031/49	Warning	RSS is limited to 4 queues. Enable Advanced Mode in the PXE BIOS to use up to 16 queues. This may require a firmware update.	Enable Advanced Mode in the PXESelect BIOS utility during boot to use up to 16 queues. This may require a firmware update. Refer to the Broadcom website for compatible firmware.

Table 21 NIC Event Log Entries (Continued)

Message ID Hexadecimal/ Decimal	Severity	Message	Recommended Resolution
0x00030/48	Warning	SR-IOV is not enabled. Update the firmware, enable SR-IOV in the server BIOS, and enable SR-IOV and Advanced Mode in the PXE BIOS.	Update the firmware, enable SR-IOV in the server BIOS, and enable SR-IOV and Advanced Mode in the PXESelect BIOS utility. Refer to the Broadcom website for compatible firmware.
0x0002f/47	Warning	VMQ offload is disabled. Disable SR/IOV support in PXE BIOS to use VMQ.	Disable SR-IOV support in the PXESelectBIOS utility to use VMQ.
0x0002e/46	Error	Device is not supported on Windows 7 Operating System.	
0x0002d/45	Error	Error recovery failed. The device is no longer operational. Update all drivers and firmware.	Refer to the Broadcom website for compatible firmware and drivers.
0x0002c/44	Warning	Error recovery is disabled on the system. The device is no longer operational.	This message is informational.
0x0002b/43	Informational	The driver successfully recovered from an error.	This message is informational.
0x0026/38	Warning	The device firmware does not support RSS functionality for this network adapter.	The firmware and the driver are not compatible versions. Refer to the Broadcom website for compatible firmware and drivers.
0x0024/36	Error	The device firmware does not support network functionality.	The firmware and the driver are not compatible versions. Refer to the Broadcom website for compatible firmware and drivers.
0x0023/35	Warning	The Ethernet link is down due to a remote fault.	The Ethernet link is down due to the remote partner signaling a fault. Check the peer device for errors.
0x0022/34	Warning	The Ethernet link is down due to a local fault.	The Ethernet link is down due to a link-down event detected at the driver.
0x0021/33	Informational	Network device is operating in Gen2 mode and installed in a 4x PCIe slot.	For best performance, install the adapter in an 8x Gen2 PCIe slot. **NOTE** A 16x slot does not provide any additional performance.
0x0020/32	Informational	The network device is operating in Gen2 mode and installed in a 1x PCIe slot.	For best performance, install the adapter in an 8x Gen2 PCIe slot. **NOTE** A 16x slot does not provide any additional performance.
0x001f/31	Informational	The network device is operating in Gen1 mode and installed in a 8x PCIe slot.	For best performance, install the adapter in an 8x Gen2 PCIe slot. **NOTE** A 16x slot does not provide any additional performance.
0x001e/30	Informational	The network device is operating in Gen1 mode and installed in a 4x PCIe slot.	For best performance, install the adapter in an 8x Gen1 PCIe slot. **NOTE** A 16x slot does not provide any additional performance.
0x001d/29	Informational	The network device is operating in Gen1 mode and installed in a 1x PCIe slot.	For best performance, install the adapter in an 8x Gen1 PCIe slot. **NOTE** A 16x slot does not provide any additional performance.

Table 21 NIC Event Log Entries (Continued)

Message ID Hexadecimal/ Decimal	Severity	Message	Recommended Resolution
0x0015/21	Warning	Firmware version does not match driver version.	The firmware version and driver must match. This is a warning message, but Broadcom recommends that you reinstall matching versions of the firmware and driver.
0x0014/20	Error	Failed to read registry configuration.	The registry is corrupted. Reinstall the driver or the operating system.
0x0013/19	Error	Resource conflict.	The operating system failed to allocate resources for the device. Check low memory conditions and operating system hardware resource conflicts.
0x0012/18	Error	Failed to enable bus mastering.	Verify that the BIOS allows bus mastering and that no resource conflicts exist.
0x0011/17	Error	The driver is incompatible with the device.	The driver is loaded on the incorrect hardware device. Verify that the correct driver is installed.
0x0010/16	Warning	The network driver was reset.	This message may indicate a system hang or hardware issue. Verify that other system devices are working properly.
0x000c/12	Informational	The Ethernet link is down.	This message is informational.
0x000b/11	Informational	The Ethernet link is up.	This message is informational.
0x000a/10	Error	The network device detected an error.	A hardware error occurred. Verify that the firmware image is not corrupted. Contact Broadcom Technical Support.
0x0009/9	Error	Failed to register interrupt service routine.	This message indicates an NDIS error. Verify that hardware resource conflicts do not exist.
0x0008/8	Error	Failed to get TCP offload handlers.	This message indicates an NDIS error. Verify that the NDIS version is valid for the driver.
0x0007/7	Warning	A memory allocation failure occurred during driver load. Performance may be reduced.	This warning occurred due to a failed memory allocation. Check low memory conditions. Use a smaller MTU to reduce driver memory requirements.
0x0006/6	Error	Driver load failed due to memory allocation failure	This failure occurred due to a failed memory allocation in the driver. Check low memory conditions.
0x0005/5	Error	Failed to register scatter gather DMA.	This failure occurred due to a failed memory allocation in the operating system. Check low memory conditions.
0x0004/4	Error	Failed to map device registers.	This failure occurred due to a failed memory allocation in the operating system. Check low memory conditions.
0x0003/3	Error	Unsupported medium.	This message indicates an internal NDIS error. Check the operating system installation.
0x0002/2	Error	The network driver initialization failed.	This error may be a firmware driver mismatch or corrupt installation. Check the firmware version, reinstall the firmware and try again. This message may also indicate a hardware issue.
0x0001/1	Informational	The driver successfully loaded.	This message is informational and indicates successful loading of the device driver.

Appendix B: Port Speed Specifications for LPe16202/OCe15100 Adapters

An adapter can support only one Ethernet port speed at a time, and the preference is always for 10 Gb/s. The type of module used (copper or optical) does not make a difference. As soon as a 10 Gb module is plugged into one of the ports, the adapter switches to 10 Gb/s no matter what speed the other port is running, or even if I/O is running on that port. This behavior is a per-adapter constraint; another adapter can be running on a different speed.

Table 22 lists negotiated speed specifications per an adapter's port connection:.

Table 22 Negotiated Speed Specification per Adapter Port Connection

Port 0	Port 1	Port Link	Status Speed
10 Gb/s	10 Gb/s	Both ports link up	10 Gb/s
10 Gb/s	1 Gb/s	Only Port 0 links up	10 Gb/s
1 Gb/s	10 Gb/s	Only Port 1 links up	10 Gb/s
1 Gb/s	1 Gb/s	Both ports link up	1 Gb/s
1 Gb/s	–	Only Port 0 links up	1 Gb/s
–	1 Gb/s	Only Port 1 links up	1 Gb/s
10 Gb/s	–	Only Port 0 links up	10 Gb/s
–	10 Gb/s	Only Port 1 links up	10 Gb/s

Appendix C: AutoPilot Installer Command Line and Configuration File Parameters

AutoPilot Installer can initiate an installation from a command prompt or script. You can run the AutoPilot Installer manually from the command line or a script, or you can run it automatically through the driver kit. When run manually from the command line or script, the command line parameters can be passed.

If you specify the `/q` switch with the driver kit installer command, the driver kit installer runs in unattended mode and automatically invokes `APInstall.exe` with its `/silent` switch. See Section 2.3.6, Unattended Driver Installation, for additional information.

C.1 APArg Driver Kit Parameter and Appending to the APInstall.exe File

If you specify a value for the `APargs` driver kit parameter, this value is appended to the `APInstall.exe` command line. For example, if you execute this installer file as:

```
elxdrv-fc-<version>.exe /q APargs=SilentRebootEnable=True
```

After installing the AutoPilot Installer, the driver kit automatically executes it as:

```
APInstall.exe /silent SilentRebootEnable=True
```

To specify more than one parameter, separate the settings by one or more spaces and put quotes around the entire `APargs` expression. For example, type the following command on one line:

```
elxdrv-fc-<version>.exe "APargs=SilentRebootEnable=True localDriverLocation =
"d:\drivers\new\Storport"
```

This results in the AutoPilot Installer being run as:

```
APInstall.exe SilentRebootEnable=True localDriverLocation =
"d:\drivers\new\Storport"
```

Parameter values that contain spaces, such as path names, must be enclosed in double quotation marks. To add such a setting to `APargs`, you must insert backslashes before the quotes around the value, and then add double quotation marks around the entire `APargs` expression. For example, the command line (all on one line):

```
elxdrv-fc-<version>.exe "APargs=ReportLocation=\"C:\Documents and
Settings\Administrator\My Documents\reports\""
```

This results in AutoPilot Installer being run as:

```
APInstall.exe ReportLocation="C:\Documents and Settings\Administrator\My
Documents\reports"
```

To pass multiple parameters to the AutoPilot Installer and minimize errors, you can run the utility kit installer interactively, delay AutoPilot Installer execution, and then run the AutoPilot Installer command. The procedure is described in Section 2.3.2.2, Option 2: Run AutoPilot Installer Separately and Section 2.3.6, Unattended Driver Installation.

You can specify a non-default directory for the driver kit by specifying an installation folder on the command line. For example:

```
elxdrv-fc-<version>.exe install:"C:\Emulex"
```

This option can be used in conjunction with the `APArgs` directive.

C.2 AutoPilot Installer Syntax

The syntax used to run AutoPilot Installer silently from a command line or script is:

`APInstall [/silent] [parameter setting][parameter setting...]`

The `/silent` switch and parameter settings can occur in any order. One or more spaces must separate the switch and each parameter setting.

The syntax of a parameter setting is:

`parameter_name =["]value["]`

Double quotation marks are required only around values that contain spaces. Spaces can separate parameters, equal signs, and values. Parameter names and values are not case-sensitive.

The `APInstall` command can contain the settings listed below. Each setting, except `ConfigFileLocation`, can also be specified in the AutoPilot configuration file. For descriptions of each parameter, see Section C.2.3, Software Configuration Parameters.

Settings specified in the `APInstall` command override those specified in the configuration file.

```
ConfigFileLocation = path-specifier
NoSoftwareFirstInstalls = { TRUE | FALSE }
SilentRebootEnable = { TRUE | FALSE }
ForceDriverUpdate = { TRUE | FALSE }
ForceDriverTypeChange = { TRUE | FALSE }
SkipDriverInstall = { TRUE | FALSE }
InstallWithoutQFE = { TRUE | FALSE }
ForceRegUpdate = { TRUE | FALSE }
LocalDriverLocation = path-specifier
ReportLocation = path-specifier
```

C.2.1 Path Specifiers

Paths can be specified as:

- An explicit path:

```
ReportLocation="C:\Program Files\Emulex\AutoPilot Installer\Reports"
```

- A relative path:

```
LocalDriverLocation="Drivers\Storport Miniport\"
```

  (assuming installation into `C:\Program Files\Emulex\AutoPilot Installer\`, this path would logically become `C:\Program Files\Emulex\AutoPilot Installer\Drivers\Storport Miniport\`)

- With the `%ProgramFiles%` environment variable:

```
LocalDriverLocation = "%ProgramFiles%\Emulex\AutoPilot Installer\Driver"
```

C.2.2 Configuration File Location

The optional `ConfigFileLocation` setting contains the path to the configuration file that should be used. If this parameter is not specified, AutoPilot Installer uses the file named `APInstall.cfg` in the same folder as `APInstall.exe`.

The format is the same as that of the other path settings.

Example:

```
APInstall /silent SkipDriverInstall=True configFileLocation=MyConfiguration.cfg
```

C.2.3 Software Configuration Parameters

C.2.3.1 DiagEnable (Running Diagnostics)

NOTE The `DiagEnable` parameter cannot be specified on the command line; it must be specified within the configuration file.

Default: True

By default, AutoPilot Installer runs its diagnostics after all driver installation tasks have been completed. To disable this function, set this parameter to False.

C.2.3.2 ForceDriverTypeChange (Forcing a Driver Type Change)

Default: False

When installing a driver, set this parameter to True to cause Silent mode installations to update or install the Storport Miniport driver on each adapter in the system, without regard for the currently installed driver type (replacing any installation of the SCSIport Miniport or FC Port driver).

C.2.3.3 ForceDriverUpdate (Forcing a Driver Version Update)

Default: False

By default, if the same version of the driver is already installed, an unattended installation proceeds with installing only the utilities. To force a driver update even if the same version of the driver is installed, set this parameter to True.

NOTE `ForceDriverUpdate` applies only to unattended installations; in interactive installations, this parameter is ignored and you are asked if the driver should be updated.

C.2.3.4 ForceRegUpdate (Forcing an Update of an Existing Driver Parameter Value)

Default: False

The `ForceRegUpdate` driver parameter setting determines whether existing driver parameters are retained or changed when you update the driver. By default, all existing driver parameter settings are retained. The `ForceRegUpdate` parameter does not affect any existing persistent bindings. To set up an installation to remove the existing driver parameters from the registry and replace them with parameters specified in the AutoPilot configuration file, set this parameter to True.

NOTE You can use this setting for attended installations with the **AutoPilot Installer** wizard if you modify the AutoPilot configuration file in AutoPilot Installer.

C.2.3.5 LocalDriverLocation (Specifying Location to Search for Drivers)

Default: Drivers (The default `Drivers` folder is located in the same folder as AutoPilot Installer.)

You can specify a local location that is to be searched for drivers during unattended installations. The location can be a local hard drive or a network shared drive. Removable media are not searched.

Example:

```
LocalDriverLocation = "d:\drivers\new\Storport"
```

NOTE On x64 and 32-bit systems, the path specified by `LocalDriverLocation` must contain at least one instance of an

FC, FCoE, and NIC driver. AutoPilot Installer automatically selects the most recent revisions that it finds.

C.2.3.6 NoSoftwareFirstInstalls (Prohibiting Software First Installations)

Default: False

If this parameter is set to True, AutoPilot Installer prevents unattended installations from performing software-first installations. This way you can run an automated installation on multiple machines in your network, but only machines with Emulex adapters actually have Emulex drivers updated or installed.

If this parameter is omitted from the configuration file or explicitly set to True, the page is not displayed. AutoPilot Installer uses configuration file parameters to determine the appropriate management mode.

C.2.3.7 ReportLocation (Setting Up an Installation Report Title and Location)

The automatically generated file name for this report is:

"report_*mm*-*dd*-*yy*.txt"

where *mm* is the month number, *dd* is the day, and *yy* indicates the year.

You can change only the installation report folder; the file name is auto-generated. In the following example, *x* could be any available drive:

```
ReportLocation = "x:\autopilot\reports\installs\"
```

C.2.3.8 SilentInstallEnable (Enabling Unattended Installation)

NOTE Setting the `SilentInstallEnable` parameter to true in the configuration file is functionally equivalent to supplying the `/silent` switch on the command line. You cannot specify the `SilentInstallEnable` parameter on the command line.

Default: False

Setting this parameter to True causes AutoPilot Installer to operate with no user interaction.

C.2.3.9 SilentRebootEnable (Enabling Silent Reboot)

Default: False

AutoPilot Installer's default behavior in unattended installations does not restart the system. AutoPilot Installer continues with the installation. Restarts often require you to log on as part of the Windows start up process. If there is no logon, the installation process would stop if the system is restarted. However, Windows can be configured to start up without requiring you to log on. You must ensure that it is safe to restart the system during unattended installations if you set this parameter to True.

C.2.3.10 InstallWithoutQFE (Enabling Installation if a QFE Check Fails)

Default: False

AutoPilot Installer checks for Microsoft's QFEs, also known as KB updates, based on the checks you have specified in the [`STORPORT.QFES`] section. By default, the installation terminates if the QFE check fails. To enable a driver installation to proceed even if a check for QFEs fails, set this parameter to True.

C.3 AutoPilot Configuration File

The AutoPilot configuration file is organized into sections, grouped according to related commands. Six main sections exist:

- [AUTOPILOT.ID] – Configuration Identification
- [AUTOPILOT.CONFIG] – Software Configuration
- [STORPORT.CONFIGURATION] – Configuration Prompts and Vendor-Specific Questions
- [STORPORT.QFES] – QFE Checks
- [STORPORT.PARAMS] – Setting Up FC Driver Parameters
- [SYSTEM.PARAMS] – Setting Up System Parameters

Each section begins with a heading. The heading is required even if there are no settings in the section. The only section not required is the Configuration Prompts section, which has the heading [STORPORT.CONFIGURATION]. That section cannot exist if AutoPilot Installer runs in Silent mode. You must delete or comment-out that entire section for unattended installation.

Lines that begin with a semicolon (;) are comments. Some of the comments are sample settings. To use the setting, remove the semicolon.

C.3.1 Using the Windows Environment Variable (%ProgramFiles%)

You can use the Windows %ProgramFiles% environment variable in the LocalDriverLocation and ReportLocation strings within the configuration file. This variable allows you to specify strings in a driver-independent manner, allowing the same configuration file to be used on different systems where Windows may have been installed on different drives. To use this option, %ProgramFiles% must be the first component specified in the string. The portion of the string that follows is appended to the contents of the %ProgramFiles% environment variable. For example:

```
ReportLocation = "%ProgramFiles%\my company\reports"
```

NOTE The contents of the %ProgramFiles% environment variable is not terminated with a slash, so you must provide one in the string. Windows environment variables are not case-sensitive.

C.3.2 Configuration Identification [AUTOPILOT.ID]

This section appears at the beginning of every AutoPilot configuration file and contains revision and label information. The revision entry identifies the file's version number and the date on which it was produced. The label entry identifies the configuration that the file supports. This section may appear only once in the APInstall.cfg file.

C.3.3 Software Configuration [AUTOPILOT.CONFIG]

This section contains settings that control and configure AutoPilot Installer and the OneCommand Manager application operation. This section can appear only once in the AutoPilot configuration file. See Section C.2.3, Software Configuration Parameters, for information about settings that can be specified in this section.

C.3.4 Configuration Prompts and Vendor-Specific Questions [STORPORT.CONFIGURATION]

NOTE You must remove or comment-out the entire [STORPORT.CONFIGURATION] section for an unattended installation.

A [STORPORT.CONFIGURATION] section can exist in the AutoPilot configuration file. The first items in this section are the driver parameters to be used regardless of how the questions are answered. These items are followed by a subsection that contains questions (these may be vendor-specific questions). A line containing [QUESTIONS] marks the start of the subsection, and the end of it is marked by a line containing [ENDQUESTIONS]. Within the question subsection there can be as many questions as needed. Each question uses the format:

```
question= "question?", "explanation", "answer0", "answer1", "answer2",.... ,
"answern"
```

Where:

- "*question*?" contains the text of the question to be asked.
- "*explanation*" contains brief text to help explain the question. The explanation displays below the question in a smaller font. If there is no explanatory text, empty quotation marks must be used in its place.
- "*answer0*" contains the first answer to be displayed in the drop-down list.
- "*answer1*" contains the second answer to be displayed in the drop-down list.
- "*answern*" contains the nth answer to be displayed in the drop-down list.

For each question there can be as many answers as needed. For each answer, there must be a corresponding "*answer =*" section with its corresponding driver parameters listed beneath it. The answer uses the format:

```
answer = 0
DriverParameter="Param1=value; Param2=value;"
answer = 1
DriverParameter="Param1=value; Param2=value;"
....
answer = n
DriverParameter="Param1=value; Param2=value;"
```

C.3.4.1 Example of [STORPORT.CONFIGURATION] Section:

```
[STORPORT.CONFIGURATION]
;The first section contains the driver parameters common to all configurations, no
matter what answers are given.
DriverParameter="EmulexOption=0;"
[QUESTIONS]
question = "What is your link speed?", "Note: select 'Auto-detect' if you are
unsure about the answer.", "4GB", "2GB", "1GB", "Auto-detect"
ANSWER = 0
DriverParameter = "LinkSpeed=4;"  ;4 GB
ANSWER = 1
DriverParameter = "LinkSpeed=2;"  ;2 GB
ANSWER = 2
DriverParameter = "LinkSpeed=1;"  ;1 GB
ANSWER = 3
DriverParameter = "LinkSpeed=0;"  ;Auto-detect question = "Describe the topology
of your storage network.", "Note: Select 'Arbitrated Loop' when directly connected
to the array (no fibre switch). Select 'Point-to-Point' when connected to a SAN
(fibre switch).", "Arbitrated Loop", "Point-to-Point"
ANSWER = 0
```

```
DriverParameter = "Topology=2;"
ANSWER = 1
DriverParameter = "Topology=3;"
[ENDQUESTIONS]
[END.STORPORT.CONFIGURATION]
```

C.3.5 QFE Checks [STORPORT.QFES]

This section specifies an additional QFE check, also known as KB updates, during installation. To add a Windows QFE check to the configuration file, edit the [STORPORT.QFES] section in the AutoPilot configuration file. You can place this section anywhere within the file as long as it is not contained within another section. This section contains a single line for each QFE that is to be checked. Up to 10 lines are checked; more than that can exist, but they are ignored. All parameters in each line must be specified. These lines have the format:

```
qfe = "qfe name", "path and file name", "file version", "applicable OS"
```

qfe name	The name of the item being checked; for example, QFE 2846340. The name should facilitate searching Microsoft's website for any required code updates.
path and file name	This string identifies the file to be checked and its location relative to the Windows home folder. In most cases, the file to check is the Microsoft Storport driver; for example: "\system32\drivers\storport.sys". This string is also used in dialogs and log file messages.
file version	This version is the minimum version that the file to be checked must have for the QFE to be considered installed. It is specified as a text string using the same format as is used when displaying the files property sheet; for example: "5.2.1390.176".
applicable OS	This is used to determine if the QFE applies to the operating system platform present.

For example:

```
[STORPORT.QFES]
qfe = "QFE 83896", "\system32\drivers\storport.sys", "5.2.1390.176","Win2012"
```

C.3.6 Setting Up FC Driver Parameters [STORPORT.PARAMS]

This section specifies driver parameters. Parameters are read exactly as they are entered and are written to the registry. To change driver parameters, modify this section of the AutoPilot configuration file. Locate the [STORPORT.PARAMS] section in the AutoPilot configuration file. This section follows Optional Configuration File Changes. Under the [STORPORT.PARAMS] heading, list the driver parameters and new values for the driver to use.

For example:

```
Driver Parameter = "LinkTimeout = 45"
```

See Table 1, Storport Miniport Driver Parameters, for a listing of driver parameters, defaults, and valid values.

C.3.7 Setting Up System Parameters [SYSTEM.PARAMS]

To change the system parameters, create a [SYSTEM.PARAMS] section in the APInstall.cfg file. Create this section under the Optional Configuration File Changes heading in the [AUTOPILOT.CONFIG] section.

For example, you can adjust the operating system's global disk timeout. The timeout is stored in the registry under the key HKML\CurrentControlSet\Services\disk and is specified with the following string:

```
TimeOutValue = 0x3C
```

where the number is the timeout value in seconds.

C.4 AutoPilot Installer Exit Codes

AutoPilot Installer sets an exit code to indicate whether an installation was successful or an error occurred. These exit codes allow AutoPilot Installer to be used in scripts with error handling. In unattended installations, AutoPilot Installer sets the following exit codes listed in Table 23.

Table 23 Unattended Installation Error Codes

Error Code	Hex	Description
0	0x00000000	No errors are reported.
2399141889	0x8F000001	An unsupported operating system detected.
2399141890	0x8F000002	The AutoPilot configuration file is not found.
2399141891	0x8F000003	Disabled adapters are detected in the system.
2399141892	0x8F000004	The selected driver is 64 bit and this system is 32 bit.
2399141893	0x8F000005	The selected driver is 32 bit and this system is 64 bit.
2399141894	0x8F000006	Installation activity is pending. AutoPilot Installer cannot run until it is resolved.
2399141895	0x8F000007	(GUI mode only) You canceled execution because you did not want to perform a software-first install.
2399141896	0x8F000008	No drivers are found.
2399141897	0x8F000009	One or more adapters failed diagnostics.
2399141904	0x8F000010	(GUI mode only) You chose to install drivers even though a recommended QFE or Service Pack was not installed.
2399141920	0x8F000020	(GUI mode only) You chose to stop installation because a recommended QFE or Service Pack was not installed.
2399141899	0x8F00000B	Unattended installation did not find any drivers of the type specified in the configuration file.
2399141900	0x8F00000C	A silent reboot was attempted, but according to the operating system a reboot is not possible.
2399141901	0x8F00000D	(GUI mode only) A driver package download was canceled.
2399141902	0x8F00000E	(Non-Enterprise) No adapters were found in the system.
2399141903	0x8F00000F	A required QFE or Service Pack was not detected on the system.
2399141836	0x8F000030	AutoPilot Installer was not invoked from an account with administrator-level privileges.
2391419952	0x8F000040	AutoPilot Installer has detected unsupported adapters on the system.
2399141968	0x8F000050	Unattended software-first installations were disallowed.
2399141984	0x8F000060	You cancelled APInstaller before any driver or utility installation occurred.
2399142000	0x8F000070	You cancelled APInstaller after driver or utility installation occurred.
2399142032	0x8F000090	APInstaller encountered an error while parsing the command line (Report file contains details).

C.5 AutoPilot Installer Installation Reports

During each installation, the AutoPilot Installer produces a report describing events that occurred during the installation. This report contains the following sections:

- The first section provides basic information including the time and date of the installation, the name of the machine on which the installation was performed, the version number of AutoPilot Installer, and the identification of the configuration file that was used.
- The second section provides an inventory of the Emulex adapters as they were before AutoPilot Installer performed any actions.
- The third section lists the tasks that AutoPilot Installer performs in the order that they are completed.
- The fourth section records the results of each task. When all driver installation tasks are completed, an updated adapter inventory is recorded.

NOTE If you cancel AutoPilot Installer, that fact is recorded along with the time you cancelled the installation. The contents of any error dialogs that are displayed are also recorded.

C.6 Command Script Example

Modify the configuration file to script the installation of a system's driver. The following example command script (batch file) assumes that you have made mandatory changes to the AutoPilot configuration file, as well as any desired optional changes. If your systems were set up with a service that supports remote execution, then you can create a command script to remotely update drivers for all of the systems on the storage network. If Microsoft's RCMD service was installed, a script similar to the following would run remote execution.

```
rcmd \\server1 g:\emulex\autopilot installer\fc\apinstall.exe
if errorlevel 1 goto server1ok
echo AutoPilot reported an error upgrading Server 1.
if not errorlevel 2147483650 goto unsupported
  echo Configuration file missing.
goto server1ok
 :unsupported
if not errorlevel 2147483649 goto older
echo Unsupported operating system detected.
 :older
if not errorlevel 2001 goto none
   echo The driver found is the same or older than the existing driver.
   goto server1ok
:none
if not errorlevel 1248 goto noreport
   echo No Emulex adapter found.
goto server1ok
:noreport
 if not errorlevel 110 goto nocfg
   echo Could not open installation report file.
 goto server1ok
:nocfg
 if not errorlevel 87 goto badcfg
   echo Invalid configuration file parameters.
   goto server1ok
:badcfg
   if not errorlevel 2 goto server1ok
   echo No appropriate driver found.
server1ok
rcmd \\server2 g:\autopilot\ApInstall
ConfigFileLocation=g:\autopilot\mysetup\apinstall.cfg
```

```
if errorlevel 1 goto server2ok
echo AutoPilot reported an error upgrading Server 2.
if not errorlevel 2147483650 goto unsupported
   echo Configuration file missing.
goto server2ok
 :unsupported
 if not errorlevel 2147483649 goto older
   echo Unsupported operating system detected.
 :older2
 if not errorlevel 2001 goto none2
   echo The driver found is the same or older than the existing driver.
   goto server2ok
:none2
 if not errorlevel 1248 goto noreport2
   echo No adapter found.
goto server2ok
:noreport
 if not errorlevel 110 goto nocfg2
   echo Could not open installation report file.
 goto server2ok
:nocfg2
 if not errorlevel 87 goto badcfg2
   echo Invalid configuration file parameters.
   goto server2ok
:badcfg2
   if not errorlevel 2 goto server2ok
   echo No appropriate driver found.
server2ok
```

Appendix D: License Notices

D.1 Secure Hash Algorithm (SHA-1) Notice


BROADCOM®

Emulex® Drivers for Linux for LightPulse® Adapters

User Guide

Version 11.4
September 6, 2017

DRVLin-LPe-UG114-100

Corporate Headquarters
San Jose, CA

Website
www.broadcom.com

Chapter 1: Introduction

1.1 Overview

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

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

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

This product supports the following Emulex LightPulse® HBAs and converged fabric adapters (CFAs):

- LPe12000-series adapters
- LPe16000-series adapters, including LPe16202/OCe15100 adapters
- LPe31000-series adapters
- LPe32000-series adapters

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

1.2 Abbreviations

AIC	adaptive interrupt coalescing
ANSI	American National Standards Institute
API	application programming interface
ASIC	application-specific integrated circuit
BIOS	basic input/output system
CLI	command line interface
CPU	central processing unit
CQ	completion queue
CQE	completion queue entry
DCBX	Data Center Bridging Capabilities Exchange
DIF	Data Integrity Field
DIMM	dual in-line memory module
DMA	direct memory access
ELS	Extended Link Service
EQ	event queue
FC	Fibre Channel
FCF	Fibre Channel over Ethernet Forwarder
FCFI	Fibre Channel Forwarder Indicator
FCoE	Fibre Channel over Ethernet
FCP	Fibre Channel Protocol
FSB	front side bus

FW	firmware
Gb/s	gigabits per second
GnuPG	GNU Privacy Guard
GRO	Generic Receive Offload
GUI	Graphical User Interface
HBA	host bus adapter
I/O	input/output
IOCB	input/output control block
IOV	I/O virtualization
IP	Internet Protocol
IRQ	interrupt request
KVM	kernel-based virtual machine
LAN	local area network
LBA	logical block address
LLC	logical link control
LOM	LAN on motherboard
LRO	large receive offload
LUN	logical unit number
MAC	Media Access Control
Mb/s	megabits per second
MDS	Multilayer Director Switch
MSI	message signaled interrupts
MSI-X	message signaled interrupts – extended
MTU	maximum transmission unit
NAA	Network Address Authority
NIC	network interface card (or controller)
NPIV	N_Port ID Virtualization
NUMA	non-uniform memory access
PCI	Peripheral Component Interconnect
PCIe	Peripheral Component Interconnect Express
PF	physical function
POST	power-on self-test
QoS	quality of service
RHEL	Red Hat Enterprise Linux
RPI	remote port indicator
RPM	resource package manager
RQ	receive queue
RSCN	registered state change notification
RSS	receive-side scaling
Rx	receive
SAN	storage area network
SCSI	Small Computer System Interface
SGE	Oracle (formerly Sun) grid engine

SLES	SUSE Linux Enterprise Server
SLI	Service Level Interface
SNMP	Simple Network Management Protocol
SR-IOV	single-root I/O virtualization
tar	tape archive
TCP	Transmission Control Protocol
TSO	TCP Segmentation Offload
Tx	transmit
UDP	User Datagram Protocol
UEFI	Unified Extensible Firmware Interface
VEB	virtual Ethernet bridging
VEPA	virtual Ethernet port aggregator
VF	virtual function
VGT	virtual guest tagging
vPort	virtual port
VLAN	virtual local area network
VM	virtual machine
VPD	vital product data
WWPN	World Wide Port Name
XRI	extensible resource indicator

Chapter 2: Installing and Uninstalling

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

http://www.broadcom.com/docs/elx-rpm-public-key

2.1 General Installation Requirements

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

Prior to driver installation, follow these general requirements:

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

2.2 Binary RPM FC and FCoE Driver Kit

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

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

  NOTE Use only officially released Linux distribution kernels. The binary RPM packages support only officially released Linux distribution kernels and do not support pre-release distribution kernels.
- An installation script, `elx_lpfc_install.sh`, that installs the FC and FCoE driver binary RPM that corresponds to the target system's architecture and kernel memory variant.
- A README file that provides a description of the kit structure, its contents, and distribution support scope.

2.2.1 Installing the Binary RPM FC and FCoE Driver Kit

NOTE You must uninstall any FC driver kits that are not part of this distribution. For example, you must uninstall any previous FC and FCoE driver kits before installing this driver kit. This installation fails if a previous version of the FC and FCoE driver kit is detected. For more information, see Section 2.2.2, Uninstalling the Binary RPM FC and FCoE Driver Kit.

To install the binary RPM FC and FCoE driver:

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

   ```
   tar xzf elx-lpfc-dd-<Linux distribution version>-<driver version>.tar.gz
   ```

3. Change to the directory into which the tarball was extracted:

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

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

   ```
   ./elx_lpfc_install.sh
   ```

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

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

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

2.2.2 Uninstalling the Binary RPM FC and FCoE Driver Kit

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

To uninstall the binary RPM FC and FCoE driver:

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

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

2.3 Ethernet Driver Kit for LPe16202/OCe15100 Adapters

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

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

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

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

2.3.1 Installing the Ethernet Driver Kit

NOTE Remove any previously installed Ethernet driver kits (that is, those that were not part of a distribution's kernel), before proceeding. For more information on uninstalling the driver, see Section 2.3.2, Uninstalling the Ethernet Driver Kit.

To install the Ethernet driver:

1. Download the appropriate driver kit from the Documents and Downloads area of http://www.broadcom.com.

2. Log in as *root* to a terminal, and unpack the tarball:

   ```
   tar xzf elx-lpnic-dd-<driver version>.tar.gz
   ```

3. Change to the directory that is extracted:

   ```
   cd elx-lpnic-dd-<driver version>/
   ```

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

   ```
   ./lpnic_install.sh
   ```

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

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

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

2.3.2 Uninstalling the Ethernet Driver Kit

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

To uninstall the Ethernet driver:

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

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

2.4 Booting from a Nonzero LUN Attached to an Emulex FC or FCoE Adapter

To configure SLES 11 SPx to boot from an FC-attached disk device other than a LUN 0 device, refer to the *Boot for the Fibre Channel Protocol User Guide* or the *Boot for LPe16202/OCe15100 Adapters User Guide*, both of which are available at the Documents and Downloads area of http://www.broadcom.com.

2.5 Emulex OneCommand Manager Application

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

Chapter 3: Configuration

3.1 FC and FCoE Driver Configuration

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

3.1.1 FC and FCoE Driver Parameters

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

For more information on driver parameters, see Section 3.1.1.1, Static FC and FCoE Driver Parameters, and Section 3.1.1.2, Dynamic FC and FCoE Driver Parameters.

3.1.1.1 Static FC and FCoE Driver Parameters

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

Table 1 Static FC and FCoE Driver Parameters

Parameter	Description	`sysfs` Visible	Applicable Adapters	Activation
`lpfc_ack0`	When enabled, ACK0 is used for Class 2. The enabled value is 1. The disabled value is 0 (default).	Yes	SLI-3 only	Driver reload
`lpfc_discovery_threads`	Specifies the maximum number of ELS commands that can be outstanding for a discovery. **NOTE** The `lpfc_discovery_threads` parameter defaults to a value of 64 for private loop topologies regardless of the configured value. If multiple ports are configured on the host, the value of 64 is used only for those ports that are connected in a private loop topology. The configured value is used for all other ports. The minimum value is 1. The maximum value is 64. The default value is 32.	No	All	Driver reload
`lpfc_enable_bg`	Enables the BlockGuard (T10-DIF) feature. The minimum value is 0 (default). The maximum value is 1.	Yes	All	Driver reload
`lpfc_enable_da_id`	When enabled, the FC and FCoE driver issues a `DA_ID CT` command to the fabric when vPorts log out of the fabric. The enabled value is 1. The disabled value is 0 (default).	No	All	Driver reload
`lpfc_enable_hba_heartbeat`	When enabled, the heartbeat logic in the FC and FCoE driver is able to detect whether the adapter is functional. If the heartbeat logic detects the adapter is not functional, the driver will shut down the adapter. The disabled value is 0 (default). The enabled value is 1.	Yes	All	Driver reload

Table 1 Static FC and FCoE Driver Parameters (Continued)

Parameter	Description	sysfs Visible	Applicable Adapters	Activation
lpfc_enable_hba_reset	When enabled, the FC and FCoE driver can pass resets to the adapter. This action is typically used for debugging purposes. The enabled value is 1 (default). The disabled value is 0.	Yes	All	Driver reload
lpfc_enable_npiv	When enabled, the FC and FCoE driver can use NPIV to create vPorts (if supported by the fabric). The enabled value is 1 (default). The disabled value is 0.	Yes	All	Driver reload
lpfc_enable_SmartSAN	Sets up FDMI support for SmartSAN. The disabled value is 0 (default). The enabled value is 1.	Yes	All	Driver reload
lpfc_EnableXLane	For Emulex LPe16000-series, LPe31000-series, LPe32000-series, and LPe16202/OCe15100 adapters, specifies whether the ExpressLane™ feature is enabled or disabled. The minimum value is 0 (default). The maximum value is 1.	Yes	SLI-4 only	Driver reload
lpfc_fcp_class	Specifies either FC Class 2 or 3 for FCP data transmission. For Class 2, the value is 2. For Class 3, the value is 3 (default).	Yes	All	Driver reload
lpfc_fcp_io_channel	For Emulex LPe16000-series, LPe31000-series, LPe32000-series, and LPe16202/OCe15100 adapters, defines the number of I/O channels supported by the driver. For more information, see Section 3.4, "NetworkManager: <warn> (bond0): slave eth1 shares a physical port with existing slave eth0"FC and FCoE Driver Performance Tuning. The default value is 4 I/O channels. The minimum value is 1 I/O channel. The maximum value is 16 I/O channels. For Emulex LPe12000-series adapters, this parameter is not applicable and has no effect.	Yes	SLI-4 only	Driver reload
lpfc_fdmi_on	Specifies if FDMI support is enabled or disabled. The enabled value is 1 (default). The disabled value is 0.	Yes	All	Driver reload
lpfc_hba_queue_depth	Limits the number of outstanding commands per HBA. The minimum value is 32. The maximum value is 8192 (also default).	Yes	All	Driver reload
lpfc_lun_queue_depth	Specifies the default maximum number of commands sent to a single logical unit (disk drive). The minimum value is 1. The maximum value is 512. The default value is 30. The value of lpfc_lun_queue_depth is dependent on the max_xri value received. The max_xri value might vary per adapter type.	Yes	All	Driver reload

Table 1 Static FC and FCoE Driver Parameters (Continued)

Parameter	Description	sysfs Visible	Applicable Adapters	Activation
lpfc_max_luns	Specifies the highest available LUN ID that is valid, per target. For example, a value of 19 means that LUN IDs from 0 to 19 are valid for the target. The SCSI layer scans each target until it reaches this specified LUN ID. The minimum value is 0. The maximum value is 65535. The default value is 255.	Yes	All	Driver reload
lpfc_max_scsicmpl_time	Uses command completion time to control queue depth. The units are in milliseconds. The minimum value is 0 (default). The maximum value is 6000.	Yes	All	Driver reload
lpfc_restrict_login	When enabled, restricts vPorts login to remote initiators. The enabled value is 1 (default). The disabled value is 0.	No	All	Driver reload
lpfc_scan_down	When enabled, selects the *scan down* method (scanning the AL_PA from high to low) to assign a SCSI ID. The enabled value is 1 (default). The disabled value is 0.	Yes	FC adapters only	Driver reload
lpfc_sg_seg_cnt	Controls the scatter/gather maximum segment count passed to the FC and FCoE driver. This variable is applicable per SCSI command. For Emulex LPe12000-series adapters, the minimum value is 64 (default), and the maximum value is 4096. For all other adapters, the minimum value is 64 (default), and the maximum value is 510.	Yes	All	Driver reload
lpfc_sli_mode	For Emulex LPe12000-series adapters, this parameter allows you to force the SLI mode requested by the adapter driver. The possible values are: ■ 0 = Auto-select (default). ■ 2 = SLI-2. ■ 3 = SLI-3. **NOTE** For all other adapters, this parameter is not applicable and has no effect.	No	SLI-3 only	Driver reload
lpfc_tgt_queue_depth	Limits the number of outstanding commands per target port. The minimum value is 10. The maximum value is 65535 (also default).	Yes	All	Driver reload
lpfc_use_msi	When enabled, determines whether the driver uses MSI or MSI-X. ■ 0 = MSI disabled; INTx mode is used (default for the FC RHEL 5.x driver). ■ 1 = MSI; allows a maximum of 32 interrupts. ■ 2 = MSI-X; allows a maximum of 2048 interrupts (default for FC RHEL 6.x/SLES 11 SPx drivers). **NOTE** The default is 2. This value reverts to 1 if the system does not support MSI-X. This value reverts to 0 if the system does not support MSI.	Yes	All	Driver reload

3.1.1.2 Dynamic FC and FCoE Driver Parameters

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

Table 2 Dynamic FC and FCoE Driver Parameters

Parameter	Description	`sysfs` Visible	Applicable Adapters	Activation
`lpfc_cr_count`	For Emulex LPe12000-series adapters, this parameter determines the value for I/O coalescing for `lpfc_cr_count` outstanding commands. The minimum value is 1 (default). The maximum value is 255. **NOTE** For all other adapters, this parameter is not applicable and has no effect.	Yes	SLI-3 only	Port reset
`lpfc_cr_delay`	For Emulex LPe12000-series adapters, this parameter determines the value for I/O coalescing for `lpfc_cr_delay` (milliseconds) outstanding commands. The minimum value is 0 (default). The maximum value is 63. **NOTE** For all other adapters, this parameter is not applicable and has no effect.	Yes	SLI-3 only	Port reset
`lpfc_devloss_tmo`	Specifies the number of seconds to hold an I/O error when a device disappears. The minimum value is 0. The maximum value is 255. The default value is 30.	Yes	All	Dynamic
`lpfc_enable_bbcr`	For Emulex LPe31000-series and LPe32000-series adapters, specifies if the Buffer to Buffer Credit Recovery feature is enabled or disabled. The minimum value is 0. The maximum value is 1 (default).	Yes	SLI-4 only	Port reset
`lpfc_enable_mds_diags`	For Emulex LPe16000-series, LPe31000-series, and LPe32000-series adapters, enables or disables the Multilayer Director Switch (MDS) Diagnostics feature. The minimum value is 0 (default). The maximum value is 1. **NOTE** The parameter should be disabled (set to 0) after diagnostics are complete.	Yes	SLI-4 only	Port reset
`lpfc_fcp_imax`	For Emulex LPe16000-series, LPe31000-series, LPe32000-series, and LPe16202/OCe15100 adapters, specifies the maximum number of fast-path FCP interrupts per second. The minimum value is 0 or OFF and 5000. The maximum value is 5,000,000. The default value is 150,000. **NOTE** For Emulex LPe12000 adapters, this parameter is not applicable and has no effect.	Yes	SLI-4 only	Dynamic

Table 2 Dynamic FC and FCoE Driver Parameters (Continued)

Parameter	Description	sysfs Visible	Applicable Adapters	Activation
lpfc_fcp_io_sched	For Emulex LPe16000-series, LPe31000-series, LPe32000-series, and LPe16202/OCe15100 adapters, determines which algorithm to use when scheduling an FCoE I/O to an I/O channel. For more information, see Section 3.4, "NetworkManager: <warn> (bond0): slave eth1 shares a physical port with existing slave eth0"FC and FCoE Driver Performance Tuning. The default value is 0, configuration by round-robin scheduling. A value of 1 sets configuration to CPU scheduling. For Emulex LPe12000-series adapters, this parameter is not applicable and has no effect. This parameter is only applicable for Emulex OneConnect adapters, LPe16000-series, LPe32000-series, LPe15000-series adapters and by default the value is configured to CPU scheduling.	Yes	SLI-4 only	Dynamic
lpfc_first_burst_size	Configures first burst size on targets that support first burst. The minimum value is 0 (default). The maximum value is 65536.	Yes	All	Dynamic
lpfc_link_speed	Specifies the FC link speed. The possible values are: ■ 0 = Auto-select (default). ■ 2 = 2Gb/s. ■ 4 = 4Gb/s. ■ 8 = 8Gb/s. ■ 16 = 16Gb/s. ■ 32 = 32Gb/s. **NOTE** For LPe32000-series adapters running 11.2 firmware or later, this parameter is ignored. **NOTE** Setting this option incorrectly can cause the adapter to fail to initialize. **NOTE** If you configure the link speed in a BIOS utility, the link speed 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.	Yes	FC adapters only	Port reset
lpfc_log_verbose	Specifies the log verbosity level of the messages posted by the driver. Extra activity logging (bit mask). The minimum value is 0x0 (default). The maximum value is 0xFFFFFFFF.	Yes	All	Dynamic
lpfc_throttle_log_cnt	Specifies the number of messages logged within lpfc_throttle_log_time. The minimum value is 1. The maximum value is 1000. The default value is 10.	Yes	All	Dynamic
lpfc_throttle_log_time	Specifies the time limit for lpfc_throttle_log_cnt. The minimum value is 1. The maximum value is 60. The default value is 1.	Yes	All	Dynamic

Table 2 Dynamic FC and FCoE Driver Parameters (Continued)

Parameter	Description	sysfs Visible	Applicable Adapters	Activation
lpfc_topology	For FC adapters, this parameter sets the link topology. The possible values are: ■ 0x0 = Loop first; if loop fails, then point-to-point (default). ■ 0x2 = Point-to-point only. ■ 0x4 = Loop only. ■ 0x6 = Point-to-point first; if point-to-point fails, then loop. **NOTE** For the LPe16202/OCe15100 adapter in FCoE mode, this parameter is not applicable and has no effect.	Yes	FC adapters only	Port reset
lpfc_use_adisc	When enabled, an ADISC is sent instead of a PLOGI for device discovery or RSCN. The enabled value is 1. The default value is 0 (disabled).	Yes	All	Dynamic
lpfc_XLanePriority	For Emulex LPe16000-series, LPe31000-series, LPe32000-series, and LPe16202/OCe15100 adapters, when enabled, this parameter defines the frame priority to be used for the ExpressLane feature. The minimum value is 0 (default). The maximum value is 0x7f.	Yes		Dynamic

3.1.2 Configuring FC and FCoE Driver Parameters

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

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

  NOTE Not all parameters visible in the sysfs directory can be modified; some are read-only.
- The Emulex OneCommand Manager application (refer to the *OneCommand Manager Application for LightPulse Adapters User Guide* for more information).

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

3.1.2.1 Temporary Configuration with modprobe

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

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

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

```
modprobe lpfc lpfc_lun_queue_depth=20
```

3.1.2.2 Persistent Configuration with the lpfc.conf file

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

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

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

If driver parameters are added to the `lpfc.conf` file, the FC and FCoE driver must be reloaded for the parameters to take effect. Also, a new ramdisk image is required if you want the changes to take effect in the next boot. For information on creating a new ramdisk, See Section 3.1.3, Creating a New Ramdisk Image.

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

3.1.2.3 Configure Parameters with a Read/Write to sysfs

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

NOTE

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

For example:

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

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

Viewing Parameters with sysfs

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

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

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

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

An example output follows:

```
-r--r--r--  1 root root 4096 Feb 28 17:03
/sys/class/scsi_host/host7/lpfc_ack0
```

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

Temporary Configuration Parameters with sysfs

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

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

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

To modify the `lpfc_log_verbose` value to 0xffffffff:

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

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

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

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

3.1.3 Creating a New Ramdisk Image

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

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

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

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

3.1.4 Dynamically Recognizing LUNs and Targets (Using scan)

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

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

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

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

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

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

3.1.5 Persistent Naming

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

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

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

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

The sample output is:

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

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

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

3.1.5.2 Configuring the System to Boot Using Persistent Names

For SLES 11 SPx and SLES 12

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

To use a persistent name for a boot device:

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

   ```
   kernel /boot/vmlinuz root=/dev/sda2 vga=0x314
   ```

2. Find the persistent name for the root partition (following `root=` on the kernel line) by using the instructions in Section 3.1.5.1, Using udev to Discover Logical to Physical Mappings for sd Devices.

3. In the same file, `/boot/grub/menu.lst`, replace the text after `root=` with the partition's persistent name. For example:

```
kernel /boot/vmlinuz root=/dev/disk/by-id/scsi-32000000c5005d6e6-part2
vga=0x314
```

4. Change any mounts listed in `/etc/fstab` that refer to this root partition by either its `/dev/sd` name or a file system label to use the persistent name as well.

For RHEL 6.x and RHEL 7.x

To use a persistent name for a boot device:

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

```
kernel /boot/vmlinuz -<kernel version> ro root=/dev/sda2
```

2. Find the persistent name for the root partition (following `root=` on the kernel line) by using the instructions in Section 3.1.5.1, Using udev to Discover Logical to Physical Mappings for sd Devices.
3. In the same file, `/boot/grub/menu.lst`, replace the text after `root=` with the partition's persistent name. For example:

```
kernel /boot/vmlinuz -<kernel version> ro
root=/dev/disk/by-id/scsi-32000000c5005d6e6-part2
```

4. Change any mounts listed in `/etc/fstab` that refer to this root partition by either its `/dev/sd` name or a file system label to use the persistent name as well.

3.1.5.3 Using udev with st Devices

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

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

FC Tape Device Examples

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

```
scsi_id [options]
```

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

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

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

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

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

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

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

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

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

FC-SCSI Tape Device Example

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

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

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

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

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

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

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

Creating the udev Persistent Name for SCSI Tape Device

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

The rule for the FC tape device is

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

The rule for the FC-SCSI tape device is:

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

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

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

3.1.5.4 Persistent Naming References

See the following references for more information on persistent naming:

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

3.1.6 ExpressLane Support

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

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

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

- The ExpressLane feature is supported only on Emulex LPe16000-series, LPe31000-series, LPe32000-series, and LPe16202/OCe15100 adapters.

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

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

3.2 vPort Configuration

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

3.2.1 vPort Configuration Prerequisites

Before configuring vPorts, consider the following points:

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

  NOTE Before performing a firmware update, driver installation is required. For more information on installing the driver, see Chapter 2, Installing and Uninstalling.
- Loop devices and NPIV are not supported on the same port simultaneously. If you are running a loop topology and you create a vPort, the vPort link state is offline.
- vPorts do not persist across system reboots.

3.2.2 Creating, Deleting, and Displaying vPorts

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

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

3.2.2.1 Creating vPorts Using the mkvport.sh Script

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

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

You must supply the physical port's host number, WWPN, and WWNN when using the `mkvport.sh` script. For example, to create a vPort with port name of 10000000c94ac63a and a node name of 20010000c94ac63a on the physical port with `scsi_host` name *host7*, type:

```
./mkvport.sh host7 10000000c94ac63a 20010000c94ac63a
```

This script fails if the vPort is not created.

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

3.2.2.2 Deleting vPorts Using the rmvport.sh Script

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

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

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

Or

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

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

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

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

3.2.2.3 Displaying vPorts Using the lsvport.sh Script

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

For example:

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

In reference to the previous example:

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

3.2.3 vPort sysfs

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

3.2.3.1 vPort sysfs Tree

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

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

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

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

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

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

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

In this example:

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

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

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

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

3.2.3.2 vPort sysfs Entries

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

Table 3 vPort sysfs Entries

vPort `sysfs` Entries	Type	Range/Input	Location and Description
`lpfc_peer_port_login`	Read/Write	0=Off (default) 1 = On	`/sys/class/scsi_host/hostX/lpfc_peer_port_login` This entry sets the port behavior when discovering targets in the SAN. The default behavior (value = 0) will log in only to `N_Ports` that are physically located on a different port. The port will still attempt to log in to targets on all other ports (including the other port in a dual-port adapter). If this parameter is turned on (value = 1), the port attempts to log in to all `N_Ports`, even if they are physically located on the same port. **NOTE** This parameter was created to reduce the amount of hardware resources (for example, RPIs) that the driver requires. In a configuration where there are many vPorts on one physical port, this capability greatly reduces the number of RPIs that the driver uses.
`lpfc_restrict_login`	Read/Write	0 = Off 1 = On (default)	`/sys/class/scsi_host/hostX/lpfc_restrict_login (vPorts only)` This entry sets the vPort behavior when discovering targets in the SAN. The default behavior (value = 1) prevents the vPort from logging into other initiator ports in the SAN. It also rejects logins from other ports in the SAN, because it assumes that all ports that send a PLOGI are initiators. If this `sysfs` entry is turned off, the driver attempts to log in to every port that it can access in the SAN, and accepts logins from all ports. **NOTE** This parameter was created to reduce the amount of hardware resources (for example, RPIs) that the driver requires. In a SAN where other initiators exist, this capability greatly reduces the number of RPIs that the driver uses.
`max_npiv_vports`	Read-only	integers	`/sys/class/fc_host/hostX/max_npiv_vports` This entry displays the maximum number of vPorts that are supported by the underlying `fc_host` hardware. This `sysfs` entry exists only if the `vport_create` and `vport_delete` `sysfs` entries exist. If an `fc_host` does not support NPIV, this `sysfs` entry might not exist. Use this `sysfs` entry with the `npiv_vports_inuse` entry to determine whether the maximum number of vPorts have been created on this `fc_host`.

`node_name`	Read-only	16-byte hexadecimal value	For the FC RHEL 5.x driver:
			`/sys/class/fc_host/hostX/node_name`
			For FC RHEL 6.x/SLES 11 SPx drivers:
			`/sys/class/fc_host/hostX/node_name`
			`/sys/class/fc_vports/vport-X:0-Z/node_name`
			These entries display the physical or vPort node name.
			You assign this value when the vPort is created, and it is transmitted to the fabric upon fabric login.

Table 3 vPort sysfs Entries (Continued)

vPort `sysfs` Entries	Type	Range/Input	Location and Description
`npiv_vports_inuse`	Read-only	integers	`/sys/class/fc_host/hostX/npiv_vports_inuse` This entry displays the number of vPorts that were created on this `fc_host`. This `sysfs` entry exists only if the `vport_create` and `vport_delete` `sysfs` entries exist. If an `fc_host` does not support NPIV, this `sysfs` entry may not exist. Use this `sysfs` entry with `max_npiv_vports` to determine whether the maximum number of vPorts have been created on this `fc_host`.

`port_name`	Read-only	16-byte hexadecimal value	`/sys/class/fc_host/hostX/port_name`
			`/sys/class/fc_vports/vport-X:0-Z/port_name`
			This entry displays the physical or vPort port name.
			You assign this value when the vPort is created, and it is transmitted to the fabric upon fabric login.

`vport_create`	Write-only	WWPN; WWNN	`/sys/class/fc_host/hostX/vport_create` This entry creates a vPort on the physical port on which hostX is located. The new vPort will have a WWPN and WWNN present on the fabric based on the WWPN and WWNN that are entered with this `sysfs` entry. This entry returns a 0 if the vPort creation was successful. A nonzero value indicates that the vPort was not created. If an `fc_host` does not support NPIV, then this `sysfs` entry may not exist. **NOTE** It is possible for the vPort creation to succeed but for the vPort to be in a failed or inoperative state. Use the new `sysfs` tree created by the new vPort to check the state of the new vPort.
`vport_delete`	Write-only	WWPN; WWNN	`/sys/class/fc_host/hostX/vport_delete` This entry deletes a vPort on the physical port on which hostX is located. The vPort matching the WWPN and WWNN is immediately deleted. This entry returns a 0 if the vPort deletion was successful. A nonzero value indicates that the vPort was not deleted. If an `fc_host` does not support NPIV, then this `sysfs` entry may not exist. **NOTE** This entry deletes the vPort even if there are mounted file systems being accessed through this vPort, or if there are open files on it.

3.2.4 Monitoring vPorts with fc_vport

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

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

```
power uevent

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

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

In this example:

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

3.2.5 vPort Configuration Limits

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

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

- Before the vPort is deleted or the driver is unloaded, I/O devices accessed through a vPort must be stopped and file systems must be unmounted.
- For all LPe12000-series, LPe16000-series (including LPe16202/OCe15100), LPe31000-series, and LPe32000-series adapters,the maximum number of vPorts configurable on a physical port is 255.
- The maximum number of LUNs supported on each driver port is 256.
- The maximum number of targets supported for each driver port is 255.
- The maximum number of driver ports in one zone is 64. This limit is based on the system's ability to recover from link events within the time constraints of the default timers.

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

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

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

3.3 Ethernet Driver Configuration for LPe16202/OCe15100 Adapters

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

3.3.1 Ethernet Driver Configuration Parameters

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

Table 4 Ethernet Driver Configuration Parameters

Parameter	Description
`num_vfs`	In systems supporting SR-IOV, when IOV is enabled, this parameter indicates the number of VFs to be enabled per PF. For configuring SR-IOV, see Section 3.3.7, SR-IOV Configuration. The default value is 0 (SR-IOV is not enabled). The possible values are 0 to 63. **NOTE** This parameter is obsolete. Use the `sysfs` method to enable or disable VFs.
`rss_on_mc`	Enables receive-side scaling (RSS) on multichannel functions that have the capability. The default value is 0 (disabled). The enabled value is 1. **NOTE** Currently, this parameter is ignored by the driver. The driver enables RSS on multichannel functions that have the capability by default.
`rx_frag_size`	The size of fragments used to DMA received data. The possible values are 2048 (default), 4096 and 8192.

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

```
options lpnic rx_frag_size=4096
```

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

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

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

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

NOTE

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

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

- The maximum number of channels supported by NIC function:
  `cat /sys/class/net/<ethx>/max_qs`
- The current number of channels configured:
  `cat /sys/class/net/<ethx>/num_qs`

- To configure the *N* number of channels requested for the NIC function:
  `echo N > /sys/class/net<ethx>/num_qs`

3.3.3 Support for Ethtool set-dump Command

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

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

NOTE

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

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

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

3.3.4 Transmit/Receive Queue Counts

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

For example:

`ethtool -l <ethx>`

Output

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

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

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

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

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

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

- Indirection table entries
- RSS hash key

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

- To show the currently configured RSS hash key:
  `cat /sys/class/net/<ethx>/rss_hashkey`
- To configure the new RSS hash key:
  `echo -e "aa:bb:cc...." > /sys/class/net/<ethx>/rss_hashkey`

3.3.6 Support for Showing Onboard Die Temperature

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

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

```
cat /sysclass/hwmon/<hwmonx>/device/temp1_input
```

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

```
cat /sys/class/net/<ethx>/device/temp1_input
```

3.3.7 SR-IOV Configuration

3.3.7.1 Introduction

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

- RHEL 6.7 – 64-bit KVM
- RHEL 6.8 – 64-bit KVM
- RHEL 6.9 – 64-bit KVM
- RHEL 7.1 – 64-bit KVM
- RHEL 7.2 – 64-bit KVM
- RHEL 7.3 – 64-bit KVM
- RHEL 7.4 – 64-bit KVM
- –
- SLES 11 SP3 – 64-bit Xen and KVM
- SLES 11 SP4 – 64-bit Xen and KVM
- SLES 12 SP2 – 64-bit Xen and KVM
- SLES 12 SP3 – 64-bit Xen and KVM
- SLES and RHEL supported with Xen

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

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

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

Known issues include:

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

3.3.7.2 Setting Up SR-IOV

Prerequisites (except FLEX10)

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

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

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

- An Emulex LPe16202/OCe15100 adapter in FCoE+NIC mode.
- If a KVM hypervisor is installed, it must contain the qemu-kvm packages.

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

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

You can use the `PXESelect` utility, the Emulex OneCommand Manager application, or the `UEFI (HII)` utility to set up SR-IOV.

To enable and configure SR-IOV using either the `UEFIBoot` utility or the `PXESelect` utility, refer to the *Boot for LPe16202/OCe15100 Adapters User Guide*.

1. Install the required Linux operating system that serves as a hypervisor.
2. For legacy boot, perform the following step:

   Update the `/boot/grub/menu.lst` file to include the following command line load parameter for the Linux kernel:

   ```
   intel_iommu=on
   ```

   For `UEFI boot`, perform the following step:

   Update the `/etc/default/grub` file to include the following command line load parameter for the Linux kernel:

   ```
   intel_iommu=on
   ```

3. Install the appropriate version of the Emulex Ethernet driver that supports SR-IOV for the operating system version that you are using.
4. Reboot the server for the new changes to become operational.
5. Use the `lspci -vvv` output command to check if SR-IOV is properly enabled.

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

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

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

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

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

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

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

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

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

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

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

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

   The total number of VFs can be distributed among available ports as required, but each port has a maximum of 64 VFs. Table 5 on page 34 lists the total number of VF counts that are supported for various adapter configurations.

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

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

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

9. The total VF count that is supported for a two-port LPe16002/OCe15100 adapter in NIC plus FCoE mode (two NIC PFs and two FCoE PFs) is 128. The VF count is symmetrically distributed across the ports, which means that the VF count per NIC port is 64. Detach VFs from the host before adding them to the guest.

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

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

   ```
   lspci | grep Emulex 16:00.0 Ethernet controller: Emulex Corporation OneConnect NIC
   (Lancer) (rev 30)<-- LOM Port 0 (Function 0) 16:00.1 Ethernet controller: Emulex Corporation
   OneConnect NIC (Lancer) (rev 30)<-- LOM Port 1 (Function 1) 16:04.0 Ethernet controller: Emulex
   Corporation OneConnect NIC (Lancer) (rev 30)<-- VF LOM Port 0 (Function 0) 16:08.0 Ethernet
   controller: Emulex Corporation OneConnect NIC (Lancer) (rev 30)<-- VF LOM Port 1 (Function 1)
   ```

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

Detach desired VFs:

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

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

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

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

List Emulex PCI device by device code:

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

10. Use the Virtual Machine Manager GUI to attach the VF (step 5 on page 30) to the guest operating system by using the add physical device option.

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

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

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

12. The Ethernet driver automatically loads with the out-of-box driver, and creates the network interfaces. Use the `ifconfig` command to verify that the interface is created.
13. After configuring the network interfaces with proper IP addresses, you can send and receive network traffic from the VM. Refer to the documentation for the host and guest operating systems for information on network configuration.

3.3.7.3 Assigning VFs to a VM on the SLES Operating System

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

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

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

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

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

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

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

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

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

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

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

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

3.3.7.4 Link State Reporting with SR-IOV

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

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

Table 5 Default Link Status of the PFs and VFs

Function	Physical Link Is Not Connected
	LPe16000
Physical function (PF)	Down
Virtual function (VF)	Down

The default behavior of VFs can be changed using the `ip link` command in distributions of Linux that support iproute version 3.11.0 or later. For details on configuring the link state for VFs, see Section 3.3.7.6, Link State Configuration.

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

3.3.7.5 Configuring VFs

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

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

MAC Address Configuration

The LPe16202/OCe15100 adapter in NIC+FCoE mode is shipped with factory-configured MAC addresses for the network interfaces corresponding to the PFs. The driver generates random MAC addresses for the network interfaces corresponding to the VFs based on the factory-configured MAC address. Other MAC addresses can be assigned for the interfaces corresponding to the VF using IP utility commands in the hypervisor.

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

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

where:

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

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

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

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

- If the VF interface is configured using `ifcfg` script:
  a. Update the `ifcfg` script for the VF, with the MAC address configured in the hypervisor in the earlier step. For example, update the `HWADDR` line in the VF `ifcfg-eth<X>` script.
     `HWADDR=00:16:88:AA:BB:AA`
  b. Restart the network service.
     `# service network restart`

Or

- If the VF interface is not configured using `ifcfg` script:
  Update the VF MAC address using the `ifconfig` or `ip link` command.`# ifconfig eth0 hw ether 00:16:88:AA:BB:AA`

VLAN Configuration

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

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

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

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

where:

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

Example:

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

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

Transmission Rate Configuration

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

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

where:

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

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

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

3.3.7.6 Link State Configuration

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

This section provides link state configuration for a VF.

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

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

where:

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

3.3.7.7 Spoof Check Configuration

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

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

where:

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

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

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

3.3.7.8 Viewing VF Properties

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

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

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

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

```
# ip link show eth0
```

Expected example output:

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

3.3.8 Bonding Considerations

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

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

```
# cat /sys/class/net/eth0/phys_port_id
0100000000000000000000383838314643
```

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

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

3.4 FC and FCoE Driver Performance Tuning

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

3.4.1 Overview

The following configurable parameters can enhance performance:

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

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

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

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

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

3.4.1.1 lpfc_fcp_io_channel

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

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

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

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

3.4.1.2 lpfc_fcp_io_sched

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

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

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

NOTE This parameter is only applicable for Emulex OneConnect adapters, LPe16000-series, LPe32000-series, LPe15000-series adapters and by default the value is configured to CPU scheduling.

3.4.1.3 lpfc_fcp_imax

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

Considerations

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

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

3.4.1.4 lpfc_vector_map.sh

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

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

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

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

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

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

Driver Mode

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

HBA Mode

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

3.5 Network Performance Tuning

This section describes optimizing network performance.

3.5.1 Memory Bandwidth Considerations

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

3.5.1.1 Enabling Optimal Bandwidth Options

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

3.5.1.2 Populating DIMM Slots

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

3.5.1.3 Disabling Memory Mirroring

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

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

Nearly any desktop or low-end server has enough memory bandwidth for Emulex LPe16202/OCe15100 adapters in NIC mode to support DMA at 20Gb/s of data (10Gb/s read, 10Gb/s write). However, most of the memory demands come from the processor accessing the data for either packet copies in the non-offloaded networking stack or application accesses. All processor memory accesses use the FSB. The clock speed of this bus is critical for allowing efficient memory bandwidth. A system with a faster processor FSB clock speed performs better than a system with a slower FSB clock speed.

3.5.2 Network Memory Limits

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

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

```
echo 4096 87380 4194304 > /proc/sys/net/ipv4/tcp_rmem
echo 4096 16384 4194304 > /proc/sys/net/ipv4/tcp_wmem
echo 64000000 > /proc/sys/net/core/rmem_default
echo 64000000 > /proc/sys/net/core/rmem_max
echo 32000000 > /proc/sys/net/core/wmem_default
echo 32000000 > /proc/sys/net/core/wmem_max
```

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

These `tcp_rmem` and `tcp_wmem` values are also the default values in recent RHEL 5 distributions. If your application requires best throughput with very small number of connections (less than four), it may help to increase the `tcp_rmem` and `tcp_wmem` to much larger values:

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

3.5.3 TCP Segmentation Offload (TSO)

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

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

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

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

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

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

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

3.5.4 Flow Control

LPe16202/OCe15100 adapters in NIC mode support IEEE 802.3x standard flow control, which uses control packets to temporarily pause the transmission of packets between two endpoints. These control messages are point-to-point and are not forwarded by switches or routers. The adapter can respond to flow control packets by temporarily pausing transmits. The adapter can send flow control pause packets when the transmitter is overwhelming the system's receive bandwidth.

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

- The adapter is installed in 4x PCIe slot or an underpowered server system.

  If the PCIe bus does not provide 10Gb/s of throughput due to chipset limitations or the bus width, the adapter cannot maintain 10Gb/s of incoming receive data. It starts dropping packets quickly. In this situation, it may be beneficial to enable receive flow control in the adapter, and enable flow control in the attached switch for all devices. This action helps to slow down the transmitters.
- The adapter transmits to 1Gb/s devices, especially when using a non-TCP protocol.

  If the adapter transmits to a 10Gb/s switch with attached 1Gb/s clients, the adapter may overwhelm the switch. The switch is then forced to start dropping packets because, although it may receive a 10Gb/s stream, the client can only sink a 1Gb/s stream. In this situation, it may be beneficial to enable transmit flow control in the adapter and enable flow control for the 10Gb/s switch port.

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

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

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

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

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

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

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

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

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

3.5.5 RX Frame Coalescing

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

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

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

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

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

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

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

3.5.6 Maximum Transmission Unit (MTU)

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

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

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

3.5.7 Interrupt Coalescing

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

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

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

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

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

For transmit interrupts, AIC is not supported.

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

3.5.7.2 Setting the Interrupt Delay Duration to a Constant Value

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

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

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

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

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

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

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

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

3.5.8 Receive-Side Scaling (RSS)

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

3.5.9 Analyzing Performance Issues

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

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

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

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

Use the `-S` option of `ethtool` to see all statistics counters maintained by the Ethernet and driver. Excessive drop or error counters are an indication of a bad link or defective hardware. See Table 1, Ethtool -S Option Statistics, and Table 2, Transmit/Receive Queue Statistics.

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

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

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

Chapter 4: Troubleshooting

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

4.1 Hardware Situations and Resolutions

Table 6 lists the hardware situations and their resolutions.

Table 6 Hardware Situations and Their Resolutions

Situation	Resolution
An unapproved optical transceiver is used to connect the adapter.	The system log can generate one or more of these events if an unapproved optical transceiver is detected: ■ Unsupported module ■ Optics faulted/incorrectly installed/not installed ■ Incompatible optics ■ Unknown port SFB status To resolve this issue, do the following: 1. Power off the system. 2. Replace the unapproved optical transceiver with an approved on. 3. Power on the system. **NOTE** For more information on replacing an optical transceiver, refer to the specific adapter hardware guide available at the Documents and Downloads area of http://www.broadcom.com.

4.2 FC and FCoE Driver Situations and Their Resolutions

Table 7 lists the FC and FCoE driver situations and their resolutions.

Table 7 FC and FCoE Driver Situations and Their Resolutions

Situation	Resolution
FC link fails to come up.	If an FC link fails to come up, verify that an 8Gb/s or 16Gb/s adapter is not attempting to connect to a 1Gb/s device. Only 2Gb/s, 4Gb/s, and 8Gb/s devices are supported on 8Gb/s adapters. Only 2Gb/s, 4Gb/s, 8Gb/s, and 16Gb/s devices are supported on 16Gb/s adapters.

`Authentication is enabled but authentication service is not running`error message.	If you see this message in `/var/log/messages` and the adapter is in an error state, the `fcauthd` daemon probably is not running. To determine whether `fcauthd` is running, run:
	`/etc/init.d/fcauthd status.`
	To start fcauthd, run
	`/etc/init.d/fcauthd start.`

If a SAN configuration has 256 targets mapped by the FC and FCoE driver, any additional added targets do not get a target ID mapping by the driver and cause target discovery to fail.	Removing targets or reinitializing the link does not solve this issue. Unload and reload the driver to reset available target IDs. Ensure that the SAN configuration is correct prior to rebooting the driver. This action clears the driver consistent binding table and frees target IDs for new target nodes.

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

Situation	Resolution
`rmmod` fails to unload FC and FCoE driver module due to `ERROR: Module lpfc is in use`.	This message can appear when you attempt to remove the driver and a Logical Volume Group is dependent on the driver. Driver unload is not supported in SAN booted environments unless the inbox driver supports SAN boot. Otherwise, use these steps to resolve this situation: 1. Make the Logical Volume Group unavailable. Type: `lvchange -a n xxxxxxx` The *xxxxxxx* parameter is the Volume Group Name. 2. Stop the Emulex OneCommand Manager application. 3. Stop Device Mapper.
`rmmod` hangs and module reference count is 0.	Due to a small race condition in the kernel, it is possible for an `rmmod` command to hang. Issue the `rmmod -w` command. If this does not help, reboot the computer.

`rmmod` fails to unload driver due to device or resource busy.	This message occurs when you attempt to remove the driver without first stopping the Emulex OneCommand Manager application or the `fcauthd` daemon when the Emulex OneCommand Manager application is installed and running, or when FC disks connected to an Emulex LightPulse adapter are mounted. To resolve this situation:
	1. Stop the Emulex OneCommand Manager application before attempting to unload the driver. The script is located in the `/usr/sbin/ocmanager` directory. Type:
	`./stop_ocmanager`
	2. Unmount any disks connected to the adapter.
	3. Unload the driver. Type:
	`rmmod lpfc`

An `lspci` shows recent Emulex adapters as unknown.	This situation occurs because of the delay of getting new product IDs into the Linux development cycle. There is no resolution at this time.
Slow targets or extended link faults on the storage side might result in storage being marked offline by the mid-level and remaining offline (not recovered) when the link faults are corrected.	This version of the driver should eliminate this issue. However, if you experience offline device issues, increase the SCSI command timeout to a value greater than or equal to 60 seconds. Emulex also provides a script that addresses this issue. To access the `lun_change_state.sh script`, go to the Documents and Downloads area of http://www.broadcom.com.
Under certain conditions of an I/O load, some targets cannot complete an I/O issued by a Linux initiator within the default timeout of 30 seconds given by the SCSI mid-level.	If the situation is not corrected, the initiator-to-target condition deteriorates into abort/recovery storms, leading to I/O failures in the block layer. These types of failures are preceded by a SCSI I/O error of hex 6000000. Emulex provides a script that addresses this issue. To access the `set_target_timeout.sh script`, go to the Documents and Downloads area of http://www.broadcom.com.
The FC and FCoE driver fails to recognize an adapter and logs `unknown IOCB` messages in the system log during driver load.	The adapter is running outdated firmware. Install the latest firmware on the adapter. **NOTE** Before performing a firmware update, driver installation is required. For more information on installing the driver, see Section 2.2.1, Installing the Binary RPM FC and FCoE Driver Kit.
Loading the FC and FCoE driver on SLES 11 SPx reports `unsupported module, tainting kernel`in system log.	This message is logged by the kernel whenever a module that is not shipped with the kernel is loaded. Ignore this message.
The system panics when it is booted with a failed adapter installed.	Remove the failed adapter and reboot the system.

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

Situation	Resolution
Unloading the FC and FCoE driver on SLES 11 SPx might cause a message to be logged in the system log, such as the following: `umount: /dev/disk/bypath/pci-0000:02:04.0-scsi-0:0:1:0: not mounted`	These messages are normal output from the SLES 11 SPx hotplug scripts. Ignore them.
Driver installation fails.	The `lpfc-install` script fails to install the driver. The install script may fail for the following reasons: ■ A previous version of the driver is installed. Run the `lpfc-install --uninstall` script and then try to install the driver. ■ The current driver is already installed. ■ Run a supported RHEL or SLES kernel.
`No module lpfc found for kernel KERNELVERSION`RPM error message when upgrading the kernel. A recently upgraded kernel cannot find the ramdisk. After upgrading the kernel, the kernel cannot find the ramdisk, which halts or panics the system. The driver is not loaded after a system reboot after upgrading the kernel.	These three situations can be resolved by upgrading the kernel. There are two ways to install the driver into an upgraded kernel. The method you use depends on whether you are updating the driver. ■ Upgrade the kernel using the same version of the driver. ■ Upgrade the kernel using a new version of the driver. For more information on upgrading the kernel, see Chapter 2, Installing and Uninstalling.
Driver uninstall fails.	The `lpfc-install --uninstall` script fails with an error. Try the following solutions: ■ Uninstall the Emulex OneCommand Manager application; refer to the *OneCommand Manager Application for LightPulse Adapters User Guide* for instructions. ■ Unmount all FC disk drives. ■ Unload the `lpfcdfc` and FC and FCoE driver. ■ Use `rpm -e lpfcdriver` and `-e ocmanager` and uninstall the new kits.

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

Situation	Resolution
`lpfc-install` script exit code.	The `lpfc-install` script contains exit codes that can be useful in diagnosing installation issues. Refer to the `lpfc-install` script for a complete listing of codes and definitions.

The Emulex driver for Linux does not load in ramdisk for a custom-built kernel.	Custom built kernels are not supported by Emulex. However, the Emulex install script attempts to install the driver into a ramdisk that follows the naming scheme used by RHEL or SLES kernels.
	■ The RHEL naming scheme for ramdisk images is:
	`/boot/initrd-`*`KERNELVERSION`*`.img.`
	■ The SLES naming scheme for ramdisk images is:
	`/boot/initrd.`
	If a custom built kernel has a ramdisk image that does not follow the appropriate naming scheme, change the name of the image using the following procedure:
	1. Change the name of the ramdisk image to match the SLES naming scheme.
	2. Update any file links to the ramdisk image.
	3. Edit the boot loader configuration file (for example, `/etc/lilo.conf`, `/etc/yaboot.conf`, `/boot/grub/grub.conf`, `/boot/grub/menu.lst`), find any references to the old ramdisk image name, and replace them with the new name.
	4. Reboot the system to verify the changes.
	5. Install the Emulex Linux driver kit.

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

4.3 Ethernet Driver Situations and Their Resolutions

The following table lists the Ethernet driver situations and their resolutions. This section applies to LPe16202/OCe15100 adapters in NIC mode only.

Table 8 Ethernet Driver Situations and Their Resolutions

Situation	Resolution
The ethtool configuration settings are not restored after system reboot.	The ethtool settings are not designed to persist across reboot. For persistence, invoke configuration commands from a boot script that is executed at system start, such as `/etc/rc.local`.
The Ethernet driver works but the transmit and receive data rates are not near a 10Gb/s line rate.	There could be several reasons for poor performance. For best performance practices, see Section 3.5, Network Performance Tuning.

When `MILI` and `SNMP` daemons start, they trigger warning messages within `SELinux` for certain operations.	This issue is known and no solution is available. However, to avoid `SELinux` warning messages, you can disable `SELinux`. To disable `SELinux`, open a terminal and enter the following command at the prompt:
	`echo 0 > /selinux/enforce`
	To enable `SELinux`, use the following command:
	`echo 1 > /selinux/enforce`
	In addition, see Section 3.5.9, Analyzing Performance Issues.

4.4 Log Messages

4.4.1 FC and FCoE Driver Log Messages

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

4.4.1.1 Retrieving FC and FCoE Driver Log Messages

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

An example of an LPFC message:

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

In this example:

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

NOTE

- If `Data:` is present in a log message, any information following `Data:` is intended only for Broadcom® Technical Support or Engineering use.
- If an error message instructs you to perform a firmware update, ensure that the driver is installed first. For more information on installing the driver, see Chapter 2, Installing and Uninstalling.

4.4.1.2 LPFC Error Log Messages and their Descriptions

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

Table 9 LPFC Error Log Messages and their Descriptions

`0111: Dropping received ELS cmd` The driver decided to drop an `ELS Response` ring entry. Data: (1) `ulpStatus` (2) `ulpWord`[4] (3) `ulpTimeout` Severity: Error Log: Always Action: This error could indicate a software driver or firmware issue. If this issue persists, report these errors to Broadcom Technical Support.
`0113: An FLOGI ELS command <elsCmd> was received from DID <did> in Loop Mode` While in Loop Mode an unknown or unsupported `ELS` command was received. Data: None Severity: Error Log: Always Action: Check device DID.
`0115: Unknown ELS command <elsCmd> received from NPORT <did>` Received an unsupported ELS command from a remote `N_Port`. Data: None Severity: Error Log: Always Action: Check the remote `N_Port` for a potential issue.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`0125: FDISC Failed (x%x). Fabric out of resources` The fabric rejected an FDISC because the switch cannot support additional virtual ports. Data: `lsRjtError` Severity: Error Log: Always Action: Reconfigure the switch to support more NPIV logins. If this issue persists, contact Broadcom Technical Support.
`0126: FDISC failed ulpStatus ulpWord4` Data: lsRjtError Severity: Error Log: Always Action: Reconfigure the switch to support more NPIV logins. If this issue persists, contact Broadcom Technical Support.
`0127: ELS timeout` An `ELS IOCB` command was posted to a ring and did not complete within ULP timeout seconds. Data: (1) `elscmd` (2) `remote_id` (3) `ulpcommand` (4) `ulpIoTag` Severity: Error Log: Always Action: If no ELS command is going through the adapter, reboot the system. If the issue persists, contact Broadcom Technical Support.
`0133: PLOGI: no memory for reg_login` Memory allocation error. Data: (1) `nlp_DID` (2) `nlp_state` (3) `nlp_flag` (4) `nlp_rpi` Severity: Error Log: LOG_ELS Action: Memory allocation error. Check system resources. Unload unused modules.
`0134: PLOGI cannot issue reg_login` The `ELS PLOGI` mailbox command has failed. Data: (1) `nlp_DID` (2) `nlp_state` (3) `nlp_flag` (4) `nlp_rpi` Severity: Error Log: LOG_ELS Action: Check the port and switch configuration.
`0135: cannot format reg_login` Could not allocate an RPI or DMA buffer for the mailbox command. Data: (1) `nlp_DID` (2) `nlp_state` (3) `nlp_flag` (4) `nlp_rpi` Severity: Error Log: LOG_ELS Action: None required.
`0136: PLOGI completes to NPort <DID> completion` A PLOGI has completed for which there is no NDLP. Data: (1) `ulpStatus` (2) `ulpWord`[4] Severity: Error Log: LOG_ELS Action: None required.
`0137: No retry ELS command <ELS_CMD> to remote` Data: (1) `ulpStatus` (2) `ulpWord`[4] Severity: Error Log: LOG_ELS Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`0138: ELS rsp: Cannot issue reg_login for <DID>` REG_LOGIN mailbox command failed. Data: (1) nlp_DID (2) nlp_state (3) nlp_flag (4) nlp_rpi Severity: Error Log: LOG_ELS Action: None required.
`0139: Ignoring ELS cmd tag <ioTag> completion Data` This ELS command was aborted. Data: (1) ulpStatus (2) ulpWord[4] (3) ulpTimeout Severity: Error Log: LOG_ELS Action: None required.
`0140: PLOGI Reject: invalid name` Invalid node WWN provided. Data: None Severity: Error Log: LOG_ELS Action: None required.
`0141: PLOGI Reject: invalid pname` Invalid port WWN provided. Data: None Severity: Error Log: LOG_ELS Action: None required.
`0142: PLOGI RSP: Invalid WWN` The PLOGI sent to the port by a remote port had an invalid WWN. Data: None Severity: Error Log: LOG_ELS Action: None required.
`0143: SLI4 Adapter Hardware Error Data: <status0>/<status1>` The HBA has encountered an unrecoverable error. Data: None Severity: Error Log: LOG_INIT Action: Use hbacmd to retrieve a dump file.
`0144: Not a valid WCQE code: <Completion Code>` The completion queue handler detected an invalid type. Data: None Severity: Error Log: LOG_SLI Action: None required.
`0147: Failed to allocate memory for RSCN event` Memory could not be allocated to send the RSCN event to the management application. Data: None Severity: Error Log: LOG_ELS Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`0148: Failed to allocate memory for LOGO event` Memory could not be allocated to send the LOGO event to the FC transport. Data: None Severity: Error Log: LOG_ELS Action: None required.
`0149: Failed to allocate memory for ELS event` Memory could not be allocated to send the ELS event to the FC transport. Data: None Severity: Error Log: LOG_ELS Action: None required.
`0154: Authentication not complete` Authentication was restarted because the previous authentication did not complete. Data: None Severity: Error Log: LOG_DISCOVERY Action: Check the switch configuration.
`0200: CONFIG_LINK bad hba state <hba_state>` A CONFIG_LINK mailbox command completed and the driver was not in the right state. Data: None Severity: Error Log: Always Action: Software driver error. If this issue persists, report these errors to Broadcom Technical Support.
`0203: Devloss timeout on WWPN <address> NPort <nlp_DID>` A remote N_Port that was discovered by the driver disappeared for more than lpfc_devloss_tmo seconds. Data: (1) nlp_flag (2) nlp_state (3) nlp_rpi Severity: Error Log: Always Action: If the device generating this message is not a target to which the HBA is connected, this error will not affect the data integrity of the I/O between the HBA and the attached storage and can be ignored.
`0206: Device discovery completion error` This indicates that an uncorrectable error was encountered during device discovery after a link up. FC devices will not be accessible if this message is displayed. Data: None Severity: Error Log: Always Action: Reboot the system. If this issue persists, report the error to Broadcom Technical Support. Run with verbose mode enabled for more information.
`0207: Device <DID> (<WWN>) sent invalid service parameters. Ignoring device.` Invalid service parameters were received from DID. This remote port is ignored. Data: DID, WWN Severity: Error Log: Always Action: Verify the remote port's configuration. If this issue persists, report the error to Broadcom Technical Support. Run with verbose mode on for more details.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`0217: Block sgl registration required DMAsize <reqlen> great than a page` The request to post SGL pages does not fit on a page. Data: None Severity: Warning Log: LOG_INIT Action: None required.
`0221: FAN timeout` A link up event was received without the login bit set, so the driver waits E_D_TOV for the fabric to send a FAN. If no FAN if received, a FLOGI will be sent after the timeout. Data: None Severity: Warning Log: LOG_DISCOVERY verbose Action: None required. The driver recovers from this condition by issuing a FLOGI to the fabric.
`0222: Initial FLOG/FDISKI timeout` The driver sent the initial FLOGI or FDISK to the fabric and never got a response back. Data: None Severity: Error Log: Always Action: Check fabric configuration. The driver recovers from this situation and continues with device discovery.
`0223: Timeout while waiting for NameServer login` Our login request to the NameServer was not acknowledged within R_A_TOV. Data: None Severity: Error Log: Always Action: Check the fabric configuration. The driver recovers from this and continues with device discovery.
`0224: NameServer Query timeout` Node authentication timeout, node Discovery timeout. A NameServer Query to the fabric or discovery of reported remote N_Ports is not acknowledged within R_A_TOV. Data: (1) fc_ns_retry (2) fc_max_ns_retry Severity: Error Log: Always Action: Check fabric configuration. The driver recovers from this situation and continues with device discovery.
`0226: Device discovery completion error` This indicates that an uncorrectable error was encountered during device discovery after a link up. FC devices will not be accessible if this message is displayed. Data: None Severity: Error Log: Always Action: Reboot the system. If this issue persists, report the error to Broadcom Technical Support. Run with verbose mode on for more details.
`0227: Node Authentication timeout` The driver has lost track of what N_Ports are being authenticated. Data: None Severity: Error Log: Always Action: None required. The driver should recover from this event.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`0228: CLEAR LA timeout` The driver issued a CLEAR_LA that never completed. Data: None Severity: Error Log: Always Action: None required. The driver should recover from this event.
`0230: Unexpected timeout, hba linkstate <link_state>` Discovery has timed out, and the HBA state is not ready. Data: None Severity: Error Log: LOG_DISCOVERY Action: None required.
`0231: RSCN timeout` The driver has lost track of which N_Ports have RSCNs pending. Data: (1) fc_ns_retry (2) lpfc_max_ns_retry Severity: Error Log: Always Action: None required. The driver should recover from this event.
`0233: Nodelist not empty` Driver unloaded or hotplug detected a node still in use. Data: None Severity: Error Log: LOG_DISCOVERY Action: None required.
`0237: Pending Link Event during Discovery: State <hba_state>` Received link event during discovery. Causes discovery restart. Data: None Severity: Warning Log: LOG_DISCOVERY verbose Action: None required, unless this issue persists. If persistent, check cabling.
`0241: NameServer rsp error` The driver received a NameServer response containing a status error. Data: (1) CommandResponse.bits.CmdRsp (2) ReasonCode (3) Explanation (4) fc_flag Severity: Error Log: LOG_DISCOVERY verbose Action: Check the fabric configuration. The driver recovers from this situation and continues with device discovery.
`0246: RegLogin failed` The firmware returned a failure for the specified RegLogin. Data: (1) Did (2) mbxStatus (3) hbaState Severity: Error Log: Always Action: This message indicates that the firmware could not do RegLogin for the specified DID. There might be a limitation on how many nodes an HBA can see.
`0249: Cannot issue Register Fabric login: Err %d\` Could not issue the fabric register login, the error value is unique for each possible failure. Data: None Severity: Error Log: LOG_ELS Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`0251: NameServer login: no memory` Could not allocate memory for the NDLP structure. Data: None Severity: Error Log: LOG_ELS Action: None required.
`0252: Cannot issue NameServer login` Could not issue an ELS PLOGI to the NameServer DID. Data: None Severity: Error Log: LOG_ELS Action: Check the port connection and the switch configuration.
`0253: Register VPI: Can't send mbox\` Could not issue the REG_VPI mailbox command for this VPort. Data: None Severity: Error Log: LOG_MBOX Action: None required.
`0254: Register VPI: no memory" goto mbox_err_exit` Could not allocate memory for the REG_VPI mailbox command. Data: None Severity: Error Log: LOG_MBOX Action: None required.
`0255: Issue FDISC: no IOCB` All of the preallocated IOCBs are in use. Data: None Severity: Error Log: LOG_ELS Action: None required.
`0256: Issue FDISC: Cannot send IOCB\` Unable to send the fabric IOCB. Data: None Severity: Error Log: LOG_ELS Action: None required.
`0257: GID_FT Query error` The GID_FT CT request for the NameServer has failed. Data: None Severity: Error Log: LOG_ELS Action: Check the switch configuration.
`0258: Register Fabric login error:` The REG_LOGIN for the fabric has failed. Data: None Severity: Error Log: LOG_MBOX Action: Check the port connection and the switch configuration.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`0259: No NPIV Fabric support` The switch to which the port is connected does not support NPIV. Data: None Severity: Error Log: LOG_ELS Action: Check the switch configuration.
`0260: Register NameServer error:` The REG_LOGIN mailbox command has failed for the NameServer. Data: None Severity: Error Log: LOG_ELS Action: Check the switch configuration
`0261: Cannot register NameServer login:` Either a memory allocation issue or an invalid parameter was sent to the REG_LOGIN. Data: None Severity: Error Log: LOG_ELS Action: At least one message (0142 0121 0133 0134 0135) should precede this message.
`0262: No NPIV Fabric support` The switch to which the port is connected does not support NPIV. Data: None Severity: Error Log: LOG_ELS Action: Check the switch configuration.
`0263: Discovery Mailbox error: state:` Either the driver could not allocate resources or it could not send sparam_mbox or cfglink_mbox. Data: (1) address of sparam_mbox command (2) address of cfglink_mbox command Severity: Error Log: LOG_MBOX Action: Attempt to unload and reload the driver when it is convenient.
`0264: No NPIV Fabric support` The switch to which the port is connected does not support NPIV. Data: None Severity: Error Log: LOG_ELS Action: Check the switch configuration.
`0266: Issue NameServer Req <cmdcode> err <rc> Data: <fc_flag> <fc_rscn_id_cnt>` The driver was unable to send the NameServer CT command. Data: (1) vports fc_flag (2) vports fc_rscn_id_cnt Severity: Error Log: LOG_DISCOVERY Action: Check the port and switch configurations.
`0267: NameServer GFF Rsp <did> Error (<ulpStatus> <un.ulpWord[4]>) Data: <fc_flag> <fc_rscn_id_cnt>` The NameServer GFF CT request failed. Data: (1) vports fc_flag (2) vports fc_rscn_id_cnt Severity: Error Log: LOG_DISCOVERY Action: Check the port and switch configurations.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`0268: NS cmd <cmdcode> Error (<ulpStatus> <un.ulpWord[4]>)` The NameServer CT request failed. Data: None Severity: Error Log: LOG_DISCOVERY Action: Check the port and switch configurations.
`0271: Illegal State Transition: node <nlp_DID> event <evt>, state <nlp_state> Data: <nlp_rpi> <nlp_flag>` The current node state does not have a handler for this event. Data: (1) nlp_rpi (2) nlp_flag Severity: Error Log: LOG_DISCOVERY Action: Verify that all targets are still visible to the SCSI mid-layer.
`0272: Illegal State Transition: node <nlp_DID> event <evt>, state <nlp_state> Data: <nlp_rpi> <nlp_flag>` The driver is completing a PLOGI but does not have the rcv_plogi flag set. Data: (1) nlp_rpi (2) nlp_flag Severity: Error Log: LOG_DISCOVERY Action: Verify that all targets are still visible to the SCSI mid-layer.
`0273: Unexpected discovery timeout, vport State x%x` The discovery process has timed out. Data: None Severity: Error Log: LOG_DISCOVERY Action: Verify that all targets are still visible.
`0274: lpfc_nlp_put: ndlp:x%pusgmap:x%x refcnt:%d, void *)ndlp, ndlp->nlp_usg_map, atomic_read(&ndlp->kref.refcount)` Data: None Severity: Warning Log: LOG_NODE Action: None required.
`0275: lpfc_nlp_put: ndlp:x%pusgmap:x%x refcnt:%d, void *)ndlp, ndlp->nlp_usg_map, atomic_read(&ndlp->kref.refcount)` A kref_put was called again after the node was already inactive. Data: None Severity: Warning Log: LOG_NODE Action: None required.
`0276: lpfc_nlp_get: ndlp:x%pusgmap:x%x refcnt:%d, void *)ndlp, ndlp->nlp_usg_map, atomic_read(&ndlp->kref.refcount)` A kref_get was attempted on a node that was being released. Data: None Severity: Warning Log: LOG_NODE Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`0277: lpfc_enable_node: ndlp:x%pusgmap:x%x refcnt:%d, void *)ndlp, ndlp->nlp_usg_map, atomic_read(&ndlp->kref.refcount)` Enable node was attempted on an inactive node. Data: None Severity: Warning Log: LOG_NODE Action: None required.
`0278: lpfc_enable_node: ndlp:x%pusgmap:x%x refcnt:%d, void *)ndlp, ndlp->nlp_usg_map, atomic_read(&ndlp->kref.refcount)` Enable node was attempted on an inactive node. Data: None Severity: Warning Log: LOG_NODE Action: None required.
`0280: lpfc_cleanup_node: ndlp:x%pusgmap:x%x refcnt:%d, void *)ndlp, ndlp->nlp_usg_map, atomic_read(&ndlp->kref.refcount)` Node clean-up was attempted on a node that has already been marked for memory free. Data: None Severity: Warning Log: LOG_NODE Action: None required.
`0281: lpfc_cleanup_node: ndlp:x%pusgmap:x%x refcnt:%d, void *)ndlp, ndlp->nlp_usg_map, atomic_read(&ndlp->kref.refcount)` Node clean-up was called to prepare the node for release. Data: None Severity: Warning Log: LOG_NODE Action: None required.
`0282: ldid:x%x ndlp:x%pusgmap:x%x refcnt:%d, ndlp->nlp_DID, (void *)ndlp, lpfc_init.c-ndlp->nlp_usg_map,` Driver clean-up has found a node that is still on the node list during driver unload or PCI hot plug removal. Data: None Severity: Error Log: LOG_NODE Action: None required.
`0283: Failed to allocate mbox cmd memory` Mailbox allocation error. Data: None Severity: Error Log: LOG_INIT Action: None required.
`0285: Allocated DMA memory size <alloclen> is less than the requested DMA memorysize<reqlen>` Memory allocation was truncated. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`0286: lpfc_nlp_state_cleanup failed to allocate statistical data buffer <nlp_DID>` Memory allocation failed for node's statistical data. Data: None Severity: Error Log: LOG_INIT Action: None required.
`0287: lpfc_alloc_bucket failed to allocate statistical data buffer DID <nlp_DID>` Memory allocation failed for node's statistical data. Data: None Severity: Error Log: LOG_NODE Action: None required.
`0288: Unknown FCoE event type <event_type> event tag <event_tag>` The firmware has detected an unknown FCoE event. Data: None Severity: Error Log: LOG_SLI Action: Check the FCoE switch configuration and the HBA DCBX mode.
`0289: Issue Register VFI failed: Err <rc>` The driver could not register the Virtual Fabric Index for the FCFI. Data: None Severity: Error Log: LOG_ELS Action: Check the switch and port configurations.
`0290: The SLI4 DCBX asynchronous event is not handled yet` The SLI-4 DCBX asynchronous event is not handled yet. Data: None Severity: Error Log: LOG_SLI Action: None required.
`0291: Allocated DMA memory size (x%x) is less than the requested DMA memory size (x%x)` The asynchronous DCBX events are not handled in the driver. Data: None Severity: Error Log: LOG_INIT Action: Check the switch configuration.
`0293: PM resume failed to start worker thread: error=<error>` The PCI resume (hot plug) could not start the worker thread for the driver. Data: None Severity: Error Log: LOG_INIT Action: Unload and reload the driver.
`0294: PM resume Failed to enable interrupt` The PCI resume (hot plug) could not get an interrupt vector. Data: None Severity: Error Log: LOG_INIT Action: Unload and reload the driver.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`0297: Invalid device group<pci_dev_grp>` While unloading the driver, the driver detected a PCI device that it should not have claimed. Data: None Severity: Error Log: LOG_INIT Action: None required.
`0299: Invalid SLI revision <sli_rev>` While processing a host attention or unrecoverable error, the driver detected an invalid SLI revision. Data: None Severity: Error Log: LOG_INIT Action: None required.
`0300: LATT: Cannot issue READ_LA: Data:<rc>` The link attention handler could not issue a READ_LA mailbox command. Data: None Severity: Error Log: LOG_MBOX Action: None required.
`0301: READ_SPARAM: no buffers` The driver attempted to issue a READ_SPARAM mailbox command to the adapter, but there were no buffers available. Data: None Severity: Warning Log: LOG_MBOX verbose Action: This message indicates one of these two issues: (1) Kernel virtual memory is depleted. Check that the system meets minimum RAM requirements for the adapter. Try closing other applications to free some memory. (2) A possible driver buffer management issue. If this issue persists, report the error to Broadcom Technical Support.
`0302: REG_LOGIN: no buffers` The driver attempted to issue a REG_LOGIN mailbox command to the adapter, but no buffers were available. Data: (1) Did, (2) flag Severity: Warning Log: LOG_MBOX verbose Action: This message indicates one of these two issues: (1) Kernel virtual memory is depleted. Check that the system meets minimum RAM requirements for the adapter. Try closing other applications to free some memory. (2) A possible driver buffer management issue. If this issue persists, report the error to Broadcom Technical Support.
`0313: Ring <ringno> handler: unexpected Rctl <Rctl> Type <Type> received` The RCTL/Type of a received frame did not match any for the configured masks for the specified ring. Data: None Severity: Warning Log: LOG_SLI verbose Action: This error could indicate a software driver, firmware, or hardware issue. Report these errors to Broadcom Technical Support.
`0303: Ring <ringno> handler: portRspPut <portRspPut> is bigger then rsp ring <portRspMax>` The port rsp ring put index is larger than the size of the rsp ring. Data: None Severity: Error Log: Always Action: This error could indicate a software driver, firmware, or hardware issue. Report these errors to Broadcom Technical Support.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`0304: Stray mailbox interrupt, mbxCommand <mbxcommand> mbxStatus <mbxstatus>` Received a mailbox completion interrupt and there are no outstanding mailbox commands. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.
`0306: CONFIG_LINK mbxStatus error <mbxStatus> HBA state <hba_state>` The driver issued a CONFIG_LINK mailbox command to the HBA that failed. Data: None Severity: Error Log: Always Action: This error could indicate a firmware or hardware issue. Report these errors to Broadcom Technical Support.
`0310: Mailbox command <mbxcommand> timeout` A mailbox command was posted to the adapter and did not complete within 30 seconds. Data: (1) hba_state (2) sli_flag (3) mbox_active Severity: Error Log: Always Action: This error could indicate a software driver or firmware issue. If no I/O is going through the adapter, reboot the system. If this issue persists, report the error to Broadcom Technical Support.
`0312: Ring <ringno> handler: portRspPut <rspPutInx> is bigger then rsp ring <numRiocb>` The IOCB command rings put pointer is ahead of the get pointer. Data: None Severity: Error Log: LOG_SLI Action: None required.
`0313: Ring <ringno> handler: unexpected Rctl <Rctl> Type <Type> received` The RCTL/Type of a received frame did not match any for the configured masks for the specified ring. Data: None Severity: Warning Log: LOG_SLI verbose Action: This error could indicate a software driver, firmware or hardware issue. Report these errors to Broadcom Technical Support.
`0315: Ring <ringno> issue: portCmdGet <local_getidx> is bigger then cmd ring <max_cmd_idx>` The port command ring get index is greater than the size of the command ring. Data: None Severity: Error Log: Always Action: This error could indicate a software driver, firmware or hardware issue. Report these errors to Broadcom Technical Support.
`0317: iotag <ulp_IoTag> is out of range: max iotag <max_iotag> wd0 <wd0>` The IoTag in the completed IOCB is out of range. Data: None Severity: Error Log: Always Action: This error could indicate a software driver, firmware or hardware issue. Report these errors to Broadcom Technical Support.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`0318: Failed to allocate IOTAG. last IOTAG is <last_allocated_iotag>` The driver cannot allocate an IoTag. Display the last value used. Data: None Severity: Error Log: Always Action: This message indicates the adapter HBA I/O queue is full. Typically this happens when heavy I/O is running on a low-end (3-digit) adapter. Upgrade to a higher-end adapter.
`0319: READ_SPARAM mbxStatus error <mbxStatus> hba state <hba_state>` The driver issued a READ_SPARAM mailbox command to the HBA that failed. Data: None Severity: Error Log: Always Action: This error could indicate a firmware or hardware issue. Report these errors to Broadcom Technical Support.
`0320: CLEAR_LA mbxStatus error <mbxStatus> hba state <hba_state>` The driver issued a CLEAR_LA mailbox command to the HBA that failed. Data: None Severity: Error Log: Always Action: This error could indicate a firmware or hardware issue. Report these errors to Broadcom Technical Support.
`0322: Ring <ringno> handler: unexpected completion IoTag <IoTag>` The driver could not find a matching command for the completion received on the specified ring. Data: (1) ulpStatus, (2) ulpWord[4], (3) ulpCommand, (4) ulpContext Severity: Warning Log: LOG_SLI verbose Action: This error could indicate a software driver or firmware issue. If this issue persists, report these errors to Broadcom Technical Support.
`0323: Unknown Mailbox command <mbxCommand> Cmpl` A unknown mailbox command completed. Data: None Severity: Error Log: Always Action: This error could indicate a software driver, firmware, or hardware issue. Report these errors to Broadcom Technical Support.
`0324: Config port initialization error, mbxCmd <mbxCommand> READ_NVPARM, mbxStatus <mbxStatus>` READ_NVPARMS mailbox command failed during port configuration. Data: None Severity: Error Log: Always Action: This error could indicate a software driver, firmware, or hardware issue. Report these errors to Broadcom Technical Support.
`0328: Rsp Ring <ring number> error: IOCB Data:` The firmware has returned an error for this IOCB. Data: (1) <iocb word[0]:iocb word[7]>, (2) <rsp word[0]:rsp[word[7]> Severity: Warning Log: LOG_SLI Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`0330: IOCB wake NOT set` The completion handler associated with the IOCB was never called. Data: (1) timeout (2) timeleft/jiffies Severity: Error Log: Always Action: This error could indicate a software driver, firmware, or hardware issue. If this issue persists, report the error to Broadcom Technical Support.
`0334: Unknown IOCB command` Received an unknown IOCB command completion. Data: (1) type (2) ulpCommand (3) ulpStatus (4) ulpIoTag (5) ulpContext) Severity: Error Log: Always Action: This error could indicate a software driver or firmware issue. If this issue persists, report these errors to Broadcom Technical Support.
`0335: Unknown IOCB command` Received an unknown IOCB command completion. Data: (1) ulpCommand (2) ulpStatus (3) ulpIoTag (4) ulpContext) Severity: Error Log: Always Action: This error could indicate a software driver or firmware issue. If this issue persists, report these errors to Broadcom Technical Support
`0336: Rsp Ring <ringno> error: IOCB` An IOCB error has occurred on the specified ring. Data: (1) ulpWord[0], (2) ulpWord[1], (3) ulpWord[2], (4) ulpWord[3], (5) ulpWord[4], (6) ulpWord[5], (7) irsp+6, (8) irsp+7 Severity: Warning Log: LOG_SLI verbose Action: If this issue persists, check the targets. If the targets are not the problem, report the error to Broadcom Technical Support.
`0340: Adapter temperature is OK now` Adapter temperature has reverted to normal range. Data: Temperature in Celsius Severity: Error Log: LOG_TEMP verbose Action: No action needed, informational.
`0341: Ring <ringno> Cannot find buffer for an unsolicited iocb tag <un.ulpWord[3]>` No more preallocated buffers are available to handle unsolicited buffers. Data: None Severity: Error Log: LOG_SLI Action: Ensure that this port is not being managed by multiple ports.
`0342: Ring <ringno> Cannot find buffer for an unsolicited iocb tag <unsli3.sli3Words>` This is a multiple IOCB unsolicited command, and sufficient buffer space cannot be allocated for it. Data: None Severity: Error Log: LOG_SLI Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`0343: Ring <ringno> Cannot find buffer for an unsolicited iocb tag <un.ulpWord[3]>` No more preallocated buffers are available to handle unsolicited buffers. Data: None Severity: Error Log: LOG_SLI Action: None required.
`0344: Ring <ringno> Cannot find buffer for an unsolicited iocb tag <unsli3.sli3Words[7]>` No more preallocated buffers are available to handle unsolicited buffers. Data: None Severity: Error Log: LOG_SLI Action: None required.
`0345: Resetting board due to mailbox timeout iocb. tag 0x%x` A mailbox command failed to complete. The driver is resetting the port. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: If the mailbox command fails again, set the lpfc_log_verbose to LOG_MBOX and retry.
`0346: Ring <ring number> handler: unexpected ASYNC_STATUS evt_code <evt code> W0 <hex w0> W1 <hex w1> W2 <hex W2> W3 <hex W3> W4 <hex W4> W5 <hex W5> W6 <hex W6> W7 <hex W7> W8 <hex W8> W9 <hex W9> W10 <hex W10> W11<hex W11>` The HBA received an asynchronous event that was not a temperature event. Data: None Severity: Error Log: LOG_SLI Action: None required.
`0347: Adapter is very hot, please take corrective action` Adapter temperature is above normal range. Data: Temperature in Celsius Severity: Error Log: LOG_TEMP verbose Action: Shutdown and remove the HBA. Contact Broadcom Technical Support.
`0348: NameServer login: node freed` The enable mode failed to free up the NameServer login. Data: None Severity: Error Log: LOG_ELSI Action: None required.
`0349: rc should be MBX_SUCCESS` The next mailbox command on the mailbox queue has failed. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`0350: rc should have been MBX_BUSY` Attempting to unregister a default RPI from an interrupt context and the mailbox state is not busy. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.
`0351: Config MSI mailbox command failed, mbxCmd <u.mb.mbxComm>, mbxStatus <u.mb.mbxStatus>` The mailbox command sent to the firmware to configure the adapter to use MSI-X has failed. Data: None Severity: Warning Log: LOG_MBOX Action: Ensure that the hardware platform supports MSI-X.
`0352: Config MSI mailbox command failed, mbxCmd <u.mb.mbxCommand>, mbxStatus <u.mb.mbxStatus>` The mailbox command sent to the firmware to configure the HBA to use MSI-X has failed. Data: None Severity: Error Log: LOG_MBOX Action: Ensure that the hardware platform supports MSI-X.
`0353: Active Mailbox cleared - mailbox timeout exiting` The mailbox timeout handler has determined that the driver is in the process of completing this mailbox command. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.
`0357: MSI-X interrupt with no EQE` SLI-4 adapter interrupt on the slow path but there is no associated EQE. Data: None Severity: Warning Log: LOG_SLI Action: None required.
`0358: MSI-X interrupt with no EQE` SLI-4 adapter interrupt on the fast path but there is no associated EQE. Data: None Severity: Warning Log: LOG_SLI Action: None required.
`0359:Not a valid slow-path completion event: majorcode=x%x, minorcode=x%x\n, bf_get(lpfc_eqe_major_code, eqe), bf_get(lpfc_eqe_minor_code, eqe));` SLI-4: The EQE is invalid. Data: None Severity: Error Log: LOG_SLI Action: None required.
`0360:Unsupported EQ count. <entry_count>` Cannot create an event queue of this size. Data: None Severity: Error Log: LOG_SLI Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`0361:Unsupported CQ count. <entry_count>` Cannot create an completion queue of this size. Data: None Severity: Error Log: LOG_SLI Action: None required.
`0362:Unsupported MQ count. <entry_count>` Cannot create MQ of this size. Data: None Severity: Error Log: LOG_SLI Action: None required.
`0364:Invalid param:` SLI-4: The post-SGL function was passed an invalid XRI. Data: None Severity: Error Log: LOG_SLI Action: None required.
`0365:Slow-path CQ identifier <CQID> does not exist:` The Completion Queue ID passed in the event queue entry does not reference a valid completion queue. Data: None Severity: Error Log: LOG_SLI Action: None required.
`0366:Not a valid fast-path completion event: majorcode=<major code hex>, minor-code=<minor code hex>` The major or minor code in the Event Queue field is invalid. Data: None Severity: Error Log: LOG_SLI Action: None required.
`0367: Fast-path completion queue does not exist` The fast-path completion queue referenced by the CQID does not exist. Data: None Severity: Error Log: LOG_SLI Action: None required.
`0368: Mis-matched fast-path completion queue identifier: eqcqid=%d, fcpcqid=%d` The CQID in the event queue entry does not match the fcp_cqid that was passed into the routine. Data: None Severity: Error Log: LOG_SLI Action: None required.
`0369: No entry from fast-path completion queue fcpcqid=<queue_id>` No completions in the completion queue were referenced by fcp_cqid. Data: None Severity: Error Log: LOG_SLI Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`0370: Invalid completion queue type <type>` The event queue entry is not for a mailbox or a work queue entry. Data: None Severity: Error Log: LOG_SLI Action: None required.
`0371: No entry from the CQ: identifier <queue_id>, type <type>` No completion queue event existed for this event queue entry. Data: None Severity: Error Log: LOG_SLI Action: None required.
`0372: iotag <iotag> is out of range: max iotag (<sli.last_iotag>)` The IOCB lookup cannot be performed because the iocb_tag is out of range. Data: None Severity: Error Log: LOG_SLI Action: None required.
`0373: FCP complete error: status=<status> hw_status=<hw status>, total_data_specified=<total data transferred>, parameter=<rsp word[4]>, word3=<wcqe word 3>` Logs the FCP failure. Status and parameter are equivalent to ulpStatus and ulpWord[4]. Data: None Severity: Warning Log: LOG_SLI Action: None required.
`0374: FCP complete with no corresponding cmdiocb: iotag <iocb iotag>` No IOCB was on the in-progress list that matched this iotag. Data: None Severity: Warning Log: LOG_SLI Action: None required.
`0375: FCP cmdiocb not callback function iotag: <iocb iotag>` The IOCB found for this iotag does not have a completion handler set in it. Data: None Severity: Warning Log: LOG_SLI Action: None required.
`0377: Error <rc> parsing vpd. Using defaults.` Could not parse the VPD data, so the driver is using the default values. Data: None Severity: Error Log: Always Action: None required.
`0378: No support for fcpi mode.` Could not configure the port to run in FCP initiator mode. Data: None Severity: Warning Log: LOG_MBOX, LOG_SLI Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`0379: Feature Mismatch Data: <req ftr word2 hex> <req_ftr word3 hex> <cfg_enable_npiv> <max vpi hex>` The features passed in to the driver as module parameters do not match what the firmware can do. Setting to default values. Data: None Severity: Warning Log: LOG_MBOX, LOG_SLI Action: None required.
`0381: Error %d during queue setup.` Could not set up all the queues that the driver requires to exchange I/Os with the HBA. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: Reload the driver.
`0382: READ_SPARAM command failed status <issue status>, mbxStatus <mailbox status>` The READ_SPARAM mailbox command has failed during initialization. The HBA has been set to error state. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI: Action: Perform a dump with hbacmd and then try reloading the driver.
`0383: Error <rc> during scsi sgl post operation` The SGL entries could not be registered with the adapter. Data: None Severity: Warning Log: LOG_MBOX, LOG_SLI Action: Reset the adapter using hbacmd.
`0384: There is pending active mailbox cmd` The mailbox commands have overlapped. This command should have been added to the mailbox queue. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.
`0385: rc should have been MBX_BUSY` The completion handler for REG_LOGIN detected the IMMED_UNREG flag and tried to issue the UNREG_LOGIN command from an interrupt level. The mailbox status should still be busy. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.
`0386: ELS complete with no corresponding cmdiocb: iotag <iotag>` The completion that the ISR is handling cannot find a tag associated with the IOTAG. Data: None Severity: Warning Log: LOG_SLI Action: None required.
`0387:Failed to allocate an iocbq` Failed to get an IOCBQ from the list of available IOCBQs. Data: None Severity: Error Log: LOG_SLI Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`0388:Not a valid WCQE code: x<hex cqe_code>` The event code is invalid. This event will be dropped. Data: None Severity: Error Log: LOG_SLI Action: Ensure that the adapter firmware is current.
`0391:Error during rpi post operation` The driver was trying to post pages to the firmware to keep target login information and encountered a failure. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: Unload and reload the driver.
`0393:Error <rc> during rpi post operation` The driver was trying to post pages to the firmware to keep target login information and encountered a failure. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: Unload and reload the driver.
`0394: Failed to allocate CQ_EVENT entry` The asynchronous event handler was unable to allocate an event queue entry to which to transfer the asynchronous event. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: This could be a V-LINK clear from the switch or a fatal error from the firmware. Perform a dump from the Emulex OneCommand Manager application.
`0395: The mboxq allocation failed` The asynchronous link event handler could not allocate a mailbox command to issue the READ_LA (read link attention) mailbox command. Data: None Severity: Error Log: LOG_SLI Action: None required.
`0396:The lpfc_dmabuf allocation failed` The asynchronous link event handler could not allocate a mailbox command to issue the READ_LA mailbox command. Data: None Severity: Error Log: LOG_SLI Action: None required.
`0397:The mbuf allocation failed` The asynchronous link event handler could not allocate DMA for the READ_LA mailbox command. Data: None Severity: Error Log: LOG_SLI Action: None required.
`0398 Invalid link fault code: <hex link_fault>` The attempt to read the link attention register has returned an unknown value. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

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

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

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

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`0427: Cannot re-enable interrupt after slot reset` The driver was unable to enable the interrupt after an HBA reset. Data: None Severity: Error Log: LOG_INIT Action: Unload and reload the driver.
`0429: MSI-X fast-path request_irq failed (<rc>)` The driver received an error for the request_irq_call. Data: None Severity: Warning Log: LOG_INIT Action: Unload and reload the driver.
`0430: PM resume Failed to enable interrupt` The driver power management resume function could not enable the interrupt. Data: None Severity: Error Log: LOG_INIT Action: Perform another PM suspend and resume or HBA reset.
`0431: Failed to enable interrupt.` The driver failed to start the interrupt. Data: None Severity: Error Log: LOG_INIT Action: Unload and reload the driver.
`0433: Wakeup on signal: rc=<rc>` A signal other than the LPFC_DATA_READY was received on the worker thread. Data: None Severity: Error Log: LOG_ELS Action: Unload and reload the driver.
`0434: PM resume failed to start worker thread: error=<error>` The driver's power management resume function could not start the worker thread. Data: None Severity: Error Log: LOG_INIT Action: Unload and reload the driver.
`0435: Adapter failed to get Option ROM version status <rc>` The driver could not read the HBA option ROM. Data: None Severity: Error Log: LOG_INIT Action: Reset the HBA. Ensure the adapter firmware is current.
`0436: Adapter failed to init, timeout, status reg <status>` The adapter failed during power-up diagnostics after it was reset. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`0437: Adapter failed to init, chipset, status reg <status>` The adapter failed during power-up diagnostics after it was reset. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.
`0438: Adapter failed to init, chipset, status reg <status>` The adapter failed during power-up diagnostics after it was reset. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.
`0439: Adapter failed to init, mbxCmd <mbxCommand> READ_REV, mbxStatus <mbxStatus>` Adapter initialization failed when issuing a READ_REV mailbox command. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.
`0440: Adapter failed to init, READ_REV has missing revision information` A firmware revision initialization error was detected. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. Install the latest firmware revision. If this issue persists, report the error to Broadcom Technical Support.
`0442: Adapter failed to init, mbxCmd <mbxCommand> CONFIG_PORT, mbxStatus <mbxStatus>` Adapter initialization failed when issuing a CONFIG_PORT mailbox command. Data: (1) hbainit Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.
`0446: Adapter failed to init, mbxCmd <mbxCommand> CFG_RING, mbxStatus <mbxStatus>, ring <num>` Adapter initialization failed when issuing a CFG_RING mailbox command. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.
`0447: Adapter failed init, mbxCmd <mbxCommand> CONFIG_LINK mbxStatus <mbxStatus>` Adapter initialization failed when issuing a CONFIG_LINK mailbox command. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.
`0448: Adapter failed to init, mbxCmd <mbxCommand> READ_SPARM, mbxStatus <mbxStatus>` Adapter initialization failed when issuing a READ_SPARM mailbox command. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`0449: lpfc_%attr attribute cannot be initialized to %d, allowed range is [%min, %max]` The sysfs attribute value written exceeds attribute range. Data: (1) attribute name (2) value written (3) minimum value (3) maximum value Severity: Error Log: Always Action: Write a value within the supported range.
`0450: lpfc_%attr attribute cannot be set to %d, allowed range is [%min, %max]` The sysfs attribute value written exceeds attribute range. Data: (1) attribute name (2) value written (3) minimum value (3) maximum value Severity: Error Log: Always Action: Write a value within the supported range.
`0451: Enable interrupt handler failed` The driver attempted to register the HBA interrupt service routine with the host operating system, but failed. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or driver issue. If this issue persists, report the error to Broadcom Technical Support.
`0453: Adapter failed to init, mbxCmd <mbxCommand> READ_CONFIG, mbxStatus <mbxStatus>` Adapter initialization failed when issuing a READ_CONFIG mailbox command. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.
`0454: Adapter failed to init, mbxCmd <mbxCommand> INIT_LINK, mbxStatus <mbxStatus>` Adapter initialization failed when issuing an INIT_LINK mailbox command. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.
`0456: Adapter failed to issue ASYNCEVT_ENABLE mbox status x%x` The mailbox command to enable an asynchronous event notification failed. Data: None Severity: Error Log: LOG_INIT Action: Ensure the adapter firmware is current. Reload the driver.
`0457: Adapter Hardware Error` The driver received an interrupt indicating a possible hardware issue. Data: (1) status (2) status1 (3) status2 Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.
`0458: Bring adapter online` The FC driver has received a request to bring the adapter online. This issue might occur when running lputil. Data: None Severity: Warning Log: LOG_INIT verbose Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`0459: Adapter heartbeat failure, taking this port offline.` The Heartbeat mailbox command failed. Data: None Severity: Error Log: LOG_INIT Action: Ensure that the adapter firmware is current. Reload the driver.
`0460: Bring adapter offline` The FC driver has received a request to bring the adapter offline. This issue might occur when running lputil. Data: None Severity: Warning Log: LOG_INIT verbose Action: None required.
`0466: Outstanding I/O when bringing Adapter offline` The I/O is still pending while attempting to stop the driver. Data: None Severity: Warning Log: LOG_INIT Action: None required.
`0467: lpfc_topology attribute cannot be set to %d, allowed range is [0, 6], phba->brd_no, val.` The lpfc_topology module parameter is invalid. Data: None Severity: Error Log: LOG_INIT Action: Use a value in the valid range.
`0468: lpfc_restrict_login must be 0 for Physical ports. "vport->cfg_restrict_login = 0;` Cannot restrict the login for the physical port. Data: None Severity: Error Log: LOG_INIT Action: None required.
`0469: lpfc_link_speed attribute cannot be set to %d, allowed range is [0, 8]` The link speed module parameter is invalid. Data: None Severity: Error Log: LOG_INIT Action: Use a link speed parameter in the valid range.
`0472: PCI channel I/O permanent failure` The PCI bus has detected an error. Data: None Severity: Error Log: LOG_INIT Action: Issue an HBA reset.
`0474: Unable to allocate memory for issuing MBOX_CONFIG_MSI command` Mailbox memory pool allocation error. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`0475: Not configured for supporting MSI-X cfg_use_msi: 0x%x` The lpfc_use_msi module parameter should have been set to 2. Data: None Severity: Error Log: LOG_INIT Action: Set module parameter lpfc_use_msi= 2.
`0476: HBA not supporting SLI-3 or later SLI Revision: <sli_rev>` The HBA does not support SLI-3 or SLI-4. Data: None Severity: Error Log: LOG_INIT Action: This HBA does not support msi. Set lpfc_use_msi = 0.
`0478: MSI request_irq failed (<rc>).` The request_irq kernel API has failed. Data: None Severity: Warning Log: LOG_INIT Action: Set lpfc_use_msi = 0.
`0479: Deferred Adapter Hardware Error` An adapter hardware error was sent to the driver. Data: (1) work_hs, (2) work_status[0], (3) work_status[1] Severity: Error Log: LOG_INIT Action: Perform a dump using hbacmd.
`0483:Invalid link-attention link speed: x%x", bf_get(lpfc_acqe_link_speed, acqe_link).` The link speed reported in the link attention interrupt is invalid. Data: None Severity: Error Log: LOG_INIT Action: Check the switch configuration.
`0485: MSI-X slow-path request_irq failed (<rc>).` The request_irq kernel API has failed. Data: None Severity: Warning Log: LOG_INIT Action: Set module parameter lpfc_use_msi = 0.
`0486: MSI-X fast-path (<index>) request_irq failed (<rc>).` The request_irq kernel API has failed. Data: None Severity: Warning Log: LOG_INIT Action: Set module parameter lpfc_use_msi = 0.
`0490: MSI request_irq failed (<rc>).` The request_irq kernel API has failed. Data: None Severity: Warning Log: LOG_INIT Action: Set module parameter lpfc_use_msi = 0.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`0492: Unable to allocate memory for issuing SLI_CONFIG_SPECIAL mailbox command` Mailbox memory pool allocation error. Data: None Severity: Error Log: LOG_INIT Action: None required.
`0493: SLI_CONFIG_SPECIAL mailbox failed with status<rc>` Mailbox command failed. Data: None Severity: Error Log: LOG_INIT Action: Ensure that the adapter firmware is current. Unload and reload the driver.
`0494: Unable to allocate memory for issuing "SLI_FUNCTION_RESET mailbox command"` Mailbox memory pool allocation error. Data: None Severity: Error Log: LOG_INIT Action: None required.
`0495: SLI_FUNCTION_RESET mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>.` Mailbox command failed. Data: None Severity: Error Log: LOG_INIT Action: Reset the HBA.
`0496: Failed allocate slow-path EQ` The event queue for the slow path was not allocated. Data: None Severity: Error Log: LOG_INIT Action: Unload and reload the driver.
`0497: Failed allocate fast-path EQ` The event queue for the fast path was not allocated. Data: None Severity: Error Log: LOG_INIT Action: Unload and reload the driver.
`0499: Failed allocate fast-path FCP CQ (<fcp_cqidx>).` The completion queue event for the fast path could not be allocated. Data: None Severity: Error Log: LOG_INIT Action: Unload and reload the driver.
`0500: Failed allocate slow-path mailbox CQ` Failed to allocate slow-path mailbox CQ. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`0501: Failed allocate slow-path ELS CQ` Failed to allocate slow-path ELS CQ. Data: None Severity: Error Log: LOG_INIT Action: None required.
`0503: Failed allocate fast-path FCP` Failed to allocate fast-path FCP. Data: None Severity: Error Log: LOG_INIT Action: None required.
`0504: Failed allocate slow-path ELS WQ` Failed to allocate slow-path ELS WQ. Data: None Severity: Error Log: LOG_INIT Action: None required.
`0505: Failed allocate slow-path ELS MQ` Data: None Severity: Error Log: LOG_INIT Action: None required.
`0506: Failed allocate receive HRQ\n` Data: None Severity: Error Log: LOG_INIT Action: None required.
`0507: Failed allocate receive DRQ` Failed to allocate receive DRQ. Data: None Severity: Error Log: LOG_INIT Action: None required.
`0520: Slow-path EQ not allocated` The slow-path EQ not allocated. Data: None Severity: Error Log: LOG_INIT Action: None required.
`0522: Fast-path EQ <fcp_eqidx> not allocated` The fast-path EQ is not allocated. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`0523: Failed setup of fast-path EQ <fcp_eqidx>, rc = <rc>` The fast-path EQ setup failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
`0526: Fast-path FCP CQ <fcp_cqidx> not allocated` The fast-path FCP is not allocated. Data: None Severity: Error Log: LOG_INIT Action: None required.
`0527: Failed setup of fast-path FCP CQ <fcp_cqidx>, rc = <rc>` The fast-path FCP CQ setup failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
`0528: Mailbox CQ not allocated` The mailbox CQ is not allocated. Data: None Severity: Error Log: LOG_INIT Action: None required.
`0530: ELS CQ not allocated` The ELS CQ is not allocated. Data: None Severity: Error Log: LOG_INIT Action: None required.
`0534: Fast-path FCP WQ <fcp_wqidx> not allocated` The fast-path FCP WQ is not allocated. Data: None Severity: Error Log: LOG_INIT Action: None required.
`0535: Failed setup of fast-path FCP WQ <fcp_wqidx>, rc = <rc>` The fast-path FCP WQ setup failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
`0536: Slow-path ELS WQ not allocated` The slow-path ELS WQ is not allocated. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`0538: Slow-path MQ not allocated` The slow-path MQ is not allocated. Data: None Severity: Error Log: LOG_INIT Action: None required.
`0540: Receive Queue not allocated` The Receive Queue is not allocated. Data: None Severity: Error Log: LOG_INIT Action: None required.
`0542: lpfc_create_static_vport failed to allocate mailbox memory` Failed to allocate mailbox memory for VPort creation. Data: None Severity: Error Log: LOG_INIT Action: None required.
`0543: lpfc_create_static_vport failed to allocate vport_info\n"))` Failed to allocate vport_info. Data: None Severity: Error Log: LOG_INIT Action: None required.
`0544: lpfc_create_static_vport failed to issue dump mailbox command ret <rc> status <mbxStatus>` Failed to issue a dump mailbox command for static VPort creation. Data: None Severity: Warning Log: LOG_INIT Action: None required.
`0545: lpfc_create_static_vport bad information header 0x%x 0x%x\n"\, le32_to_cpu(vport_info->signature), le32_to_cpu(vport_info->rev) & VPORT_INFO_REV_MASK);` Invalid information header; the signature or revision is invalid. Data: None Severity: Error Log: LOG_INIT Action: None required.
`0546: lpfc_create_static_vport failed to create vport` Failed to create a VPort. Data: None Severity: Warning Log: LOG_INIT Action: None required.
`0582: Error <rc> during sgl post operation` The SGL post operation failed. Data: None Severity: Error Log: LOG_MBOX, LOG_IP verbose Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`0602: Failed to allocate CQ_EVENT entry` Failed to allocate a CQ_EVENT entry. Data: None Severity: Error Log: LOG_SLI Action: None required.
`0603: Invalid work queue CQE subtype (x%x)\n", cq-<subtype>` Invalid work queue CQE. Data: None Severity: Error Log: LOG_SLI Action: None required.
`0700: Bus Reset on target <i> failed` The bus reset for the specified target failed. Data: None Severity: Error Log: LOG_FCP Action: None required.
`0704: At limitation of <total> preallocated command buffers` The maximum number of command buffers have already been allocated. Data: None Severity: Warning Log: LOG_FCP verbose Action: None required.
`0705: Allocation request of <num> command buffers will exceed max of <hba_queue_depth>. Reducing allocation request to <size>` The number of command buffers requested will exceed the maximum so a smaller quantity will be allocated. Data: None Severity: Warning Log: LOG_FCP verbose Action: None required.
`0708: Allocation request of <num_to_alloc> command buffers did not succeed. Allocated <num_allocated> buffers.` The allocation request for the specified command buffers did not succeed. However, the specified number of buffers has been allocated. Data: None Severity: Warning Log: LOG_FCP Action: None required.
`0711: detected queue full - lun queue depth adjusted to%d` The driver detected a queue full status on a SCSI command response. New LUN queue depth is reported. Data: (1) New LUN queue depth Severity: Warning Log: LOG_FCP verbose Action: This may indicate an oversubscribed target array. Check your SAN configuration and I/O workload.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`0713: SCSI layer issued Device Reset (%d, %d)` A device reset was issued. Data: None Severity: Error Log: LOG_FCP Action: None required.
`0714: SCSI layer issued bus reset` The SCSI layer is requesting the driver to abort all I/Os to all targets on this HBA. Data: (1) ret Severity: Error Log: Always Action: Check the state of the targets in question.
`0720: FCP command <cmnd[0]> residual overrun error` A residual overrun error has occurred while processing the specified FCP command. Data: (1) request_bufflen (2) resid Severity: Warning Log: LOG_FCP verbose Action: If this issue persists, check the targets for errors.
`0721: Device Reset rport failure: rdata <rdata>` The reset of the R_Port failed. Data: None Severity: Error Log: LOG_FCP Action: None required.
`0722: Target Reset rport failure: rdata <rdata>` The reset of the target failed. Data: None Severity: Error Log: LOG_FCP Action: None required.
`0723: SCSI layer issued Target Reset (%d, %d)` The SCSI layer issued a target reset. Data: None Severity: Error Log: LOG_FCP Action: None required.
`0724: I/O flush failure for context <"LUN","TGT","HOST","Unknown">: cnt <cnt>` The I/O flush to the LUN, target, or host has failed. Data: None Severity: Error Log: LOG_FCP Action: None required.
`0727: TMF <cmd> to TGT <TGT#> LUN <LUN#> failed (<ulpStatus>, <ulpWord[4]>)` The task management command failed. Data: None Severity: Error Log: LOG_FCP Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`0729: FCP cmd <cmnd> failed <target>/<lun> status: <status> result: <result>` The specified device failed an FCP command. Data: (1) ulpContext (2) iotag Severity: Warning Log: LOG_FCP verbose Action: Check the state of the target in question.
`0730: FCP command failed: RSP` The FCP command failed with a response error. Data: (1) resp_info (2) scsi_status (3) ResId (4) SnsLen (5) RspLen (6)rspInfo3 Severity: Warning Log: LOG_FCP verbose Action: Check the state of the target in question.
`0734: FCP read check error` The issued FCP command returned a read check error. Data: (1) fcpDl (2) rspResId (3) fcpi_parm (4) cmd[0] Severity: Warning Log: LOG_FCP verbose Action: Check the state of the target in question.
`0735: FCP Read Check Error and Underrun Data` HBA reported under run from storage array. Data: (1) vpi (2) fcpDl (3) res_id (4) fcpi_parm Severity: Warning Log: LOG_FCP_ERROR verbose Action: No action needed, informational.
`0748: Abort handler timed out waiting for abort to complete:ret <status> D <target id> LUN <lun id>` The abort handler timed out waiting for abort to complete. Data: None Severity: Error Log: Always Action: None required.
`0749: SCSI layer issued abort device` The SCSI layer aborted a device. Data: (1) ret, (2) id, (3) lun, (4) snum Severity: Warning Log: LOG_FCP verbose Action: None required.
`0915 Register VPI failed:<mbxStatus>` Could not register the VPI. Data: None Severity: Error Log: LOG_MBOX Action: None required.
`1019: Request tranid <tran_id> timed out` A transaction with storage array could not complete due to timeout. Data: (1) tran_id Severity: Warning Log: LOG_SECURITY verbose Action: Software driver warning. If this issue persists, report these errors to Broadcom Technical Support.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`1021: ERROR: attempted to queue security work, when no workqueue created` Driver encountered missing queue required for processing security information. Data: None Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report these errors to Broadcom Technical Support.
`1022: Security request does not exist` A security request operation failed because no match was found for such request. Data: None Severity: Warning Log: LOG_SECURITY Action: Software driver warning. If this issue persists, report these errors to Broadcom Technical Support.
`1023: Warning - data may have been truncated. Data: <data> reqdl: <data_len> mesdl:<data_len>` A security message exchange operation failed because the response was missing or unreliable. Data: None Severity: Warning Log: LOG_SECURITY Action: Software driver warning. If this issue persists, report these errors to Broadcom Technical Support.
`1028: Start Authentication: No buffers` The authentication failed because some memory resources were not allocated. Data: None Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.
`1029: Reauthentication Failure` The driver encountered errors and there was a failure to re-authenticate. Data: None Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.
`1031: Start Authentication: Get config failed` The authentication failed due to some error during port configuration. Data: None Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.
`1032: Start Authentication: get config timed out` The node authentication was aborted because waiting for port configuration to complete, timed out. Data: None Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.
`1033: Received auth_negotiate from Nport: < nlp_DID>` Unsolicited authentication negotiation message received from a port. Data: (1) nlp_DID Severity: Warning Log: LOG_SECURITY Action: No action, this message is informational.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`1034: Not Expecting Challenge - Rejecting Challenge` Unsolicited authentication challenge received from a port was rejected. Data: None Severity: Warning Log: LOG_SECURITY Action: Software driver warning. If this issue persists, report errors to Broadcom Technical Support.
`elx_mag1036: Authentication transaction reject - re-auth request reason <reason> exp <explanation>` An authentication was rejected and requested again due to reason as displayed with explanation. Data: (1) reason (2) explanation. Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.
`1037: Authentication transaction reject - restarting authentication, reason <reason> exp <explanation>` An authentication process was rejected then restarted and an authentication requested was again due to reason as displayed with the explanation. Data: (1) reason (2) explanation. Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.
`1039: Not Expecting Reply - rejecting. State <state>` An unanticipated reply was received during authentication and was subsequently rejected. Data: (1) auth_state. Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.
`1040: Bad Reply trans_id- rejecting. Trans_id < trans_id > Expecting: < trans_id>` Unexpected transaction id was received during authentication and was subsequently rejected. Data: (1) auth_state Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.
`1043: Authentication LS_RJT` The authentication request was rejected. Data: None Severity: Error Log: LOG_ELS Action: None required.
`1045: Issue AUTH_NEG failed Status:%x` The authentication negotiation failed. Data: None Severity: Error Log: LOG_ELS Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`1048: Issue AUTH_REJECT failed` Could not issue the reject for the authentication request. Data: None Severity: Error Log: LOG_ELS Action: None required.
`elx_ msg1050: Authentication mode is disabled, but is required by the fabric` Discovery failed because the switch fabric required authentication, but either authentication was not configured or the authentication mode for this port pair is disabled. Data: None Severity: Error Log: LOG_SECURITY Action: Configure the driver to authenticate with the switch or disable authentication on the switch to this port.
`1053: Start Authentication: Security service offline` The authentication failed because security service was unavailable. Data: None Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.
`1055: Authentication parameter is disabled, but is required by the fabric` FLOGI failed because the fabric has indicated that authentication is required, but authentication has not yet been configured or enabled on the HBA. Data: None Severity: Error Log: LOG_SECURITY Action: Configure authentication on this HBA.
`1057: Authentication transaction reject. reason <reason> exp <explanation>` An authentication was rejected and requested again due to the reason as displayed with the explanation. Data: (1) reason (2) explanation. Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.
`1058: Waiting for authentication service` A delay occurred when the authentication service was not initially available as expected. Data: None Severity: Warning Log: LOG_SECURITY Action: Software driver warning. If this issue persists, report these errors to Broadcom Technical Support.
`1059: Authentication became available` The authentication service came online, but it was not initially available as expected. Data: None Severity: Warning Log: LOG_SECURITY Action: Software driver warning. If this issue persists, report these errors to Broadcom Technical Support.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`1201: Failed to allocate dfc_host` Failed to allocate memory for the dfc_host_struct. Data: None Severity: Error Log: LOG_ELS Action: None required.
`1210: Invalid cmd size: cmd <cmd> cmdsz <cmdsize> rspsz <rspsize>` The management command for LPFC 2100 has failed. Data: None Severity: Error Log: LOG_LIBDFC Action: None required.
`1211: genreq alloc failed\n");` Resource allocation failure. Data: (1) return code Log: LOG_LIBDFC Severity: Error Action: Kernel memory resources are too low.
`1213: FCoE cmd overflow: off <#> + cnt <#> > cmdsz <#>` Application has tried to read more data than originally requested. Data: (1) response offset (2) size (3) cmd size Log: LOG_LIBDFC Severity: Error Action: The application might have sent a invalid command.
`1214: Cannot issue FCoE cmd SLI not active: <#> rc= -EACCESS` The SLI layer has not been initialized. Data: (1) offset Log: LOG_LIBDFC Severity: Error Action: Restart the HBA.
`1215: Cannot issue FCoE cmd: not ready or not in maint mode` Either the external link is unplugged, link down, and the FCoE is not in maintenance mode. Data: (1) current offset (2) return code. Log: LOG_LIBDFC Severity: Error Action: Plug external cable in or set FCoE in maintenance mode.
`1216: FCoE IOCB failed: off <#> rc <#>` FCoE command generated by the application has failed. Data: (1) offset (2) return code. Log: LOG_LIBDFC Severity: Error Action: Application should retry the command.
`1223: menlo_write: couldn't alloc genreq` Resource allocation failure. Data: None Log: LOG_LIBDFC Severity: Error Action: Kernel memory resources too low.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`1224: FCoE iocb failed off <#> rc=<#>",` FCoE command failed in SLI. Data: (1) offset (2) return code Log: LOG_LIBDFC Severity: Informational. Action: Retry the command. Iif it fails again, reset HBA when convenient.
`1227: FCoE IOCB TMO: handler set for <context3>` The management command for the LPFC 2100 has timed out. Data: None Severity: Warning Log: LOG_LIBDFC Action: None required.
`1228: FCoE IOCB TMO: handler set for <context3>` A management IOCB for the LPFC 2100 has timed out. Data: None Severity: Warning Log: LOG_LIBDFC Action: None required.
`1229: Waiting for menlo mnt` Waiting for the LPFC 2100 to enter maintenance method. Data: None Severity: Warning Log: LOG_LIBDFC Action: None required.
`1230: Could not find buffer for FCoE cmd:off <#> indmp <addr> off <#>` Could not find resources associated with this FCoE command. Data: (1) current offset (2) buffer desc pointer (3) size Severity: Error Log: LOG_LIBDFC Action: Try reloading the driver when convenient.
`1231: bad bpl:` An invalid buffer list was detected upon completion. Data: None Severity: Error Log: LOG_LIBDFC Action: None required.
`1235: Could not find buffer for FCoE cmd: off:<#> poff:<#> cnt:<#> mlastcnt:<#> addl:<x> addh:<x> mdsz:<#>` FCoE command failed because it could not find the resource. Data: (1) current offset (2)previous offset (3) count (4) last count (5) address low (6) address high Severity: Error Log: LOG_LIBDFC Action: No action needed, informational.
`1238: FCoE IOCB failed: off <#> rc=<#>` The command generated by the driver to check the FCoE has failed. Data: (1) offset (2) return code Log: LOG_LIBDFC Severity: Error Action: Make sure link is up or the adapter has set menlo in maintenance mode.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`1240: Unable to allocate command buffer memory` Could not allocate memory for the command buffer. Data: None Severity: Error Log: LOG_LINK_EVENT Action: None required.
`1243: Menlo command error. code=%d.\n", mlorsp->code` The Menlo maintenance command failed. Data: None Severity: Error Log: LOG_LINK_EVENT Action: None required.
`1244: Unable to allocate response buffer memory.` Could not allocate memory for the management command response. Data: None Severity: Error Log: LOG_LINK_EVENT Action: None required.
`1246: FCoE chip is running golden firmware. Update FCoE chip firmware immediately <fw_type>` The FCoE is running the golden firmware. Data: (1) firmware-type Severity: Error Log: LOG_LINK_EVENT Action: Try resetting the FCoE to operational mode and disable maintenance mode.
`1247: FCoE chip is running diagnostic firmware. Operational use suspended. <fw_type>` The FCoE is running a diagnostic. Data: (1) firmware-type Severity: Error Log: LOG_LINK_EVENT Action: Try resetting the FCoE to operational mode.
`1248: FCoE chip is running unknown firmware. <fw_type>` The FCoE is running an unknown firmware version. Data: (1) firmware-type Severity: Error Log: LOG_LINK_EVENT Action: Try resetting the FCoE to operational mode. Try loading the latest FCoE firmware.
`1249: Invalid FRU data found on adapter. Return adapter to Emulex for repair.` The FRU data on the FCoE chip is invalid. Data: (1) firmware-type Severity: Error Log: LOG_LINK_EVENT Action: Try resetting the FCoE to operational mode. Try loading the latest FCoE firmware or send the HBA back to Broadcom for repair.
`1250: Menlo command error. code=<#>` The IOCB driver sent to check FCoE state has bad header size. Data: (1) return code Log: LOG_LINK_EVENT Severity: Error Action: Try resetting the FCoE to operational mode.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

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

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

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

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`1312: Link Down Event <eventTag> received` Maintenance mode link up notification received without entering link down. Data: (1) fc_eventTag (2) port_state (3) vport fc_flag Severity: Error Log: LOG_LINK_EVENT Action: None required.
`1400: Failed to initialize sgl list.` Failed to initialize SGL list during initialization. Data: None Severity: Error Log: LOG_INIT Action: None required.
`1401: Failed to enable pci device.` Failed to enable PCI device during initialization. Data: None Severity: Error Log: LOG_INIT Action: None required.
`1402: Failed to set up pci memory space.` PCI initialization failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
`1403: Failed to set up driver resource.` Driver resource initialization failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
`1404: Failed to set up driver resource.` Driver resource initialization failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
`1405: Failed to initialize iocb list.` Driver resource initialization failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
`1406: Failed to set up driver resource.` Initialization failed to set up driver resource. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`1407: Failed to create scsi host.` Initialization failed to create SCSI host. Data: None Severity: Error Log: LOG_INIT Action: None required.
`1408: Failure HBA POST Status: sta_reg=0x%x, ""perr=x%x, sfi=x%x, nip=x%x, ipc=x%x, xrom=x%x, ""dl=x%x, pstatus=x%x\n", sta_reg.word0, bf_get(lpfc_hst_state_perr, &sta_reg),` The HBA POST has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
`1409: Failed to enable pci device.` Failed to enable PCI device during initialization. Data: None Severity: Error Log: LOG_INIT Action: None required.
`1410: Failed to set up pci memory space.` Initialization failed to set up PCI memory space. Data: None Severity: Error Log: LOG_INIT Action: None required.
`1411: Failed to set up driver resource.` Data: None Severity: Error Log: LOG_INIT Action: None required.
`1412: Failed to set up driver resource.` Initialization failed to set up driver resource. Data: None Severity: Error Log: LOG_INIT Action: None required.
`1413: Failed to initialize iocb list.` Initialization failed to initialize the IOCB list. Data: None Severity: Error Log: LOG_INIT Action: None required.
`1414: Failed to set up driver resource.` Initialization failed to set up driver resource. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`1415: Failed to create scsi host.` Initialization failed to create SCSI host. Data: None Severity: Error Log: LOG_INIT Action: None required.
`1416: Failed to allocate sysfs attr` Initialization failed to sysfs attribute. Data: None Severity: Error Log: LOG_INIT Action: None required.
`1418: Invalid HBA PCI device group: <dev_grp>` Invalid HBA PCI device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.
`1419: Invalid HBA PCI device group: <dev_grp>` Invalid HBA PCI device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.
`1420: Invalid HBA PCI device group:<dev_grp>` Invalid HBA PCI device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.
`1421: Failed to set up hba` Initialization failed to set up the HBA. Data: None Severity: Error Log: LOG_INIT Action: None required.
`1422: HBA Unrecoverable error: uerr_lo_reg=<ue lo>, uerr_hi_reg=<ue hi>, online0_reg=<Online0>, online1_reg=<Online1>` The HBA has notified the driver that it has encountered an unrecoverable error. Data: None Severity: Error Log: LOG_INIT Action: Perform a dump from the Emulex OneCommand Manager application. Then, unload and reload the driver.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`1423: HBA Unrecoverable error: uerr_lo_reg=<ue lo>, uerr_hi_reg=<ue hi>, online0_reg=<Online0>, online1_reg=<Online1>` The HBA has notified the driver that it has encountered an unrecoverable error. Data: None Severity: Error Log: LOG_INIT Action: Perform a dump from the Emulex OneCommand Manager application. Then, unload and reload the driver.
`1424: Invalid PCI device group:<pci_dev_grp>` Invalid HBA PCI device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.
`1425: Invalid PCI device group: <pci_dev_grp>` Invalid HBA PCI device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.
`1426: Invalid PCI device group: <pci_dev_grp>` Invalid HBA PCI device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.
`1427: Invalid PCI device group: <pci_dev_grp>` Invalid HBA PCI device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.
`1428: Invalid PCI device group: <pci_dev_grp>` Invalid HBA PCI device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.
`1429: Invalid PCI device group: <pci_dev_grp>` Invalid HBA PCI device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.
`1430: Failed to initialize sql list` Failed to initialize SQL list. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`1431: Invalid HBA PCI device group: <pci_dev_grp>` Invalid HBA PCI device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.
`1432: Failed to initialize rpi headers.` Failed to initialize RPI headers. Data: None Severity: Error Log: LOG_INIT Action: None required.
`1476: Failed to allocate sysfs attr` Failed to allocate sysfs attrlbutes. Data: None Severity: Error Log: LOG_INIT Action: None required.
`1477: Failed to set up hba` Failed to set up the HBA. Data: None Severity: Error Log: LOG_INIT Action: None required.
`1603: Loopback test did not receive expected data length. actual length <len>expected length <full_size>.` The loopback test did not receive the same amount of data that it transmitted. Data: None Severity: Error Log: LOG_LIBDFC Action: None required.
`1800: Could not issue unreg_vpi` Driver attempt to unregister VPI failed. Data: None Severity: Error Log: LOG_VPORT verbose Action: Software driver error. If this issue persists, report these errors to Broadcom Technical Support.
`1801: Create vport work array FAILED: cannot do scsi_host_get` The driver was unable to get a reference to a SCSI host. Data: None Severity: Warning Log: LOG_VPORT verbose Action: Software driver warning. If this issue persists, report these errors to Broadcom Technical Support.
`1816: FLOGI NPIV supported, response data <port>` The fabric reports support for NPIV upon FLOGI. Data: (1) response_multiple_NPort Severity: Warning Log: LOG_VPORT verbose Action: No action needed, informational.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`1817: Fabric does not support NPIV - configuring single port mode` The fabric reports no support for NPIV upon FLOGI. Data: None Severity: Warning Log: LOG_VPORT verbose Action: No action needed, informational.
`1818: VPort failed init, mbxCmd <mailbox command> READ_SPARM mbxStatus <mailbox status> , rc = <status>` A pending mailbox command that was issued to initialize the port failed. Data: (1) mbxCommand (2) mbxStatus (3) rc Severity: Error Log: LOG_VPORT verbose Action: Software driver error. If this issue persists, report these errors to Broadcom Technical Support.
`1819: Unrecognized lpfc_sli_mode parameter: <mode>` An attempt was made to set the SLI mode to an invalid value. The valid values for the SLI mode are 0, 2, and 3. Data: (1) lpfc_sli_mode Severity: Error Log: LOG_VPORT verbose Action: Correct the lpfc_sli_mode driver parameter setting. Valid values are 0, 2, and 3.
`1820: Unable to select SLI-3. Not supported by adapter.` The HBA is incapable of operating in a given mode. Data: None Severity: Error Log: LOG_VPORT verbose Action: SLI-3 mode is only available on some HBAs. Do not attempt to force the SLI mode to 3 on HBAs that do not support SLI-3 mode. This is an informational message. HBAs that do not support SLI-3 will be configured to run in SLI-2 mode, but it is recommended to use the auto setting (0).
`1821: Create VPORT failed. Invalid WWN format` The port could not be created due to an invalid WWNN or WWPN format. Data: None Severity: Error Log: LOG_VPORT verbose Action: Provide a valid WWN when creating Vports.
`1822: Invalid <name>: <xx: xx: xx: xx: xx: xx: xx: xx>` An invalid WWN was used when creating a VPort. Data: (1) type_name (2) wwn[1] (3) wwn[3] (3) wwn[5] (4) wwn[7] Severity: Error Log: LOG_VPORT verbose Action: When creating a VPort, you must furnish a valid WWN.
`1823: Create VPORT failed. Duplicate WWN on HBA.` The port could not be created because it would duplicate an existing WWNN HBA address. The resources for the port were discarded. Data: None Severity: Error Log: LOG_VPORT verbose Action: Provide a unique WWN.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`1824: NPIV enabled: Override lpfc_sli_mode parameter (<mode>) to auto(0)` The lpfc_enable_npiv and lpfc_sli_mode driver parameter settings conflict. The HBA must be configured for SLI-3 mode to support NPIV. Data: (1) lpfc_sli_mode Severity: Error Log: LOG_VPORT verbose Action: This is an informational message that indicates that the lpfc_enable_npiv and lpfc_sli_mode parameter settings are not compatible. Resolve the parameter conflict by setting the SLI mode to 0 or 3 or, if SLI-2 mode is required, disable NPIV.
`1825: Vport Created.` This message indicates that a port was created in the system. It is displayed at this level to ensure that it always appears at all log levels. Data: None Severity: Error Log: LOG_VPORT verbose Action: No action, informational.
`1826: Vport Disabled.` The port had to be disabled in the system. Data: None Severity: Error Log: LOG_VPORT verbose Action: No action, informational.
`1827: Vport Enabled` The port was enabled after possible recovery from some errors. Data: None Severity: Error Log: LOG_VPORT verbose Action: No action, informational.
`1828: Vport Deleted` A Vport was deleted. Data: None Severity: Error Log: LOG_VPORT verbose Action: No action, informational.
`1829: CT command failed to delete objects on fabric.` A command issued to the fabric to delete an associated resource for an object, such as for a port, failed. Data: None Severity: Error Log: LOG_VPORT verbose Action: Software driver error. If this issue persists, report these errors to Broadcom Technical Support.
`1830: Signal aborted mbxCmd <command>` A pending mailbox command was aborted because the thread received a signal. Data: None Severity: Error Log: LOG_VPORT verbose Action: Retry the attempted command.
`1831: Create VPORT Interrupted` The port creation process was unexpectedly interrupted at a critical time, and the operation was unsuccessful. Data: None Severity: Error Log: LOG_VPORT verbose Action: The process was interrupted while creating a VPort. Retry the command.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`1832: No pending MBOX command to handle` Data: None Severity: Error Log: LOG_MBOX Action: None required.
`1835: Vport discovery quiesce failed: state <port_state> fc_flags <fc_flag> wait msecs <jiffies_to_msecs(jiffies - start_time)>` Could not pause discovery on this VPort. Data: None Severity: Error Log: LOG_VPORT Action: None required.
`1836: Could not issue unreg_login(all_rpis) status <rc>` The unreg_login cannot be issued. Data: None Severity: Error Log: LOG_MBOX, LOG_VPORT Action: None required.
`1837: Vport_delete failed: Cannot delete static vport` Static VPorts cannot be deleted. Data: None Severity: Error Log: LOG_VPORT Action: None required.
`1838: Failed to INIT_VPI on vpi <vpi> status <rc>` Failed to INIT_VPI. Data: None Severity: Error Log: LOG_VPORT Action: None required.
`2000: Failed to allocate mbox for read_FCF cmd` Failed to allocate mailbox for READ_FCF command. Data: None Severity: Error Log: LOG_INIT Action: None required.
`2001: Unable to allocate memory for issuing SLI_CONFIG_SPECIAL mailbox command` Unable to allocate memory for issuing the SLI_CONFIG_SPECIAL mailbox command. Data: None Severity: Error Log: LOG_SLI Action: None required.
`2002: Error Could not grow rpi count` An error occurred because the RPI count could not be increased. Data: None Severity: Error Log: LOG_SLI Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`2004: Failed to allocate XRI.last XRITAG is <XRI> Max XRI is <MAX_XRI>, Used XRI is <USED_XRI>.` All XRIs are in use. Data: None Severity: Warning Log: LOG_SLI Action: None required.
`2005: Unable to deregister pages from HBA: <rc>` The SGL pages could not be unregistered from the firmware. Data: None Severity: Error Log: LOG_SLI Action: None required.
`2007: Only Limited Edition cmd Format supported <iocb.ulpCommand>` The SGL pages could not be unregistered from the firmware. Data: None Severity: Error Log: LOG_SLI Action: None required.
`2008: Error <rc> posting all rpi headers` The RPI headers could not be posted to the firmware. Data: None Severity: Error Log: LOG_SLI Action: None required.
`2009: Failed to allocate mbox for ADD_FCF cmd` Failed to allocate mailbox for ADD_FCF command. Data: None Severity: Error Log: LOG_INIT Action: None required.
`2010: Resume RPI Mailbox failed status <status>, mbxStatus <mbx status>` Data: None Severity: Error Log: LOG_SLI Action: None required.
`2011: Unable to allocate memory for issuing SLI_CONFIG_SPECIAL mailbox command` Data: None Severity: Error Log: LOG_SLI Action: None required.
`2012: Mailbox failed , mbxCmd <mbx_cmd> READ_CONFIG, mbxStatus <mbx status>` The READ_CONFIG mailbox command failed. Data: None Severity: Error Log: LOG_SLI Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`2013: Could not manually add FCF record 0, status <rc>` Could not add an FCF record to the FCF list. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.
`2014: Invalid command <iocb.ulpCommand>` The IOCB command is invalid. Data: None Severity: Error Log: LOG_SLI Action: None required.
`2015: Invalid CT %x command <iocb.ulpCommand>` Invalid command type in the IOCB is not supported. Data: None Severity: Error Log: LOG_SLI Action: None required.
`2017: REG_FCFI mbxStatus error <mbx status> HBA state <port_state>` The REG_FCFI mailbox command has failed. Data: None Severity: Error Log: LOG_MBOX Action: None required.
`2018: REG_VFI mbxStatus error <mbx status> HBA state <port_state>` The REG_VFI mailbox command has failed. Data: None Severity: Error Log: LOG_MBOX Action: None required.
`2020: Failed to allocate mbox for ADD_FCF cmd` Failed to allocate mailbox for ADD_FCF command. Data: None Severity: Error Log: LOG_INIT Action: None required.
`2022: VPI Mailbox failed status <status>, mbxStatus <mbxStatus>` The INIT VPI mailbox command has failed. Data: None Severity: Error Log: LOG_SLI Action: None required.
`2401: Failed to allocate memory for ELS XRI management array of size <els_xri_cnt>.` Initialization failed to allocate memory for the ELS XRI management array. Data: None Severity: Error Log: LOG_SLI Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`2500: EQ_CREATE mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>` The mailbox command sent to create the event queue has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
`2501: CQ_CREATE mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>` The mailbox command sent to create the completion queue has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
`2502: MQ_CREATE mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>` The mailbox command sent to create the mailbox queue has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
`2503: WWQ_CREATE mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>` The mailbox command sent to create the work queue has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
`2504: RQ_CREATE mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>` The mailbox command sent to create the receive queue has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
`2505: EQ_DESTROY mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>` The mailbox command sent to delete the event queue has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
`2506: CQ_DESTROY mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>` The mailbox command sent to delete the completion queue has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

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

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`2514: POST_RPI_HDR mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>` The mailbox command sent to post the RPUI header pages to the firmware has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
`2515: ADD_FCF_RECORD mailbox failed with status <rc>` The mailbox command to add the FCF record has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
`2516: DEL FCF of default FCF Index failed mbx status <rc>, status <shdr_status> add_status<shdr_add_status>` The mailbox command to delete the FCF record has failed. Data: None Severity: Error Log: LOG_SLI Action: None required.
`2517: Unregister FCFI command failed status %d, mbxStatus x%x", rc, bf_get(lpfc_mqe_status, &mbox->u.mqe)` The driver was unable to unregister the FCFI from the firmware. Data: None Severity: Error Log: LOG_SLI Action: None required.
`2518: Requested to send 0 NOP mailbox cmd` Data: None Severity: Warning Log: LOG_INIT Action: None required.
`2519: Unable to allocate memory for NOP mailbox command` Memory allocation for this mailbox command has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
`2520: NOP mailbox command failed status x%x add_status x%x mbx status x%x, shdr_status, shdr_add_status, rc.` The NOP mailbox command has failed. Data: None Severity: Warning Log: LOG_INIT Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`2521: READ_FCF_RECORD mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx` The READ_FCF_RECORD mailbox command failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
`2523: Allocated DMA memory size (<alloc_len>) is less than the requested DMA memory size (<req_len>)` The ADD_FCF_RECORD mailbox command failed to retrieve the length required from the firmware. Data: None Severity: Error Log: LOG_INIT Action: None required.
`2524: Failed to get the non-embedded SGE virtual address` The READ_FCF_RECORD mailbox command could not retrieve the SGE that was requested. Data: None Severity: Error Log: LOG_MBOX Action: None required.
`2527: Failed to allocate non-embedded SGE array.` Failed to allocate the non-embedded SGE array. Data: None Severity: Error Log: LOG_MBOX Action: None required.
`2528: Mailbox command <vpi> cannot issue` The mailbox command could not be issued because the mailbox interrupt is disabled. Data: (1) mbxCommand (2) sli_flag (3) flag Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.
`2529: Mailbox command <vpi> cannot issue` Data: (1) mbxCommand (2) sli_flag (3) flag Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.
`2530: Mailbox command <vpi> cannot issue` The SLI layer in the driver is inactive. Data: (1) mb.mbxCommand (2) sli_flag (3) flag Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.
`2531: Mailbox command <cpi> cannot issue` Data: (1) mb.mbxCommand (2) sli_flag (3) flag Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`2532: Mailbox command <vpi> (<mbxCommand>) cannot issue` The mailbox bootstrap code detected that the SLI layer is active. Data: (1) sli4_mbox_opcode (2) sli_flag,(3) MBX_POLL Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.
`2533: Mailbox command <vpi> (<mbxCommand>) cannot issue` Data: (1) sli4_mbox_opcode (2) sli_flag (3) MBX_NOWAIT Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.
`2535: Unsupported RQ count. (<entry_count>)` The receive queue ring can only be 512, 1024, 2048, or 4096. Data: None Severity: Error Log: LOG_SLI Action: None required.
`2536: Unsupported RQ count. (<entry_count>)` The receive queue ring can only be 512, 1024, 2048, or 4096. Data: None Severity: Error Log: LOG_SLI Action: None required.
`2537: Receive Frame Truncated!` The receive unsolicited handler detected a truncated frame. Data: None Severity: Error Log: LOG_SLI Action: None required.
`2539: Dropped frame rctl:%s type:%s\n` An unsupported frame was received by the port and dropped. Data: (1) rctl_names[fc_hdr->fh_r_ctl], (2) type_names[fc_hdr->fh_type] Severity: Error Log: Always Action: No action needed, informational.
`2540: Ring <ring #> handler: unexpected Rctl <fh_rctl> Type <fh_type>` The received frame has an unsupported RCTL or FH_TYPE. Data: None Severity: Warning Log: LOG_SLI Action: None required.
`2541: Mailbox command <vpi> (<mbxCommand>) cannot issue` Data: (1) sli_mbox_opcode (2) sli_flag (3) flag Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`2542: Try to issue mailbox command <vpi> (<mbxCommand>) synchronously ahead of async mailbox command queue` Attempting to send a synchronous mailbox command ahead of the asynchronous mailbox commands. Data: (1) sli4_mbx_opcode or sli_mbox_opcode, (2) sli_flag, (3) flag Severity: Warning Log: LOG_MBOX, LOG_SLI Action: None required.
`2543: Mailbox command <vpi> (<mbxCommand>) cannot issue` The mailbox command does not have all of the fields set correctly. Data: (1) sli_mbox_opcode (2) sli_flag (3) flag Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.
`2544: Mailbox command <vpi> (<mbxCommand>) cannot issue` The HBA cannot be accessed on the PCI bus. Data: (1) sli_mbox_opcode (2) sli_flag (3) flag Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.
`2546: New FCF found index <index> tag <event_tag>` A new FCF has been found. Data: None Severity: Error Log: LOG_DISCOVERY Action: None required.
`2547: Read FCF record failed` Could not read the FCF record from the firmware. Data: None Severity: Error Log: LOG_DISCOVERY Action: None required.
`2548: FCF Table full count <count> tag <event_tag>` The FCF table is full. Data: None Severity: Error Log: LOG_SLI Action: None required.
`2549: FCF disconnected from network index <index> tag <event_tag>` The FCF has disconnected from the network. Data: None Severity: Error Log: LOG_DISCOVERY Action: None required.
`2550: UNREG_FCFI mbxStatus error <u.mb.mbxStatus> HBA state <port_state>.` The UNREG_FCFI mailbox command has failed. Data: None Severity: Error Log: LOG_DISCOVERY. LOG_MBOX Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`2551: UNREG_FCFI mbox allocation failed HBA state <port_state>` The allocation for the UNREG_FCFI mailbox command has failed. Data: None Severity: Error Log: LOG_DISCOVERY, LOG_MBOX Action: None required.
`2552: UNREG_FCFI issue mbox failed rc <rc> HBA state <port_state>.` The UNREG_FCFI mailbox command has failed. Data: None Severity: Error Log: LOG_DISCOVERY, LOG_MBOX Action: None required.
`2553: lpfc_unregister_unused_fcf failed to read FCF record HBA state.` Data: None Severity: Error Log: LOG_DISCOVERY, LOG_MBOX Action: None required.
`2554: Could not allocate memory for fcf record` Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.
`2555: UNREG_VFI mbxStatus error <u.mb.mbxStatus> HBA state <port_state>` The UNREG_VFI mailbox command has failed. Data: None Severity: Error Log: LOG_DISCOVERY, LOG_MBOX Action: None required.
`2556: UNREG_VFI mbox allocation failed HBA state <port_state>` Could not allocate memory for UNREG_VFI mailbox command. Data: None Severity: Error Log: LOG_DISCOVERY, LOG_MBOX Action: None required.
`2557: UNREG_VFI issue mbox failed rc <rc> HBA state <port_state>` Could not issue the UNREG_VFI mailbox command. Data: None Severity: Error Log: LOG_DISCOVERY, LOG_MBOX Action: None required.
`2558: ADD_FCF_RECORD mailbox failed with status<shdr_status> add_status <shdr_add_status>` The ADD_FCF_RECORD mailbox command has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`2559: Block sgl registration required DMA size <reqlen> great than a page.` Attempting to register more SGEs with the firmware than can fit in a page. Data: None Severity: Warning Log: LOG_INIT Action: None required.
`2560: Failed to allocate mbox cmd memory\n` Failed to allocate mailbox command memory. Data: None Severity: Error Log: LOG_INIT Action: None required.
`2561: Allocated DMA memory size (<alloclen>) is less than the requested DMA memory size (<reqlen>)` Could not get the memory required for the number of XRIs that are attempting to be posted. Data: None Severity: Error Log: LOG_INIT Action: None required.
`2562: No room left for SCSI XRI allocation max_xri=<sli4_hba.max_cfg_param.max_xri>, els_xri=<els_xri_cnt>n` The number of allocated XRIs has reached the max_xri value. Data: None Severity: Error Log: LOG_SLI Action: None required.
`2563: Failed to allocate memory for SCSI XRI management array of size <sli4_hba.scsi_xri_max>.` Initialization could not allocate memory to hold the XRIs. Data: None Severity: Error Log: LOG_SLI Action: None required.
`2564: POST_SGL_BLOCK mailbox command failed status <shdr_status> add_status <shdr_add_status> mbx status <rc>` The list of XRI SGEs failed to be registered with the firmware. Data: None Severity: Error Log: LOG_SLI Action: None required.
`2566: Failed to allocate table entry` Failed to allocate connection table entry. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`2567: Config region 23 has bad signature` The driver was unable to read Config Region 23 because it has an invalid signature. Data: None Severity: Error Log: LOG_INIT Action: None required.
`2568: Config region 23 has bad version` The driver was unable to read Config Region 23 because it is an invalid version. Data: None Severity: Error Log: LOG_INIT Action: None required.
`2569: lpfc_dump_fcoe_param: memory allocation failed` Memory allocation has failed. Data: None Severity: Warning Log: LOG_MBOX Action: None required.
`2570: Failed to read FCoE parameters` The driver failed to read FCoE parameters. Data: None Severity: Error Log: LOG_MBOX, LOG_INIT Action: None required.
`2572: Failed allocate memory for fast-path per-EQ handle array` Failed to allocate memory for the fast-path per-EQ handle array. Data: None Severity: Error Log: LOG_INIT Action: None required.
`2573: Failed allocate memory for msi-x interrupt vector entries` The driver was unable to allocate memory during initialization of the MSI-X interrupt array. Data: None Severity: Error Log: LOG_INIT Action: None required.
`2574: Not enough EQs (<sli4_hba.max_cfg_param.max_eq>) from the pci function for supporting FCP EQs (<cfg_fcp_eq_count>)` Failed to create the minimum fast-path event queues. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`2575: Not enough EQs (<max_eq>) from the pci function for supporting the requested FCP EQs (<cfg_fcp_eq_count>), the actual FCP EQs can be supported: <eq_count>` The driver was not configured with enough fast-path event queues. Data: None Severity: Warning Log: LOG_INIT Action: None required.
`2576: Failed allocate memory for fast-path EQ record array` Failed to allocate memory for the fast-path EQ record array. Data: None Severity: Error Log: LOG_INIT Action: None required.
`2577: Failed allocate memory for fast-path CQ record array` Failed to allocate memory for the fast-path EQ record array. Data: None Severity: Error Log: LOG_INIT Action: None required.
`2578: Failed allocate memory for fast-path WQ record array` Failed to allocate memory for the fast-path EQ record array. Data: None Severity: Error Log: LOG_INIT Action: None required.
`2579: Slow-path wqe consume event carries miss-matched qid: wcqe-qid=<wcqe_quid>, sp-qid=<sp_quid>` The consumed entry does not have the queueID of the slow-path. Data: None Severity: Warning Log: LOG_SLI Action: None required.
`2580: Fast-path wqe consume event carries miss-matched qid: wcqe-qid=<fcp_wqid>.` The consumed entry does not have the queueID of the fast path. Data: None Severity: Warning Log: LOG_SLI Action: None required.
`2581: Not enough WQs (<sli4_hba.max_cfg_param.max_wq>) from the pci function for supporting FCP WQs (<cfg_fcp_wq_count>)` The driver was not configured with the minimum number of fast-path work queues. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`2582: Not enough WQs (<max_wq>) from the pci function for supporting the requested FCP WQs (<cfg_wq_count>), the actual FCP WQs can be supported: <wq_count>` The driver was not configured with enough fast-path work queues. Data: None Severity: Warning Log: LOG_INIT Action: None required.
`2593: The FCP EQ count(<cfg_fcp_eq_count>) cannot be greater than the FCP WQ count(<cfg_fcp_wq_count>), limiting the FCP EQ count to <cfg_fcp_wq_count>` The fast-path event queue cannot be greater than the fast-path work queue count. Data: None Severity: Warning Log: LOG_INIT Action: None required.
`2597: Mailbox command <vpi> (<mbxCommand>) cannot issue Synchronou(2) sli_flag (3) flag` Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.
`2598: Adapter Link is disabled.` The adapter link has been disabled. Data: None Severity: Error Log: LOG_INIT Action: None required.
`2599: Adapter failed to issue DOWN_LINK mbox command rc <rc>` The driver was unable to issue the DOWN_LINK mailbox command. Data: None Severity: Error Log: LOG_INIT Action: None required.
`2600: lpfc_sli_read_serdes_param failed to allocate mailbox memory` Failed to allocate mailbox memory. Data: None Severity: Error Log: LOG_INIT Action: None required.
`2605: lpfc_dump_static_vport: memory allocation failed` Failed to allocate mailbox memory. Data: None Severity: Error Log: LOG_MBOX Action: None required.
`2606: No NPIV Fabric support` No NPIV Fabric support. Data: None Severity: Error Log: LOG_ELS Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`2607: Failed to allocate init_vpi mailbox` Failed to allocate the INIT_VPI mailbox command. Data: None Severity: Error Log: LOG_MBOX Action: None required.
`2608: Failed to issue Init VPI mailbox` The driver was unable to send an INIT_VPI mailbox command. Data: None Severity: Error Log: LOG_MBOX Action: None required.
`2609: Init VPI mailbox failed <u.mb.mbxStatus>` The INIT_VPI mailbox command failed. Data: None Severity: Error Log: LOG_MBOX Action: None required.
`2610: UNREG_FCFI mbox allocation failed` Failed to allocate mailbox memory. Data: None Severity: Error Log: LOG_DISCOVERY, LOG_MBOX Action: None required.
`2611: UNREG_FCFI issue mbox failed` Could not issue the UNREG_FCFI mailbox command. Data: None Severity: Error Log: LOG_DISCOVERY, LOG_MBOX Action: None required.
`2619: Config region 23 has bad signature` Configuration region 23 has an invalid signature. Data: None Severity: Error Log: LOG_INIT Action: None required.
`2620: Config region 23 has bad version` Configuration region 23 has an invalid version. Data: None Severity: Error Log: LOG_INIT Action: None required.
`2621: Failed to allocate mbox for query firmware config cmd` Failed to allocate mailbox memory. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 9 LPFC Error Log Messages and their Descriptions (Continued)

`2622: Query Firmware Config failed mbx status <rc>, status <shdr_status> add_status <shdr_add_status>` Could not read the firmware configuration. Data: None Severity: Error Log: LOG_SLI Action: None required.
`2623: FCoE Function not supported by firmware. Function mode = <function_mode>>` FCoE is not supported by this firmware. Data: None Severity: Error Log: LOG_SLI Action: Use the Emulex OneCommand Manager application to update to the latest firmware.
`2707: Ring <Ring#> handler: Failed to allocate iocb Rctl <fh_rctl> Type <fh_type> received` The driver was unable to allocate memory to send a query config mailbox command. Data: None Severity: Error Log: LOG_SLI Action: None required.
`2717: CT context array entry [<index>] over-run: oxid:<fh_ox_id>, sid:<fh_SID>` All of the array slots to hold buffers that are passed to the application are in use. Data: None Severity: Warning Log: LOG_ELS Action: None required.
`2718: Clear Virtual Link Received for VPI <index> tag <event_tag>` A Clear virtual link was received from the fabric for this VPI. Data: None Severity: Error Log: LOG_DISCOVERY Action: None required.
`2719: Invalid response length: tgt <TGT_ID> lun <LUN> cmnd <CMD> rsplen <RSPLEN>` The response length for this FCP command is not supported. Data: None Severity: Error Log: LOG_FCP Action: None required.
`2721: ndlp null for oxid %x SID %x\n, icmd->ulpContext, dfchba->ct_ctx[tag].SID);` The Node value for this SID is not in the node list. Data: None Severity: Warning Log: LOG_ELS Action: None required.
`2726: READ_FCF_RECORD Indicates empty FCF table` The driver requested the firmware provide a list of FCF entries to connect to and the firmware responded that the FCF table is empty. Data: None Severity: Error Log: LOG_INIT Action: None required.

4.4.2 Ethernet Driver Log Messages for LPe16202/OCe15100 Adapters

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

4.4.2.1 Retrieving Ethernet Driver Log Messages

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

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

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

A message is displayed similar to the following:

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

4.4.2.2 Ethernet Driver Log Messages and Their Descriptions

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

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

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

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

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

Table 10 Ethernet Driver Log Messages of Type

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

Table 10 Ethernet Driver Log Messages of Type (Continued)

Ethernet Log Message	Description	Type
`Failed to enable VLAN promiscuous mode`	Firmware command failed to enable VLAN promiscuous mode.	Error
`Could not get CRC for <DD> region`	The driver could not get enough information from the adapter to decide whether a region from a firmware image should be flashed. The driver skips updating this section. This is a very unlikely error.	Error
`Could not get <DD> MSI-x vector(s)` `Using <DD> vector(s)`	Enabled only a subset of MSI-x vectors requested.	Information
`Could not set PCI DMA Mask`	The operating system call to set the DMA mask failed.	Error
`created <DD> RX queue(s)`	An informational message logging number of receive rings was created.	Information
`created <DD> TX queue(s)`	An informational message logging number of transmit rings was created.	Information
`<DD> VFs are already enabled`	When unloading the driver while a VF is still assigned to VM, the VFs remain enabled. When the driver is reloaded, this message is logged with the number of VFs that are enabled and active.	Information
`Debug data event - <DDDD>,` `<DDDD>, <DDDD>, <DDDD>`	These log messages are generated by firmware to identify certain infrequent hardware and firmware events and aid further analysis if required. These messages by themselves are not indicative of any problem. Data from these logs might help Broadcom troubleshoot other problems, if any, that are observed along with these messages.	Information
`enabled <D> (out of <DD>) MSI-x vector(s) for NIC`	The number of MSI-x vectors enabled for the NIC function.	Information
`Enable VLAN promiscuous mode` `Disabling VLAN promiscuous mode`	This is an informational log message when the driver enables or disables VLAN promiscuous mode.	Information

`Error detected in the card`
`ERR: sliport status <DDDDDD>`
`ERR: sliport error1 <DDDDDD>`
`ERR: sliport error2 <DDDDDD>`
	An SLIPORT error was detected in the adapter, followed by status and error messages.	Error

`Error in mapping PCI BARs`	Initialization of the adapter failed due to an error while mapping PCI BARs.	Error
`Failed to optimize SR-IOV resources`	When the adapter is in an SRIOV-capable configuration, driver failed to redistribute the PF pool resources among the PF and requested number of VFs.	Error
`FW config: function_mode=<DD>, function_caps=<DD>`	A message indicating the function mode and function capabilities set on the adapter during initialization.	Information
`FW dump deleted successfully`	Previously generated FW dump is successfully deleted.	Information
`FW dump initiated successfully`	FW dump is successfully initiated on the adapter.	Information
`FW dump not generated`	Failed to initiate FW dump on the adapter.	Error
`Firmware flashed successfully`	This is an informational message that the firmware on the adapter has been updated.	Information
`Firmware load error`	Updating the adapter with new firmware failed. Usually this message is accompanied by a detailed message on the failure.	Error
`Firmware load not allowed (interface is down)`	Flashing firmware failed because the status of the interface is down.	Error
`Firmware update in progress`	Firmware update is in progress on the adapter.	Information
`Flashing firmware file <filename>`	This is an informational message that the firmware in the adapter is being updated with the firmware image in the file indicated.	Information

Table 10 Ethernet Driver Log Messages of Type (Continued)

Ethernet Log Message	Description	Type
`Flashing section type <DD> failed`	The driver could not get enough information from the FW UFI image to flash a section in the UFI image, after which the FW download will be aborted.	Error
`FW image size should be multiple of 4`	The FW UFI image size is invalid. Download the correct FW image and try again.	Error
`FW version is <DD.DD.DDD.DDD>`	This is an informational message to log the FW version loaded on the adapter.	Information
`INTx request IRQ failed - err <DDD>`	The request for INTx interrupt registration failed. The driver is nonfunctional if the INTx interrupt cannot be registered.	Error
`Invalid Cookie. FW image may be corrupted`	The firmware image under `/lib/firmware/<filename>` does not have the expected cookie. The firmware in this file will not be flashed. Copy the proper file and try flashing again.	Error
`Invalid digital signature`	Updating the adapter with new firmware failed. The firmware image being updated does not have a valid digital signature.	Error
`invalid mailbox completion`	The driver received an unexpected completion status for a firmware command.	Error
`Lancer FW download timed out`	FW command timed out, while downloading FW image.	Error
`Link down event type: <DD>`	The reason that the ASIC signaled the link status as down. Possible values are: 0 = Link down due to reasons other than those listed here. 1 = Link down caused by Dynamic Control channel protocol. 3 = Link down triggered by Virtual NIC configuration (for example, a zero bandwidth is assigned to a VNIC). 4 = Link down caused by Ethernet Pause frame flooding. 5 = Link down due to physical thermal temperature going up.	Information
`Link down/Link up`	This is an informational message about a change in link status.	Information
`LPVID <DD>`	VLAN filter configured in BIOS in multichannel configuration.	Information
`Mac address assignment failed for VF <DD>` `MAC <S>M set on VF <DD> Failed`	Firmware command fails to add a MAC address to a virtual function during initialization or when requested.	Error
`MAC address change to <S> failed`	Failed to change to new MAC address as requested.	Warning
`Mac hash table alloc failed`	The hash table to remember learned MAC addresses could not be allocated.	Error
`Max: txqs <DD>, rxqs <DD>, rss <DD>, eqs <DD>, vfs <DD>` `Max: uc-macs <DD> , mc-macs <DD>, vlans <DD>`	An informational message logging maximum resources available to the function. Resources include TX rings, RX rings, RX rings with RSS capability, number of VFs supported, unicast MACs, multicast MACs and VLAN filters, respectively.	Information
`Memory allocation failure`	The driver could not allocate the memory required for the requested operation.	Error
`Memory allocation failure during GET_MAC_LIST`	Failed to allocate memory to issue GET_MAC_LIST FW command.	Error
`Missing digital signature`	Updating the adapter with new firmware failed. The firmware image being updated does not have a digital signature.	Error
`Module param rx_frag_size must be 2048/4096/8192.` `Using 2048`	An unsupported receive buffer size was passed for the rx_frag_size module parameter. The driver ignores the specified value and uses the default RX buffer size of 2048.	Warning
`MSIx enable failed`	A request for enabling MSIx interrupts registration failed. The driver will then use INTx interrupts.	Warning

Table 10 Ethernet Driver Log Messages of Type (Continued)

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

`Unrecoverable Error detected in the adapter`
`Please reboot server to recover`
`UE LOW: <S> bit set`
`UE HIGH: <S> bit set`
	An unrecoverable error is detected in the adapter that requires a reboot to recover. Low and high bits are set in the data path in which error occurred.	Error

`User has aborted FW download`	The user requested abort when FW download is in progress.	Error

Table 10 Ethernet Driver Log Messages of Type (Continued)

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

Appendix A: Ethtool -S Option Statistics for LPe16202/OCe15100 Adapters in NIC Mode

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

Table 11 Ethtool -S Option Statistics

Name	Description
dma_map_errors	The number of packets dropped due to DMA mapping errors.
eth_red_drops	Received packets dropped due to the ASIC's Random Early Drop policy.
forwarded_packets	The number of packets generated by the ASIC internally. These packets are not handed to the host. This counter is shared across ports and all functions (NIC/FCoE).
jabber_events	The number of jabber packets received. Jabber packets are packets that are longer than the maximum size Ethernet frames and that have bad CRC.
link_down_reason	The reason that the ASIC signaled the link status as down. The various values are: ■ 0 – Link down due to reasons other than those listed here. ■ 1 – Link down caused by Dynamic Control channel protocol. ■ 3 - Link down triggered by Virtual NIC configuration (for example: zero bandwidth assigned to a VNIC). ■ 4 – Link down caused by Ethernet Pause frame flooding. ■ 5 – Link down due to physical thermal temperature going up.
pmem_fifo_overflow_drop	Received packets dropped when an internal FIFO going into main packet buffer tank (PMEM) overflows.
rx_address_filtered	Received packets dropped when they do not pass the unicast or multicast address filtering.
rx_alignment_symbol_errors	The number of packets dropped due to L1 alignment errors. This counter is on a per-port basis.
rx_compl_err	The number of RX completion errors received.
rx_control_frames	The number of control frames received.
rx_crc_errors	The number of packets dropped due to CRC errors.
rx_dropped_header_too_small	Received packets dropped when the IP header length field is less than 5.
rx_dropped_runt	Dropped receive packets due to runt packets (for example, packets shorter than the Ethernet standard).
rx_dropped_tcp_length	Received packets dropped when the TCP header length field is less than 5 or the TCP header length + IP header length is more than IP packet length.
rx_dropped_too_short	Received packets dropped when IP length field is greater than the actual packet length.
rx_dropped_too_small	Received packets dropped when IP packet length field is less than the IP header length field.
rx_drops_mtu	Received packets dropped when the frame length is more than 9018 bytes.
rx_drops_no_erx_descr	Received packets dropped due to the input receive buffer descriptor FIFO overflowing.
rx_drops_no_pbuf	Packets dropped due to lack of available HW packet buffers used to temporarily hold the received packets.
rx_drops_no_tpre_descr	Packets dropped because the internal FIFO to the offloaded TCP receive processing block is full. This could happen only for offloaded FCoE traffic.
rx_drops_too_many_frags	Received packets dropped when they need more than 8 receive buffers. This counter is always 0.
rx_frame_too_long	Received packets dropped when they are longer than 9216 bytes.
rx_in_range_errors	Received packets dropped when the Ethernet length field is not equal to the actual Ethernet data length.
rx_ip_checksum_errs, rx_tcp_checksum_errs, rx_udp_checksum_errs	Packets dropped due to TCP, IP, or UDP checksum errors.

Table 11 Ethtool -S Option Statistics (Continued)

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

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

Table 12 Transmit/Receive Queue Statistics

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

Appendix B: lpfc Driver BlockGuard Functionality

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

B.1 Overview

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

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

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

Figure 1 Data Block Showing Type 1 DIF

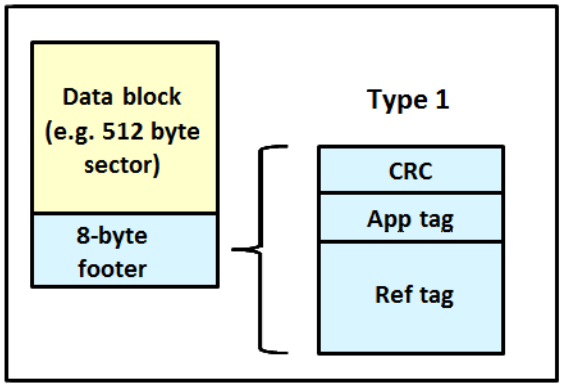

When data is written, the DIF is generated by the Host, or by the adapter, based on the block data and the logical block address. The DIF field is added to the end of each data block, and the data is sent through the SAN to the storage target. The storage target validates the CRC and Ref tag and, if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

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

B.2 Enabling BlockGuard

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

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

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

```
options lpfc lpfc_enable_bg=1
```

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

B.3 SCSI Command Support – SCSI Operation Codes

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

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

Figure 2 Data Transfer between Initiator Host and the Adapter

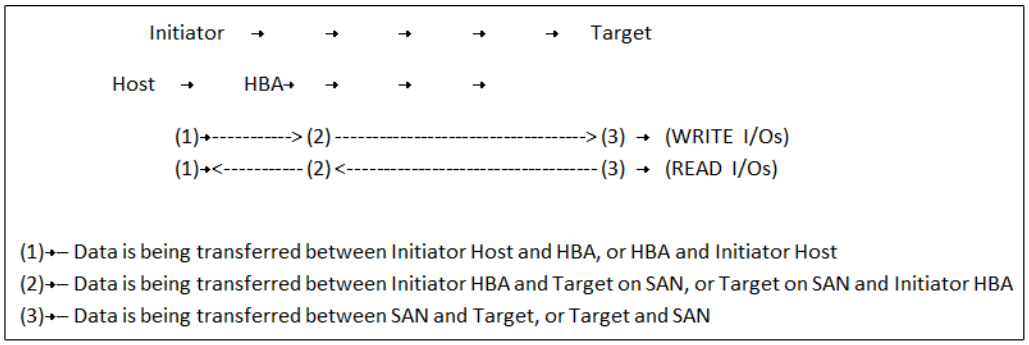

The initiator operations are listed in the following table.

Table 13 Initiator Operations

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

B.4 lpfc Driver Module Parameters

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

B.4.1 lpfc_prot_mask

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

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

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

Table 14 `lpfc_prot_mask` **Protection Types**

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

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

Table 15 Protection Data Placement for Supported Profiles

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

B.4.2 lpfc_prot_guard

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

Table 16 `lpfc_prot_guard` **Guard Types**

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

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

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


BROADCOM®

Emulex® Drivers for VMware ESXi for LightPulse® Adapters

User Guide

Version 11.4
September 6, 2017

DRVVM-LPe-UG114-100

Corporate Headquarters
San Jose, CA

Website
www.broadcom.com

Chapter 1: Introduction

This product supports the following Emulex® LightPulse® HBAs and converged fabric adapters (CFAs):

- LPe12000-series adapters
- LPe16000-series adapters, including LPe16202/OCe15100 adapters
- LPe31000-series adapters
- LPe32000-series adapters

The VMware ESXi drivers support the FC, FCoE, and Ethernet NIC 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 and FCoE
  - Adding LUNs and targets
  - Configuring driver parameters
  - Creating an FC remote boot disk
  - Managing devices through the CIM interface
  - Working with 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

Refer to the *Emulex OneCommand Manager for LightPulse Adapters User Guide* and *Emulex OneCommand Manager for VMware vCenter for LightPulse Adapters User Guides* for complete lists of supported technology.

1.1 ESXi Compatibility

For compatible firmware versions, go to the Documentation and Downloads page at http://www.broadcom.com for the specific adapter.

1.2 Supported Guest Operating Systems

The Emulex drivers support the following guest operating systems:

- CentOS 6.x
- CentOS 7.x
- RHEL 6.7, 6.8, and 6.9
- RHEL 7.1, 7.2, 7.3, and 7.4 (inbox only)
- SLES 11 SP3 and SP4
- SLES 12 SP2 and SP3
- XenServer Citrix 7.0
- OVM 3.3.3
- OVM 3.4.1

- Oracle UEK R3 and R4
- Oracle S10 and S11
- Windows Server 2012 and R2
- Windows Server 2016

1.3 Abbreviations

ACL	access control list
AL_PA	Arbitrated Loop Physical Address
API	application programming interface
BIOS	basic input/output system
CFA	converged fabric adapter
CLI	command line interface
CPU	central processing unit
CT	Common Transport
DCBX	Data Center Bridging Capabilities Exchange
DID	device ID
DMA	direct memory access
DRSS	default queue receive-side scaling
EEPROM	electrically erasable programmable read-only memory
ELS	extended link service
FA-PWWN	Fabric Assigned Port WWN
FC	Fibre Channel
FCF	Fibre Channel over Ethernet Forwarder
FCFI	Fibre Channel Forwarder Indicator
FCoE	Fibre Channel over Ethernet
FCP	Fibre Channel Protocol
FDMI	Fabric-Device Management Interface
FIP	FCoE Initialization Protocol
FLOGI	fabric login
FW	firmware
GbE	gigabit Ethernet
Gb/s	gigabits per second
GFC	gigabit Fibre Channel
GUI	Graphical User Interface
HBA	host bus adapter
HBAAPI	Host Bus Adapter Application Programming Interface
HBACMD	host bus adapter command
HW	hardware
I/O	input/output

IOCB	input/output control block
IOCTL	input/output control
IODM	I/O Device Management
IP	Internet Protocol
KV	Key Value
LPFC	Light Pulse Fibre Channel
LUN	logical unit number
MAC	Media Access Control
MB	megabyte
MDS	Microprocessor Development System
MSI	message signaled interrupt
MSI-X	message signaled interrupt-extended
MTU	maximum transmission unit
NIC	network interface card
NLP	Natural Language Processing
NPIV	N_Port ID Virtualization
N_Port	node port
PCI	Peripheral Component Interconnect (interface)
PCIe	PCI Express
PE	Protocol Endpoints
PF	physical function
PLOGI	port login
POST	power-on self-test
R_A_TOV	resource allocation timeout
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
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
VMID	Virtual Machine Identifier
VPD	vital product data
vPort	virtual port
VST	VLAN Switch Tagging
vSwitch	virtual switch
VVols	Virtual Volumes
VXLAN	virtual extensible 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, and NIC protocols.

2.1 Native Mode Overview

With the release of ESXi 5.5 (vSphere 2013) and subsequent operating systems, VMware introduced a new driver model called *native mode*. Emulex has endorsed the native mode driver model with the following drivers:

- For FC and FCoE functionality, the out-of-box driver is the native mode Light Pulse Fibre Channel (*lpfc)* driver.
- For Ethernet (NIC) functionality, the out-of-box driver for LPe16202/OCe15100 adapters is the native mode *lpnic* driver.

ESXi 5.5 and subsequent operating systems also require changes to the installation process and tools:

- **Driver parameters migration** – If you upgrade to ESXi 5.5 or later, 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 Later, for additional information.
- **Command line tool transition** – For the ESXi 5.5 release and all subsequent operating system releases, VMware has transitioned from `esxcfg` commands to `esxcli` commands. 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 4.3, Native Mode NIC Driver Troubleshooting Support, and Section 4.4, Native Mode NIC Driver Support for ethtool Commands, for additional information.

2.1.1 Migrating Driver Parameters When Upgrading to ESXi 5.5 or Later

VMware's ESXi 5.5 and later 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 6.0 and later 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 later.

If you install ESXi 5.5 or later 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 later, you must manually set all of the driver parameter values.

2.1.2 Preparing to Upgrade to ESXi 5.5 or Later

Upgrading your server to ESXi 5.5 or later 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 and 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 later, the server administrator must run the following command.

NOTE The driver binary name is now the native mode Emulex FC and 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 LPe16202/OCe15100 adapter in the host, because the `max_vfs` parameter receives the number of virtual functions (VFs) for each physical function (PF).

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

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

After upgrading to ESXi 5.5 or later, when the server administrator runs this command for the out-of-box native mode Ethernet driver for the LPe16202/OCe15100 adapter, there are no module parameters.

NOTE The driver binary name is now *lpnic* for the native mode Emulex Ethernet driver for the LPe16202/OCe15100 adapter.

```
~# esxcfg-module -g lpnic
lpnic 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 lpnic
```

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

```
~ # esxcfg-module -g lpnic
lpnic 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 and FCoE Driver

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

2.2.1 Installing the FC and FCoE Driver and Management Software

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

For VMware ESXi 5.5, 6.0, and 6.5 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. Refer to the *Emulex OneCommand Manager Application for LightPulse Adapters User Guide* and the *Emulex OneCommand Manager Application for VMware vCenter for LightPulse Adapters User Guide* for instructions on installing the respective applications. The installation requires that the CIM Provider be installed. For more information on installing the CIM Provider, refer to the *Emulex CIM Provider for LightPulse Adapters Installation Guide*.

NOTE Before installing the management application, you must install the FC or FCoE driver from the VMware website and then reboot the server.

Before installing the FC or FCoE driver 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 and FCoE Driver

Go to the VMware Patch Download page on the VMware website for instructions.

2.2.3 Installing the Native Mode FC and FCoE Driver esxcli Plug-In

This section describes the installation and usage of the `esxcli` plug-in for the native mode FC and FCoE (lpfc) driver on ESXi systems. This diagnostic tool is for the FC and FCoE driver module.

NOTE You can download the `esxcli` plug-in from the Broadcom® website.

To install the `esxcli elxfc` plug-in for ESXi 5.5, perform these 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 vib 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`
   After the driver has restarted, the `lpnic` namespace appears under `elxfc`.

Supported commands:

`esxcli elxfc`

Usage:

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

Table 1 Available Namespaces

Namespace	Description
adapter	Emulex FC/FCoE HBA instances adapter information
bginjerr	Emulex FC/FCoE HBA instances blockguard error injection page information
event	Emulex FC/FCoE HBA instances events information
eventsnap	Emulex FC/FCoE HBA instances events information with snap after display
fcf	Emulex FC/FCoE HBA instances FCF table information
heap	Emulex FC/FCoE HBA instances heap usage statistics and information
iostat	Emulex FC/FCoE HBA instances I/O statistics information
kv	Emulex FC/FCoE HBA instances Key Value (KV) page information
lun	Emulex FC/FCoE HBA instances SAN LUN information
memlog	Emulex FC/FCoE HBA instances memlog information
memory	Emulex FC/FCoE HBA instances memory statistics and information
mgmt	Emulex FC/FCoE HBA instances management statistics and information
node	Emulex FC/FCoE HBA instances nodes in SAN information
param	Emulex FC/FCoE HBA instances dynamic configuration parameter information
port	Emulex FC/FCoE HBA instances port information and statistics
rsnn	Emulex FC/FCoE HBA instances symbolic node name information
sli4q	Emulex FC/FCoE HBA instances sli4q information
target	Emulex FC/FCoE HBA instances SAN targets information
vmid	Emulex FC/FCoE HBA VMID page information
xlane	Emulex FC/FCoE HBA instances ExpressLane™ information

Table 2 Available Commands

Command	Description
list	Lists 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.

2.3 Installing the NIC Driver

This section provides installation information for the NIC driver for LPe16202/OCe15100 adapters in NIC+FCoE mode. 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

The latest Broadcom NIC and FCoE driver is available out-of-box. Refer to the operating system guide for installation instructions.

NOTE 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. Refer to the *Emulex OneCommand Manager Application for VMware vCenter for LightPulse Adapters User Guide* for additional information. The installation requires that the CIM Provider be installed. For more information on installing the CIM Provider, refer to the *Emulex CIM Provider for LightPulse Adapters Installation Guide*.

2.3.2 Uninstalling the NIC Driver

Go to the VMware Patch Download page on the VMware website for instructions.

2.3.3 Installing the Native Mode Ethernet Driver esxcli Plug-in

This section describes the installation and usage of the `esxcli` plug-in for the native mode Ethernet (`lpnic`) driver.

NOTE You can download the `esxcli` plug-in from the Broadcom website.

To install the `esxcli lpnic` plug-in for ESXi 5.5, perform these steps:

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

2.3.3.1 Supported Commands

NOTE The following command can be used to get the *<pci dev name>* value:
`#esxcli network nic list`

- **stats** – The `stats` command is invoked as follows:
  `esxcli lpnic stats get -p <pci dev name>`
  For example:
  `lpnic stats get -p 0000:003:00.1`
- **regdump (fatdump)** – The `regdump` command is invoked as follows:
  `esxcli lpnic regdump get -f <absolute file pathname> -p <pci device name>`

For example:

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

- **debugMask** – The `debugMask` command is invoked as follows:
  - To retrieve the `debugMask`, type:

    ```
    esxcli lpnic dbgmask get -p <pci dev name>
    ```
  - To set the `debugMask`, type:

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

For ESXi 6.0 and ESXi 6.5, follow the same procedure as described previously, but use the following command in step 2.

```
esxcli software vib install -v /<pathname>/ EMULEX-ESX-lpniccli-<version>.vib
```

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, type:

```
esxcli software vib list | grep lpnic
```

2.4 Updating Drivers with VMware Offline Bundle Media

VMware recommends using the offline bundle to update software on VMware ESXi platforms.

NOTE For more information about the ESXi patch management activities, go 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>/Emulex-FCoE-FC-lpfc-<driver-version>-offline-bundle-<OS
   version>.zip
   ```

   where `<driver-version>` represents the FC and FCoE or NIC driver.

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

   ```
   esxcli software vib install --maintenance-mode -d
   <pathname>/Emulex-FCoE-FC-lpfc-<driver-version>-offline-bundle-<OS
   version>.zip
   ```
2. Reboot the VMware ESXi Server to activate the drivers.

Chapter 3: Configuration

This section provides configuration information for the ESXi driver.

3.1 ESXi 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. 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 5.5, 6.0, and 6.5 Implementation

For ESXi 5.5 and subsequent operating systems, Broadcom uses the `esxcli` version of the command sequence, which has a different command syntax than did earlier releases

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>
```

3.1.1.1 FC and FCoE Driver Examples

To set extended logging for the Emulex FC and 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.

3.1.1.2 Ethernet Driver Examples

To list the module parameters supported by the `lpnic` driver:

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

To set the number of virtual functions for the Emulex `lpnic` driver:

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

To verify the altered parameter value in the `lpnic` driver module:

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

To disable message signaled interrupt-extended (MSI-X) support for the Emulex `lpnic` driver:

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

To verify the altered parameter value in the `lpnic` driver module:

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

To enable `emi_canceller` for the Emulex `lpnic` driver:

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

To verify the altered parameter value in the `lpnic` driver module:

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

To enable receive-side scaling (RSS) for the Emulex `lpnic` driver:

```
# esxcli system module parameters set -p RSS=4,4,4,4 -m lpnic
```

NOTE If RSS is to be disabled for a function, it must be specified as zero for that function.

To change the global `debugMask` value of the Emulex `lpnic` driver:

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

See Section 4.2.3, NIC Informational Log Groups, for more information.

To verify the altered parameter value in the `lpnic` driver module:

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

To enable default queue receive-side scaling (DRSS) for the Emulex `lpnic` driver:

```
~ # esxcli system module parameters set -p DRSS=4,4,4,4 -m lpnic
```

NOTE If DRSS is to be disabled for a function, it must be specified as zero for that function.

To enable the Dynamic NetQueue feature for the Emulex `lpnic` driver:

```
~ # esxcli system module parameters set -p dyn_netq=1 -m lpnic
```

To enable the VxLAN Offload feature for the Emulex `lpnic` driver:

```
~ # esxcli system module parameters set -p vxlan_offload=1 -m lpnic
```

NOTE Among the Emulex family, only the OCe14000-series adapters support the VxLAN Offload feature. The LPeOCe11000-series, LPe16000-series, LPe31000-series, LPe32000-series adapters do not support VxLAN Offload feature. Even if you specify `vxlan_offload = 1` for the LPeOCe11000-series, LPe16000-series, LPe31000-series, or LPe32000-series adapters, the driver ignores it.

3.2 FC and FCoE Driver Configuration

You can configure driver parameters using native ESXi tools, the 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, go to VMware's public website. If you have further questions, contact a VMware technical support representative.

Refer to the *Emulex OneCommand Manager Application for LightPulse Adapters User Guide* and the *Emulex OneCommand Manager Command Line Interface for LightPulse Adapters User Guide* for information about the OneCommand Manager application.

Refer to the *Emulex OneCommand Manager for VMware vCenter for LightPulse Adapters User Guide* for information about the OneCommand Manager for VMware vCenter Server application.

3.2.1 FC and 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 *Emulex OneCommand Manager Application for LightPulse Adapters User Guide* or the *Emulex OneCommand Manager for VMware vCenter for LightPulse Adapters User Guide* 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 and 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.

See Section 3.2.5, FC and 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 the *Emulex CIM Provider for LightPulse Adapters 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, 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.

2. To reboot the server, type:

   ```
   reboot
   ```

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

   N_Port ID Virtualization (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 LightPulse adapters in your system.

1. Locate the parameter `lpfc_lun_queue_depth` in Table 3.
2. Set the permanent value by typing

   ```
   esxcli system module parameters set -p lpfc_lun_queue_depth=20 -m lpfc
   ```
3. To reboot the server, type:

   ```
   reboot
   ```

The new setting is used when the driver reloads.

To verify the setting, type:

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

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 by typing:

   ```
   esxcli system module parameters set -p lpfc1_lun_queue_depth=20 -m lpfc
   ```
2. To reboot the server, type:

   ```
   reboot
   ```

The new setting is used when the driver is reloaded.

To verify the setting, type:

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

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.

1. Set the adapter-specific value by typing :

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

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, type:

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

3.2.2 Emulex ExpressLane Support

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

The ExpressLane feature enables special priority queuing for selected LUNs (ExpressLane LUNs). 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. Switches can provide up to three priority levels; Low, Medium, and High.

The following requirements must be met to use ExpressLane:

- ExpressLane is not supported on LPe12000-series adapters.
- 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 driver through the appropriate driver parameters.
- ExpressLane is not supported for LUNs attached to virtual ports.

For additional information, refer to the *Emulex OneCommand Manager for VMware vCenter for LightPulse Adapters User Guide*.

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 access control lists (ACLs), the same link bounce requirement applies after the security ACLs are corrected.

To force the ESXi server to rescan all devices, perform one of these actions:

- Run the following command:

  ```
  esxcfg-rescan vmhbaX
  ```

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

3.2.5 FC and FCoE Driver Configuration Parameters

Table 3 lists the FC and FCoE driver module parameters, their descriptions, and their corresponding values in ESXi native mode.

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

Table 3 FC and FCoE Driver Parameters

Module Parameter	Description and Whether Dynamic or Reboot Parameter	ESXi 5.5, ESXi 6.0, and ESXi 6.5 Native Mode Driver Model Values	Comments
`throttle_log_cnt`	Do not exceed this number of messages logged within `throttle_log_time`. Dynamic	Default (Def) = 10 Minimum (Min) = 1 Maximum (Max) = 1000	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). Dynamic	Def = 1 Min = 1 Max = 60	Works with `throttle_log_cnt`.
`compression_log`	Define how often the compression logs are written (in seconds). Dynamic	Def = 300 Min = 5 Max = 86400	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 Reboot	Def = 0 Min = 0 Max = 2	Enable this parameter to assist with SAN issues during ESX boot.
`max_targets`	The maximum number of discovered targets allowed. Reboot	Def = 256 Min = 0 Max = 4096	Driver parameter to adjust the supported target count.
`max_multiq`	Set how many completion queues the driver is requesting from ESXi for each HBA instance. Each completion queue uses an MSI-X vector. 0 indicates disabled Multi-Queue. Reboot	Def = 4 Min = 0 Max = 8	Sets the number of MultiQ FCP IO channels. The driver sets this parameter to determine how many completion queues to use per HBA port. Each completion queue consumes an MSI-X vector and defaults to a low number to ensure systems with large amount of HBAs do not run out of vectors.
`lpfc_delay_discovery`	Delay NPort discovery when the Clean Address bit is cleared. Reboot	Def = 0 Min = 0 Max = 1	When the parameter is set to `0`, no delay is added to the initial discovery. When this parameter is set to `1`, initial Nport discovery is delayed by RA_TOV seconds when the Clean Address bit is cleared in FLOGI/FDISC accept and the FCID/Fabric name/Fabric portname is changed.
`enable_da_id`	Control the DA_ID Common Transport (CT) command. Unregisters objects with the fabric nameserver.	Def 1 = Enable Min 0 = Disable Max 1 = Enable	

Table 3 FC and FCoE Driver Parameters (Continued)

Module Parameter	Description and Whether Dynamic or Reboot Parameter	ESXi 5.5, ESXi 6.0, and ESXi 6.5 Native Mode Driver Model Values	Comments
`enable_SmartSAN`	Enable SmartSAN functionality. The function works with FDMI-2 to provide enhanced Fabric diagnostics. Reboot	Def = 0 disable Min = 0 disable Max =1 enable	
`enable_fcp_priority`	Enable (1) or disable (0) FCP priority. Reboot	Def = 0 Min = 0 Max = 1	0 = Do not use the nodelist priority table (default) 1 = Use the nodelist priority table.
`iocb_cnt`	IOCBs allocated for extended link service, common transport, and abort sequence in 1024 increments. Reboot	Def = 1 Min = 1 Max = 5	
`devloss_tmo`	The number of seconds the driver holds I/O waiting for a lost device to return. Dynamic	Def = 10 Min = 1 Max = 255	
`log_verbose`	Verbose logging bit-mask. Dynamic	Def = 0 Min = 0 Max = 0x7fffffff	
`lun_queue_depth`	The maximum number of FCP commands that can queue to a specific LUN. **NOTE** 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. Dynamic	Def = 30 Min = 1 Max = 254	
`tgt_queue_depth`	The maximum number of FCP commands queued to a specific target port. Dynamic	Def = 65535 Min = 10 Max = 65535	

Table 3 FC and FCoE Driver Parameters (Continued)

Module Parameter	Description and Whether Dynamic or Reboot Parameter	ESXi 5.5, ESXi 6.0, and ESXi 6.5 Native Mode Driver Model Values	Comments
`hba_queue_depth`	The maximum number of FCP commands queued to an FC and FCoE adapter. The driver automatically adjusts the `hba_queue_depth` to match adapter capabilities. This setting may be overridden. Reboot	Def = 8192 Min = 32 Max = 8192	
`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 Emulex driver gradually reduces the LUN queue depth with each Task Set Full status completion, and then ramps back up as the /IO successfully completes. When disabled (value = 0), the driver takes no action when an I/O completes with a Task Set Full status. Reboot	Def = 1 Min = 0 Max = 1	Disabling this parameter has a potential impact on the overall VM performance. Consult your target vendor for guidance.
`scan_down`	Start scanning for devices from the highest Arbitrated Loop Physical Address (AL_PA) to the lowest. Reboot	Def = 1 Min = 0 Max = 1	
`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 Dynamic	Def = 0 Min = 0 Max = 6	The LPe31000-series and LPe32000-series adapters' topology cannot be changed and is set to point-to-point. **NOTE** The LPe31000 and LPe32000-series adapters are not manageable. Supports point-to-point only.

Table 3 FC and FCoE Driver Parameters (Continued)

Module Parameter	Description and Whether Dynamic or Reboot Parameter	ESXi 5.5, ESXi 6.0, and ESXi 6.5 Native Mode Driver Model Values	Comments
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 Dynamic	Def = 0 Min = 0 Max = 32	Updated for LPe32000-series adapters. For ESXi 5.5, link_speed can be changed on all adapters. For ESXi 6.0 and above, link_speed can only be changed using the driver parameter on LPe12000 series adapters. **NOTE** Setting this option incorrectly can cause the adapter to fail to initialize.
fcp_class	Select the FC class of service for FCP sequences. Reboot	Def = 3 Min = 2 Max = 3	
use_adisc	Use address discovery on rediscovery, initiated by registered state change notification (RSCN), to authenticate FCP devices instead of port login. Dynamic	Def = 0 Min = 0 Max = 1	
first_burst_size	First burst size for targets that support first burst. Dynamic	Def = 0 Min = 0 Max = 65536	
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. Dynamic	Def = 0 Min = 0 Max = 60000	
fdmi_on	Controls FDMI support. 0 = FDMI support off. 1 = FDMI support on. Reboot	Def = 1 Min = 0 Max = 1	If enable_SmartSAN is set 1, the driver automatically supports FDMI-2. If enable_SmartSAN is set 0, the driver uses the current value of fdmi_on to provide FDMI support – 0 meaning no support or 1 meaning FDMI-1 support. If FDMI-2 fails, the driver falls back to FDMI-1. If enable_SmartSAN is set to 1, the driver ignores the fdmi_on value and goes directly to FDMI-2 support. Traditional FDMI support means the driver will assume FDMI-2 support; however, if that fails, it falls back to FDMI-1.

Table 3 FC and FCoE Driver Parameters (Continued)

Module Parameter	Description and Whether Dynamic or Reboot Parameter	ESXi 5.5, ESXi 6.0, and ESXi 6.5 Native Mode Driver Model Values	Comments
discovery_threads	The maximum number of extended link service commands that can be outstanding during discovery. Reboot	Def = 32 Min = 1 Max = 64	
max_luns	The maximum number of LUNs allowed. Reboot	Def = 256 (ESXi 5.5) Def = 65535 (ESXi 6.0 and 6.5) Min = 1 Max = 256 (ESXi 5.5) Max = 65535 (ESXi 6.0 and 6.5)	Setting in ESXi 6.0 and 6.5 allows for sparse LUN IDs above 256.
task_mgmt_tmo	The maximum time to wait for task management commands to complete. Dynamic	Def = 60 Min = 5 Max = 180	
use_msi	Use preferred MSI-X interrupt mode, if possible. ■ 0 = Message signaled interrupt (MSI) disabled (INTx mode) ■ 1 = MSI enabled ■ 2 = MSI-X enabled Reboot	Def = 2 Min = 0 Max = 2	
fcf_failover_policy	■ 1 = Fibre Channel over Ethernet Forwarder (FCF) Fast failover ■ 2 = Priority failover Reboot	Def = 1 Min = 1 Max = 2	
enable_rrq	Enable Reinstate Recovery Qualifier functionality. ■ 0x0: Disabled, extensible resource indicator (XRI)/OXID use not tracked. ■ 0x1: XRI/OXID reuse is timed with resource allocation timeout (R_A_TOV), Reinstate Recovery Qualifier sent. ■ 0x2: XRI/OXID reuse is timed with R_A_TOV; No reinstate recovery qualifier is sent. Reboot	Def = 2 Min = 0 Max = 2	

Table 3 FC and FCoE Driver Parameters (Continued)

Module Parameter	Description and Whether Dynamic or Reboot Parameter	ESXi 5.5, ESXi 6.0, and ESXi 6.5 Native Mode Driver Model Values	Comments
EnableXLane	Enable ExpressLane. Reboot	Def = 0 Min = 0 Max = 1	1 = Enable ExpressLane 0 = Disable ExpressLane **NOTE** If the lpfc_vmid_priority_tagging parameter is enabled, EnableXLane is disabled as both the features conflict in using CS_CTL field in FC header.
XLanePriority	ExpressLane CS_CTL Priority. Sets the CS_CTL field in the FC header. Refer to the switch vendor administration guide for additional information. Reboot	Def = 0x0 Min = 0x0 Max = 0x7F	
sg_seg_cnt	The maximum scatter gather segment count for DMA. The maximum data allowed in one SG element is 0x80000000. Reboot	Def = 64 Min = 64 Max = 4096	
nlp_slab_cnt	NLP slab entries. Reboot	Def = 128 Min = 32 Max = 256	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_slab_cnt	Receive buffer slab entries. Reboot	Def = 256 Min = 32 Max = 256	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. Reboot	Def = 128 MB Min = 64 MB Max = 512 MB	
lpfc_mem_tracker	This turns memory tracking on (1) or off (0). Reboot	0 = Disabled (default) 1 = Enabled	Keeps track of driver heap and slab allocation/free. It is used only when debugging driver memory issues.
lpfc_external_dif	Enable external DIF support on select devices. Reboot	0 = Disabled 1 = Enabled (default)	External DIF is only available on FC HBAs (LPe16000-series, LPe31000-series, and LPe32000-series adapters).
lpfc_max_vmid	Maximum number of VMs to be tagged. Range 4 to 255. Reboot	Def = 8 Min = 4 Max = 255	This value indicates the number of VMIDs supported.
lpfc_vmid_inactivity_timeout	Inactivity timeout duration in hours. Range 0 to 24. Reboot	Def = 4 Min = 0 Max = 24	VMID is an ID assigned per VM. The VMID is removed when a VM is inactive for lpfc_vmid_inactivity_timeout duration. The lpfc_vmid_inactivity_timeout value is in hexadecimal.
lpfc_enable_bb_credit_recovery	Toggle the Buffer-to-Buffer Credit Recovery feature. Reboot	Def = 1 Min = 0 Max = 1	Applicable to SLI-4 adapters only.

Table 3 FC and FCoE Driver Parameters (Continued)

Module Parameter	Description and Whether Dynamic or Reboot Parameter	ESXi 5.5, ESXi 6.0, and ESXi 6.5 Native Mode Driver Model Values	Comments
`lpfc_enable_mds_diags`	Enable MDS diagnostics.	0 = disabled (default) 1 = enabled	Enables or disables the Cisco Fiber Channel Link Diagnostics feature. **NOTE** The parameter should be disabled (set to 0) after diagnostics are complete.
`lpfc_vmid_priority_tagging`	VMID CS_CTL tagging. Reboot	Min = 0 Max = 2 Def = 0	Cisco VMID uses the priority tagging field to communicate the capability. `vmid_priority_tagging` value `0` means disable priority tagging. Priority tagging is also disabled when the parameter is set to anything other than `1` or `2`. `vmid_priority_tagging` value `1` means enable priority tagging only for targets that support it in their port login (PLOGI) `LS_ACC` response. `vmid_priority_tagging` value `2` means enable priority tagging for all targets, whether or not they support it in their `PLOGI LS_ACC` response.
`lpfc_vmid_app_header`	VMID Application ID tagging. Reboot	0 = Disabled (default) 1 = Enabled	Brocade VMID uses the application service header field to communicate the capability.

NOTE The values in Table 4 and Table 5 are taken from the FC-GS FC 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 4 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 5 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 5.5, 6.0, and 6.5 hosts, you can manage adapters using the OneCommand Manager application on Windows, but you must install and use the appropriate Emulex CIM Provider.

NOTE If advanced adapter management capabilities are required (for example, disabling a port), 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 HBAAPI interface. To manage LightPulse adapters (including updating the firmware) on an ESXi 5.5, 6.0, or 6.5 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, refer to the *Emulex OneCommand Manager for VMware vCenter for LightPulse Adapters User Guide*.

3.2.8 Installing the Emulex CIM Provider

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

3.2.9 Creating, Deleting, and Displaying VPorts

The Emulex driver for VMware supports NPIV by default. ESXi provides the only management API for creating and deleting a vPort and creating an NPIV-enabled VM. 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 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 behavior 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 Broadcom 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 8, 16, and 32 GFC LightPulse 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 and 6.5

The Emulex native mode FC and FCoE driver supports the Virtual Volumes (VVols) feature released with ESXi 6.0 and 6.5. VMware's VVols 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 number depends upon the array. One storage container can be simultaneously accessed through multiple PEs. When 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 vSphere APIs for Storage Awareness (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 for LPe16202/OCe15100 adapters in NIC+FCoE mode.

Table 6 lists the Ethernet driver module parameters and descriptions..

Table 6 NIC Ethernet Driver Parameters

Module Parameter	Description	ESXi 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 Virtual extensible LAN (VXLAN) groups)	
`emi_canceller`	Enable or disable the EMI Canceller.	Def = 0	
`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 = 32	
`msix`	Enable or disable MSI-X. 0 = Disabled. 1 = Enable.	Def = 1	For the native mode driver.
`rss`	Enable or disable RSS queues. 0 = Disabled. 1 = Enable with 4 RSS queues. 2–N = Enable with N RSS queues.	Def = 0	For the native mode driver.

In the `lpnic` driver, there is a new module parameter:

- **Enable or disable MSI-X support** – The driver default is to have MSI-X 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 lpnic
```

And to verify that the value has been reprogrammed:

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

3.3.1 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.

3.3.2 Receive-Side Scaling

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

RSS can be selectively enabled for a particular vNIC in a VM, which can be configured in the `.vmx` file. There can be VMs that require RSS and VMs that do not require RSS. RSS is disabled by default.

To use RSS, the VM must be configured properly:

- The VM must have at least four cores and 4 GB to 8 GB 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 preceding line for the vNIC in the `.vmx` file will not force the VM or ESXi to use the RSS feature, even though RSS is enabled and queues are created.

  RSS limitations include:

  — MSI-X must be enabled.

  — The number of regular NetQueues is reduced by 1 when RSS is enabled.
- RSS is supported when the ESXi host has at least 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).

LPe16202/OCe15100 adapters support up to eight NetQueues (including the default queues) and up to four RSS queues.

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 lpnic
```

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 lpnic
```

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

For ESXi 6.0 and 6.5 systems, `vsish` can be used to view the contents of the new node created for RSS in each of the uplinks (when RSS is enabled) after 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 at least one filter is placed on the RSS queue (`<qid#>`). If no filter is 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 lpnic stats get -p <<device_id>> | grep rx_compl >
```

where `<device_id>` can be obtained from the `esxcli network nic list` command.

In normal configurations, the RSS queues will be queues with IDs 8, 9, 10, and 11. 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.

3.3.3 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).

All vNICs are fully functional and support the same capabilities as a standard NIC. The vNICs can also be linked to a vSwitch 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 10GbE bandwidth.

3.3.4 Enabling SR-IOV

Single root I/O virtualization (SR-IOV) capability can be enabled for LPe16202/OCe15100 adapters if your system BIOS supports SR-IOV.

ESXi 6.0 and 6.5 fully support configuration networking options of a virtual function that is assigned to a VM with compatibility for ESXi 5.1 or later, including default VLAN Tagging, virtual guest tagging (VGT), and configuring a static MAC for a virtual function using the ESXi driver.

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

ESXi 6.0 and 6.5 have limited support for extended configuration networking options of a virtual function that is assigned to a VM with compatibility for ESXi 6.0 or later.

Supported extended SR-IOV options include:

- Enabling VGT
- Enabling VLAN Switch Tagging (VST) mode

Extended SR-IOV option limitations include:

- It cannot change the size of the maximum transmission unit (MTU) (cannot enable jumbo frames)
- It cannot accept or drop incoming frames for a new address with the MAC address change option
- It cannot enable global promiscuous mode for VM network adapters

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

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 the following command:

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

   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 32 VFs per PF for LPe16202/OCe15100 adapters.

For multiple adapters, use the following command:

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

where:

x1 is the number of VFs to be enabled on the first port of the first adapter

y1 is the number of VFs to be enabled on the second port of the first adapter

x2 is the number of VFs to be enabled on the first port of the second adapter

y2 is the number of VFs to be enabled on the second port of the second adapter

x3 is the number of VFs to be enabled on the first port of the third adapter

y3 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 lpnic
   ```

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 virtual function (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.

3.3.5 Performance Tuning

3.3.5.1 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.

3.3.5.2 Enabling TSO

Some 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 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.

3.3.5.3 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. 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**, select **Networking**.
3. Under **Networking**, select **Properties**.
4. Select the vSwitch you want to edit and select **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
```

3.3.5.3.1 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.

3.3.5.3.2 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.

3.3.5.4 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, run the following command:

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

See Table 14 for the commands to get driver private statistics.

3.4 Emulex Drivers for VMware NSX 6.2

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

3.4.1 Module Parameters

`rss`	Creates *x* RSS (receive-side scaling) queues on one of the NetQueues (*x* ≤ 8)
`drss`	Creates *y* RSS queues on the default queue (*y* ≤ 8, excluding the default queue)

3.4.2 Functionality

- RSS and DRSS (default queue receive-side scaling) are enabled by default, unless explicitly disabled, which is an NSX requirement.
- `rss = 1` – Creates four NetQueue RSS rings.
- `drss = 1` – Creates four default queue RSS rings.
- Priority order:
  - If `drss` is not specified, *AND* `rss` is not specified or is ≤ 4

    The NetQueues 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 NetQueue count is reduced to a minimum of four.

  If both NetQueue 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 7 RSS and DRSS Rings

RSS	DRSS	Number of RSS Rings Created[1]	Number of DRSS Rings Created[2]
U[3] | 1	U[c] | 1	4	4
0	0	0	0
1	0	4	0
0	1	0	4
x ≥ 2	0	x, max 8	0
x ≥ 2	U[c] | 1	x, max 8	4
U[c] | 1	y ≥ 2	4	y, max 8
0	y ≥ 2	0	y, max 8
x ≥ 2	1	x, max 8	4
x ≥ 2	y ≥ 2	x, max 8	y, max 8

1. NetQueue RSS queues include leading NetQueue.
2. DRSS rings exclude default Q.
3. Unspecified.

3.4.3 Usage

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

NOTE 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, 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 lpnic -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 lpnic -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 lpnic -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 lpnic -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 lpnic -s "RSS=8,8,8,8 DRSS=4,4,4,4"
```

NOTE In case of insufficient resources, the maximum possible RSS queues are created after reducing NetQueues to a minimum of four.

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

```
esxcfg-module lpnic -s "RSS=0,0,0,0 DRSS=0,0,0,0"
```

3.4.4 Limitations

RSS with dynamic NetQueue is not supported.

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 and FCoE Driver

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

Table 8 identifies some of the common situations and their potential resolutions.

Table 8 Troubleshooting the FC and FCoE Driver

Situation	Resolution
Port link fails to come up.	If an FC link fails to come up, verify that the adapter is connected to a supported device. The supported 8GFC adapters are: ■ 2GFC ■ 4GFC ■ 8GFC The supported 16GFC adapters are: ■ 4GFC ■ 8GFC ■ 16GFC The supported 32GFC adapters are: ■ 8GFC ■ 16GFC ■ 32GFC 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 Emulex driver is not loaded and all paths are down.	Use the `lspci` utility to determine whether the 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 log messages indicating an error. If you specified driver logging (see Section 3.2, FC and FCoE Driver Configuration), 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 Technical 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.[1] 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.[2] 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 Technical Support if you are unable to resolve missing zone members.

Table 8 Troubleshooting the FC and FCoE Driver

Situation	Resolution
The FC and FCoE driver fails to recognize an adapter and logs `unknown IOCB` 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 Broadcom website at http://www.broadcom.com.
The system panics when booted with a failed adapter installed.	Remove the failed adapter and reboot.

The FC and 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. See Section 3.2, FC and FCoE Driver Configuration, for more information on this driver parameter.
	Globally:
	`esxcli system module parameters set -p lpfc_nlp_slab_cnt=200 -m lpfc`
	Per instance:
	`esxcli system module parameters set -p lpfc0_nlp_slab_cnt=200 -m lpfc`
	A reboot is required.

1. 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
       XRIs 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
```

2. Data from KV command 2:

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

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

4.1.1 FC and FCoE Driver Log Messages

Log messages have traditionally been organized into logical groups based on code functionality in the FC or FCoE driver. With the introduction of the latest 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.log` file 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 Technical Support.

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 9 lists the groups and defines the associated number ranges.

Table 9 Message Log Table

LOG Message Verbose Mask Definition	Verbose Bit	Verbose Description
LOG_ELS	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

Table 9 Message Log Table

LOG Message Verbose Mask Definition	Verbose Bit	Verbose Description
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 trace logging
LOG_ALL_MSG	0x7fffffff	Log all messages

The following is an example of a LOG message:

```
#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 preceding LOG message:

- lpfc – Driver binary
- `lpfc_mbx_cmpl_read_topology` – The function generating the log
- 1 – Identifies 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 or 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 ID (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 node port (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 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 extended link service (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) ulpIoTag

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 remote port indicator (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: 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 received a response.

DATA: None

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

4.1.3.8 elx_mes0227: Node Authentication timeout

DESCRIPTION: The driver has lost track of which 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 which 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: The driver unloaded or a hotplug is 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 REG_LOGIN command.

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 command; 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 command for the fabric has failed.

DATA: None

ACTION: Check the port connection 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 occurred 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 the `sparam_mbox` or `cfglink_mbox` command.

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 memorysize <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 Fibre Channel Forwarder Indicator (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 memorysize <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 mailbox 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 cannot 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 <rspPutInx> 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 mailbox 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 nvparams 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) ulpIoTag, (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) ulpIoTag, (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: The driver that 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 an 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: Shut down and remove the adapter. Contact Broadcom Technical 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 an 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 fcp_cqid 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 vital product data (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: An 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 maximum 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_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.

4.1.5.36 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.

4.1.5.37 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.

4.1.5.38 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.

4.1.5.39 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.

4.1.5.40 elx_mes0451: Failed to enable interrupt

DESCRIPTION:

DATA: None.

ACTION: None required.

4.1.5.41 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.

4.1.5.42 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.

4.1.5.43 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.

4.1.5.44 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.

4.1.5.45 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.

4.1.5.46 elx_mes0474: Unable to allocate memory for issuing "MBOX_CONFIG_MSI command”

DESCRIPTION: Mailbox memory pool allocation error.

DATA: None

ACTION: None required.

4.1.5.47 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`.

4.1.5.48 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`.

4.1.5.49 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.

4.1.5.50 elx_mes0482: Illegal interrupt mode

DESCRIPTION: The driver could not set the MSI-X, MSI, or INTx interrupt modes.

DATA: None

ACTION: This error could be a server issue. Reboot. If this issue persists, report the error to Broadcom Technical Support.

4.1.5.51 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.

4.1.5.52 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.

4.1.5.53 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.

4.1.5.54 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.

4.1.5.55 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.

4.1.5.56 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.

4.1.5.57 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.

4.1.5.58 elx_mes0498: Adapter failed to init, mbxCmd <cmd> INIT_LINK, mbxStatus <status>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.5.59 elx_mes0499: Failed allocate fast-path FCP CQ (<fcp_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.

4.1.5.60 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.

4.1.5.61 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.

4.1.5.62 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.

4.1.5.63 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.

4.1.5.64 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.

4.1.5.65 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.

4.1.5.66 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.

4.1.5.67 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.

4.1.5.68 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.

4.1.5.69 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.

4.1.5.70 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.

4.1.5.71 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.

4.1.5.72 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.

4.1.5.73 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.

4.1.5.74 elx_mes0529: Failed setup of slow-path mailbox CQ: rc = <value>

DESCRIPTION: The driver failed to set up the Completion Queue. A failure code is reported.

DATA: None

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

4.1.5.75 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.

4.1.5.76 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.

4.1.5.77 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.

4.1.5.78 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.

4.1.5.79 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.

4.1.5.80 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.

4.1.5.81 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.

4.1.5.82 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.

4.1.5.83 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.

4.1.5.84 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.

4.1.5.85 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.

4.1.5.86 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.

4.1.5.87 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.

4.1.5.88 elx_mes0582: Error <rc> during sgl post operation

DESCRIPTION: The SGL post operation failed.

DATA: None

ACTION: None required.

4.1.5.89 elx_mes0602: Failed to allocate CQ_EVENT entry

DESCRIPTION: The driver failed to allocate a CQ_EVENT entry.

DATA: None

ACTION: None required.

4.1.5.90 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 <cntx> 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 nor 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 FC 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 FC 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 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 – portsmphr=<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 POST 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: SLImode <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 vPort. 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 the 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 the 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 (<alloclen>) is less than the requested DMA memory size (<reqlen>)

DESCRIPTION: The driver could not get the memory required for the number of XRIs 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_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.

4.1.14.95 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.

4.1.14.96 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.

4.1.14.97 elx_mes2598: Adapter Link is disabled.

DESCRIPTION: The adapter link is disabled.

DATA: None

ACTION: None required.

4.1.14.98 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.

4.1.14.99 elx_mes2600: failed to allocate mailbox memory

DESCRIPTION: The adapter failed to allocate mailbox memory.

DATA: None

ACTION: None required.

4.1.14.100 elx_mes2605: lpfc_dump_static_vport: memory allocation failed

DESCRIPTION: Memory allocation failed.

DATA: None

ACTION: None required.

4.1.14.101 elx_mes2606: No NPIV Fabric support

DESCRIPTION: No NPIV Fabric support.

DATA: None

ACTION: None required.

4.1.14.102 elx_mes2607: Failed to allocate init_vpi mailbox

DESCRIPTION: The adapter failed to allocate the init_vpi mailbox.

DATA: None

ACTION: None required.

4.1.14.103 elx_mes2608: Failed to issue init_vpi mailbox

DESCRIPTION: The driver failed to issue the init_vpi mailbox.

DATA: None

ACTION: None required.

4.1.14.104 elx_mes2609: Init VPI mailbox failed <u.mb.mbxStatus>

DESCRIPTION: Initialization of the VPI mailbox has failed.

DATA: None

ACTION: None required.

4.1.14.105 elx_mes2610: FCF <value> reached driver's book keeping dimension: <value>.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.14.106 elx_mes2619: Config region 23 has bad signature

DESCRIPTION: Configuration region 23 has an invalid signature.

DATA: None

ACTION: None required.

4.1.14.107 elx_mes2620: Config region 23 has bad version

DESCRIPTION: Configuration region 23 has an invalid version.

DATA: None

ACTION: None required.

4.1.14.108 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.

4.1.14.109 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.

4.1.14.110 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.

4.1.14.111 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 resets it to the default value of 1.

DATA: None

ACTION: None required.

4.1.14.112 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.

4.1.14.113 elx_mes2719: Invalid response length: tgt <TGT_ID> lun <LUN> cmnd <CMD> rsplen <RSPLEN>

DESCRIPTION: The response length for this FCP command is not supported.

DATA: None

ACTION: None required.

4.1.14.114 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.

4.1.14.115 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.

4.1.14.116 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.

4.1.14.117 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.

4.1.14.118 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 it failed to send the mailbox command due to the state of the adapter.

DATA: None

ACTION: None required.

4.1.14.119 elx_mes2745: Failed to allocate mbox for requesting FCF rediscover

DESCRIPTION: The driver is trying to rediscover the FCF table, but it 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.

4.1.14.120 elx_mes2746: Requesting for FCF rediscovery failed status <value> add_status <value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.14.121 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 it failed to send the mailbox command due to the state of the adapter.

DATA: None

ACTION: None required.

4.1.14.122 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.

4.1.14.123 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.

4.1.14.124 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.

4.1.14.125 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.

4.1.14.126 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.

4.1.14.127 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.

4.1.14.128 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.

4.1.14.129 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.

4.1.14.130 elx_mes2757: Protocol failure detected during processing of FCP I/O op: tgt <tgt ID> lun <LUN> cmnd <CMD> rspInfo3 <rspInfo3>

DESCRIPTION: The FCP response from a target indicated that the response length is valid, but rspInfo3 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.

4.1.14.131 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.

4.1.14.132 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.

4.1.14.133 elx_mes2762: FCF <value> reached driver's book keeping dimension: <value>

DESCRIPTION:

DATA: None

ACTION: Report this error to Broadcom Technical Support.

4.1.14.134 elx_mes2763: Failed to allocate mbox for READ_FCF cmd.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.14.135 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: None

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

4.1.14.136 elx_mes2772: Issue FCF rediscover mailbox command failed, failt through to FCF dead event

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.14.137 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 vPort 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 through the `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 or 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 notification message is 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 issue 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 writes 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 no node association exists.

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: lpc_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.

Table 10 identifies some of the common situations and their potential resolutions.

Table 10 Troubleshooting the NIC Driver

Issue	Resolution
Unable to connect 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, connection 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 connectivity between these VMs does not work.
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 ESXi Server NIC Event Log Entries

Table 11 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 11, <D>, <DD>, or <DDD> in the Displayed Message column refers to decimal values that appear in the error messages.

Table 11 ESXi Server NIC Event Log Entries

Severity	Displayed Message	Description
Error	OneConnect POST failed	The POST 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	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 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 GB to 8 GB 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.
Warning	WARNING: Found a OneConnect card in Gen <D> x<D> PCI-e 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	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.

Table 11 ESXi Server NIC Event Log Entries (Continued)

Severity	Displayed Message	Description
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

The `lpnic` 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 12 lists the `lpnic` driver informational log categories.

Table 12 NIC Informational Log Groups

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. See Table 11, 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 Emulex NIC PFs.

Use either of these commands:

```
esxcfg-module -s "debugMask=0x0120" lpnic
```

- Or -

```
esxcli system module parameters set -p "debugMask=0x0120" -m lpnic
```

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 Emulex NIC PF (PF with SBDF is specified with `-p`).

Command to Get:

```
esxcli lpnic dbgmask get -p 0000:03:00.1
```

Command to Set:

```
esxcli lpnic dbgmask set -p 0000:03:00.1 0x000000ff
```

4.3 Native Mode NIC Driver Troubleshooting Support

Table 13 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 13 Native Mode NIC Driver Troubleshooting Support

Troubleshooting Items from Earlier `vmklinux` Driver using the `proc` Interface	Description	Native Mode Driver Support		Usage
		(`VmkMgmtKeyVal` interface)	(`esxcli` plug-in)	Native Mode Driver

`csr_read`	Read the 32-bit register value from the CSR space at the offset set through the `csr_read` file.	X	—	`#/usr/lib/vmware/vmkmgmt_keyval/vmkmgmt_keyval -i "vmnic0-KeyValue/emulex" -k "CsrRead" -s "<offset = 0x2d90>"`
				`#/usr/lib/vmware/vmkmgmt_keyval/vmkmgmt_keyval -i "vmnic0-KeyValue/emulex" -k "CsrRead" -g`
				`Key 'CsrRead': CSR Offset:0x2d90 ==> 0x126`

`csr_write`	Set the CSR space offset for `csr_write`.	X	—	`#/usr/lib/vmware/vmkmgmt_keyval/vmkmgmt_keyval -i "vmnic0-KeyValue/emulex" -k "CsrWrite" -s "<offset = 0x2d90> <value = 3>"`
`drvr_stat`	Read the driver statistics.	—	X	Available in custom `esxcli` plug-in in the asynchronous driver: `esxcli lpniccli stats get -p <pci dev name>`
`eth_ring`	Get the various Ethernet ring properties.	—	—	Not implemented.
`misc_stat`	Read the miscellaneous counters and format them with description. These counters are clear on read.	—	Partially implemented	Error statistics available in a custom `esxcli` plug-in in the asynchronous driver: `esxcli lpniccli stats get -p <pci dev name>`
`pci_read`	Read the 32-bit register value from the PCI space at the offset set through the `pci_read` file.	X	—	`#/usr/lib/vmware/vmkmgmt_keyval/vmkmgmt_keyval -i "vmnic0-KeyValue/emulex" -k "PciRead" -s "<offset = 0x1f0>"` `#/usr/lib/vmware/vmkmgmt_keyval/vmkmgmt_keyval -i "vmnic0-KeyValue/emulex" -k "PciRead" -g`

Table 13 Native Mode NIC Driver Troubleshooting Support (Continued)

Troubleshooting Items from Earlier `vmklinux` Driver using the `proc` Interface	Description	Native Mode Driver Support		Usage
		(`VmkMgmtKeyVal` interface)	(`esxcli` plug-in)	Native Mode Driver
`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 format 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 lpniccli stats get -p <pci dev name>`
`vlan_stat`	Get the driver-specific VLAN status.	—	Not implemented	Not relevant because the driver does not use VLAN filtering.

4.4 Native Mode NIC Driver Support for ethtool Commands

Table 14 lists the ethtool support provided by the earlier vmklinux driver and the equivalent support provided through the `esxcli` and `vsish` commands.

Table 14 ESXi 5.5 and 6.0 Native Mode NIC Driver Support for ethtool Commands

`ethtool` Command	Description	`vmklinux` Driver Support (`VmkMgmtKeyVal` Interface)	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 14 ESXi 5.5 and 6.0 Native Mode NIC Driver Support for ethtool Commands (Continued)

`ethtool` **Command**	**Description**	`vmklinux` **Driver Support (**`VmkMgmtKeyVal` **Interface)**	**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 vmnic`*X*)
			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 lpnic regdump get -p <`*pcidevname*`> -f <`*filepath*`>`)
`ethtool -e|--eeprom-dump DEVNAME`	Dump device electrically erasable programmable read-only memory (EEPROM).	✓	Not implemented (`esxcli network nic eeprom dump -n vmnic`*X*)
`ethtool -E|--change-eeprom DEVNAME`	Change bytes in device EEPROM.	Not implemented	Not implemented (`esxcli network nic eeprom change -n vmnic`*X*)
`ethtool -r|--negotiate DEVNAME`	Restart N-WAY negotiation.	Not implemented	Not implemented (`esxcli network nic negotiate restart -n vmnic`*X*)
`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 vmnic`*X*)
`ethtool -S|--statistics DEVNAME`	Show adapter statistics.	✓	✓ (`esxcli network nic stats get`) The standard `esxcli` statistics expose only a limited set of statistic counters. More extensive statistics are available through the `esxcli` plug-in available in the asynchronous driver: `esxcli lpniccli stats get -p <`*pci dev name*`>`

Appendix A: esxcli Management Tool

A.1 Installing the esxcli Management Tool

The esxcli management tool is delivered as a VIB or as an offline-bundle for the `esxcli` plug-in.

To install the esxcli management tool, perform the following steps:

1. Copy the VIB or offline-bundle to the ESXi host.
2. On the ESXi host, install the VIB as follows:
   `esxcli software vib install -v=<vib with complete path> --no-sig-check`
3. Restart the hostd using the following command:
   `/etc/init.d/hostd restart`

A.2 Usage

The esxcli management tool supports both local and remote management.

- Local Management
  `esxcli elxmgmt <Command> <Parameter(s)>`
- Remote Management
  To manage the ESX system directly, use the following command:
  `esxcli-s <server> -u <username> -p <password> -d <Thumbprint> elxmgmt <Command> <Parameter(s)>`
  To managing the ESX system using a vCenter server, use the following command:
  `esxcli -s <vCenter Server> -u <username> -p <password> -h <VI_HOST> -d <thumbprint> elxmgmt <Command> <Parameter(s)>`

A.3 esxcli Management Commands

The commands listed are for local management only. You can add the remote management parameters to the local commands for remote management.

NOTE The commands are supported on LPe31000-series and LPe32000-series adapters only.

The following table contains a list of esxcli management commands and their descriptions.

Table 15 esxcli Management Commands

Command	Description
`version`	Displays the version of different components.
`listhbas`	Lists all the manageable cards in the system.
`portattributes -w <wwpn>`	Prints the port attributes of the port specified by WWPN. A failure to retrieve the port attributes is displayed as an error.

Table 15 esxcli Management Commands (Continued)

Command	Description
`portstatistics -w <wwpn>`	Displays the different statistical parameters of the port specified by WWPN.
`resetportstatistics -w <wwpn>`	Resets the port statistics on FC functions.
`hbaattributes -w <wwpn>`	Prints the adapter attributes of the port specified by WWPN. A failure to retrieve the port attributes is displayed as an error.
`getxcvrdata -w <wwpn> [-t <type>]`	Displays transceiver data in raw or formatted output.
`download -w <wwpn> -f <fully qualified path to firmware file>`	Flashes the specified firmware file on to the given adapter.
`enablebootcode -w <wwpn> -s <state>`	Enables or disables boot code on the adapter port.
`getbeacon -w <wwpn>`	Shows the current beacon state of the port (either ON or OFF).
`setbeacon -w <wwpn>`	Sets the beacon state of the port (either ON or OFF).
`getdumpdirectory`	Displays the dump directory for the adapters in the host.
`setdumpdirectory -d <DumpDirectory>`	Sets the dump directory. A dump directory has to be set before a dump can be taken. **NOTE** The dump directory must have a sub-directory under the `/vmfs/volumes` directory AND the directory has to exist at the time of setting. The directory path must not contain spaces.
`elxmgmt dump -w <wwpn>`	Create a dump file for a selected adapter. Dump files contain information, such as firmware version, driver version, and operating system information. This information is useful when troubleshooting an adapter. Text (`.txt` extension) and binary files (`.bin` extension) files are created with the `dump` command. **NOTE** The dump directory has to be explicitly set before a dump operation can succeed. If a dump operation is performed without setting a dump directory, an error message is displayed. For information on setting a dump directory, see the `setdumpdirectory` command.
`elxmgmt reset -w <wwpn>`	Sets the port speed of the FC function. **NOTE** This feature requires that the complete stack (firmware and driver) is at version 11.2 or higher. If either the firmware or the driver is earlier than version 11.2, then this command returns an error. The port speed setting must be done using VMware's system module plug-in. A reboot is required for the setting to be activated in the driver.
`reset -w <wwpn>`	Resets the FC function.
`dport -w <wwpn>`	D_Port, also called ClearLink, is a set of diagnostic tests that allows the user to detect physical cabling issues that result in increased error rates and intermittent behavior. This command is only supported for adapters connected to D_Port-enabled Brocade switches.
`targetmapping -w <wwpn>`	Displays the targets connected to a particular port along with the LUNs that are hosted by the target.
`echotest -w <initiator wwpn> -d <target wwpn> -c <number of cycles> -s <stop on error flag> [-p pattern]`	Runs the echo test on FC functions. The EchoTest command fails if the target WWPN does not support the ECHO ELS command. **NOTE** The number of cycles must be between 1 and 99999. **NOTE** The <*target wwpn*> is the WWPN of the target port and can be obtained by running the `targetmapping` command.
`getfwparams -w <wwpn>`	Gets the firmware parameters that can be set from the management layer
`setfwparam -w <wwpn>`	Sets the firmware parameters. The only firmware parameter that is currently supported is FA-PWWN (Fabric Assigned Port WWN).

Table 15 esxcli Management Commands (Continued)

Command	Description
`getdriverparams -w <wwpn>`	Displays the driver parameters associated with the WWPN specified with the `-w` flag.
`getdriverparamsglobal -w <wwpn>`	Displays the global driver parameters.
`setdriverparam -s [G,L] -t [T,P] -p <param name> -v ,<value> -w <wwpn>`	Sets the driver parameters. ■ `s` specifies the scope – G for Global and L for Local. ■ `t` specifies the nature of setting – T for Temporary and P for Permanent. ■ `p` specifies the parameter name. ■ `v` specifies the value to be set for the parameter.


BROADCOM®

Emulex® CIM Provider Package for LightPulse® Adapters

Installation Guide

Version 11.4
September 6, 2017

CIMPP-LPe-IG114-100

Corporate Headquarters
San Jose, CA

Website
www.broadcom.com

Chapter 1: Introduction

The Emulex® Common Information Model (CIM) provider enables comprehensive management of Emulex HBAs and converged fabric adapters (CFAs). It uses an industry-standard API, Common Manageability Programming Interface (CMPI) v2.0, to manage various Emulex adapters.

The CIM provider supports basic inventory and active management of the following LightPulse® FC adapters:

- LPe12000-series adapters
- LPe16000-series adapters, including LPe16202/OCe15100 adapters
- LPe31000-series adapters
- LPe32000-series adapters

During installation, the Emulex CIM provider registers with a Web-based Enterprise Management (WBEM) server running a CIM Object Manager (CIMOM) service. The CIM provider performs the following functions:

- Internally communicates with the Emulex management API.
- Internally communicates with the Emulex drivers.
- Handles inquiries and requests from various CIM clients.

1.1 Supported Platforms

The following table shows the platforms supported with the Emulex CIM provider kits.

Table 1 Emulex CIM Provider Supported Platforms

Operating Systems	Emulex CIM Provider Kits
VMware ESXi 6.0	`VMW-ESX-6.0.0-emulex-cim-provider-<`*`kit version`*`>-offline_bundle-<`*`vmware_version`*`>.zip` `vmware-esx-provider-emulex-cim-provider-6.0.0-<kit version>.vib`
VMware ESXi 6.5	`VMW-ESX-6.5.0-emulex-cim-provider-<`*`kit version`*`>-offline_bundle-<`*`vmware_version`*`>.zip` `vmware-esx-provider-emulex-cim-provider-6.5.0-<kit version>.vib`

NOTE The CIM Provider kits are available in both `.zip` and `.vib` formats.

1.2 Supported CIM Provider Profiles

1.2.1 LPe16202/OCe15100 Adapters in NIC+FCoE Mode

Access Points	Storage Networking Industry Association (SNIA) 1.3.0
Boot Control	Distributed Management Task Force (DMTF) 1.1.0
Diagnostics (CDMv2)	DMTF DSP1002
Ethernet Port	DMTF DSP1014
FCoE Initiator Ports	SNIA SMI-S 1.5 Part 2 Clause 22
Host Discovered Resources	SNIA SMI-S 1.5 Part 6 Clause 7
Host LAN Network Port	DMTF DSP1035
IP Interface	DMTF DSP1036
PCI Device	DMTF DSP1075
Physical Assets Profile	DMTF DSP1011
Profile Registration	DMTF DSP1033
Record Log	DMTF DSP1010
Software Inventory	DMTF DSP1023
Software Update	DMTF DSP1025
Storage HBA	SNIA SMI-S 1.5 Part 6 Clause 6

1.2.2 LPe12000-Series, LPe16000-Series, LPe31000-Series, and LPe32000-Series Adapters

Profile Registration	DMTF DSP1033
Software Inventory	DMTF DSP1023
Physical Package	SNIA SMI-S 1.5 Part 2 Clause 31
Host Discovered Resources	SNIA SMI-S 1.5 Part 6 Clause 7
Storage HBA	SNIA SMI-S 1.5 Part 6 Clause 6
Software Update	DMTF DSP1025
Record Log	DMTF DSP1010
Software	SNIA 1.2.0
Access Points	SNIA 1.3.0
PCI Device	DMTF1.0.0
Physical Asset Profile	DMTF DSP1011
Boot Control	DMTF 1.1.0
FC HBA	SNIA SMI-S 1.5 Part 6 Clause 5
FC HBA Diagnostic Profile	DMTF DSP1104
FC Initiator Ports Profile	SNIA SMI-S 1.5 Part 2 Clause 17

1.3 Abbreviations

API	Application Programming Interface
CIM	Common Information Model
CIMOM	CIM Object Manager
CMPI	Common Manageability Programming Interface
DMTF	Distributed Management Task Force
FC	Fibre Channel
FCoE	Fibre Channel over Ethernet
HBA	host bus adapter
IP	Internet Protocol
LAN	local area network
NIC	network interface card (or controller)
PCI	Peripheral Component Interconnect
SFCB	Small Footprint CIM Broker
SMI	Storage Management Initiative
SNIA	Storage Networking Industry Association
URI	Uniform Resource Identifier
VIB	vSphere Installation Bundle
vNIC	virtual network interface card
WBEM	Web-based Enterprise Management

Chapter 2: Installing the Emulex CIM Provider

The following items must be installed before you can install the Emulex CIM provider.

- One of the following adapters:
  - LPe12000-series adapter
  - LPe16000-series adapter, including LPe16202/OCe15100 adapters
  - LPe31000-series adapter
  - LPe32000-series adapter
- The appropriate adapter drivers

NOTE Adapters on an ESXi host running Emulex CIM providers can be managed by a Windows server using the following applications (installed on Windows operating systems):

- The Emulex OneCommand® Manager application for Windows
- The Emulex OneCommand Manager application for VMware vCenter

Go to the download page on the Broadcom website, at http://www.broadcom.com, or to the vendor website to verify the driver version or the OneCommand Manager application version that must be installed on your system.

2.1 Installing the VIB in the VMware Operating System

Use one of the standard `esxcli` commands to install the VIB or the offline bundle.

- To install the VIB, type:

```
esxcli software vib install -v <provider.vib> --maintenance-mode
```

- To install the signed offline bundle, type:

```
esxcli software vib install -d <offline-bundle.zip> --maintenance-mode
```

- To install an unzipped file, type:

```
esxcli software vib install --viburl=<file:/vmware-esx-provider-emulex.vib>
--maintenance-mode
```

2.2 Using Adapters

This section describes updating firmware, enabling logs, running diagnostics, and discovering vNICs on Emulex adapters.

If you require additional information, contact an authorized Broadcom Technical Support representative at ecd-tech.support@broadcom.com or request assistance online at https://oemsupportportal.emulex.com/web2tech/ecd.html.

2.2.1 Updating Firmware on Emulex Adapters

Use the Software Update profile to update the firmware on Emulex adapters. The following methods are implemented in the Emulex-specific Software Update profile:

- **Install from Byte Stream** – Requires a custom CIM client that can read the firmware file and create a ByteStream used to update the firmware.
- **Install from URI** – Supports two different types of URI:
  - The firmware file to be updated is available locally on the machine hosting the Emulex adapter.
  - The firmware file to be updated is available on a remote machine, such as an http or https server. In this case, the Emulex CIM Provider uses the `libcurl` library available on the host machine, where the CIM provider is running, to download the firmware file.

The Emulex `CIM_SoftwareInstallationServiceCapabilities` class has the attribute `SupportedURISchemes` that identifies the supported URI schemes.

To update the firmware, perform these steps:

1. List the `CIM_SoftwareInstallationServices` in the Emulex namespace. Select the `CIM_SoftwareInstallationService` specific to the adapter on which the firmware is to be updated. For example, if an LPe16202/OCe15100 adapter needs a firmware update, select the `ECDUCNA_SoftwareInstallationService` class instance:

   ```
   wbemcli  -noverify ein
   'https://root:<password>@<IP>/root/emulex:ecducna_softwareinstallationservi
   ce' -nl
   ```

2. List the `CIM_EthernetPorts` for LPe16202/OCe15100 adapters, or list the `CIM_FCPorts` for all other LightPulse adapters. Select the desired port from the listed instances. For example, for LPe16202/OCe15100 adapters, type:

   ```
   wbemcli  -noverify ein
   'https://root:<password>@<IP>/root/emulex:ecducna_ethernetport' -nl
   ```

   For all other LightPulse adapters, type:

   ```
   wbemcli  -noverify ein
   'https://root:<password>@<IP>/root/emulex:elxhba_fcport' -nl
   ```

3. Run `InstallFromURI` using the output from step 1 and step 2.

   For example:

   ```
   wbemcli cm -noverify 'https://root:<password>@<IP>/root/emulex:<Output of
   step 1>
   InstallFromURI.URI=<Full path of fw file>,Target=<Output of step 2>
   ```

2.2.2 Enabling Logs and Collecting Symptoms

To enable provider logs, perform these steps:

1. List either the `ELXHBA_RecordLog` (for an FC HBA) or the `ECDUCNA_RecordLog` (for an LPe16202/OCe15100 adapter in FCoE+NIC mode) class and note the provider log instance. If the list command fails, perform the steps detailed in Section 2.2.3, Generating Provider Logs if Listing Them Fails.
2. Set the provider log settings:

   ```
   wbemcli -noverify cm 'https://root:<password>@<IP>/root/emulex:<Instance
   from step 1> ' SetLogParams.LogLevel=5,logmode=2,tracepath='<file path>'
   ```

3. Set the provider log state:

   ```
   wbemcli -noverify cm 'https://root:<password>@<IP>/root/emulex:<Instance
   from step 1> ' RequestStateChange.RequestedState=2
   ```

4. Perform the operation that is not working as expected. The provider logs are available in the *<file path>* specified in step 2.
5. Contact a Broadcom Technical Support representative with the provider logs and the system logs (`/var/log/syslog.log`).

To disable provider logs, enter the following command:

```
wbemcli -noverify cm 'https://root:<password>@<IP>/root/emulex:<Instance
from step 1> ' RequestStateChange.RequestedState=3
```

2.2.3 Generating Provider Logs if Listing Them Fails

To enable provider logs if listing them fails, perform these steps:

1. Stop the CIMOM.
2. Create a `.dmp` file.
   a. To create a provider log for the Emulex FC provider, create an `emulex_fc_provider.dmp` file in the `/etc/cim/emulex` location.
   b. To create a provider log for an LPe16202/OCe15100 adapter in FCoE+NIC mode, create an `emulex_ucna_provider.dmp` file in the `/etc/cim/emulex` location.

   NOTE If the FC and NIC+FCoE logs must be created together, use different file names for each dmp file. For example:
   - `/tmp/fc.txt`
   - `/tmp/nicfcoe.txt`
3. Manually enter the following two lines (without spaces):

   ```
   2,5,1
   /tmp/providerlogs.txt
   ```
4. Start the CIMOM and list the Emulex classes in the `root/emulex` namespace.
5. Perform the operation that is not working as expected. The provider logs are available in *</tmp/providerlogs.txt>* specified in step 3.
6. Contact a Broadcom Technical Support representative with the provider logs and the system logs (`/var/log/syslog.log`).

   NOTE Dead dump is not supported.

2.2.4 Running Diagnostic Tests

Running a Diagnostic Test on an LPe16202/OCe15100 Adapter in NIC+FCoE Mode

To run a diagnostic test on the Emulex CIM provider, perform these steps:

1. Get the instance of the managed element (Ethernet port). For example:

   ```
   wbemcli  -noverify ein
   'https://root:<password>@<IP>/root/emulex:ecducna_ethernetport' -nl
   ```
2. Get the instance of the `ECDUCNA_DiagnosticTest` class. For example:

   ```
   wbemcli  -noverify ein
   'https://root:<password>@<IP>/root/emulex:ecducna_diagnostictest' -nl
   ```

3. Run the RunDiagnosticservice on the ecducna_diagnostictest. For example:

```
wbemcli -noverify cm 'https://root:<password>@<IP>/root/emulex:<Diagnostic
Test instance from step 2>'
RunDiagnosticservice.ManagedElement=<ManagedElement instance from step 1>'
```

A CIM_ConcreteJob instance is created for each diagnostic test run. For example:

```
wbemcli  -noverify ein
'https://root:<password>@<IP>/root/emulex:ecducna_concretejob' -nl
```

Results of the diagnostic test runs are available in ECDUCNA_DiagnosticCompletionRecord class instances. For example:

```
wbemcli  -noverify ein
'https://root:<password>@<IP>/root/emulex:ecducna_diagnosticcompletionrecor
d' -nl
```

The diagnostic logs can be cleared using the ClearLog function of ECDUCNA_DiagnosticsLog class. For example:

```
wbemcli -noverify cm
'https://root:<password>@<IP>/root/emulex:<Corresponding Diagnostic log
instance>' ClearLog
```

2.2.4.1 Running a Diagnostic Test on an Emulex FC Adapter

To run a diagnostic test on an FC adapter, perform these steps:

1. Get the instance of the managed element (ELXHBA_PortController). For example:

```
wbemcli  -noverify ein
'https://root:<password>@<IP>/root/emulex:elxhba_portcontroller' -nl
```

2. Get the instance of the ELXHBA_FCHBADiagnostictest class. For example:

```
wbemcli  -noverify ein
'https://root:<password>@<IP>/root/emulex:elxhba_diagnostictest' -nl
```

3. Run the RunDiagnosticservice on the ELXHBA_FCHBADiagnostictest. For example:

```
wbemcli -noverify cm 'https://root:<password>@<IP>/root/emulex:<Diagnostic
Test instance from step 2>'
RunDiagnosticservice.ManagedElement=<ManagedElement instance from step 1>'
```

A CIM_ConcreteJob instance is created for each diagnostic test run. For example:

```
wbemcli  -noverify ein
'https://root:<password>@<IP>/root/emulex:elxhba_concretejob' -nl
```

Results of the diagnostic test runs are available in ELXHBA_DiagnosticCompletionRecord class instances. For example:

```
wbemcli  -noverify ein
'https://root:<password>@<IP>/root/emulex:elxhba_diagnosticcompletionrecord
' -nl
```

The diagnostic logs can be cleared using the ClearLog function of ELXHBA_DiagnosticsLog class. For example:

```
wbemcli -noverify cm
'https://root:<password>@<IP>/root/emulex:<Corresponding Diagnostic log
instance>' ClearLog
```

2.2.5 Discovering vNICs (LPe16202/OCe15100 Adapters Only)

You can discover vNICs in three ways:

- Determine vNICs using the CIM provider
- Determine vNICs hosted by the adapter in FCoE+NIC mode in a specific host
- Determine vNICs operating over the same physical port

2.2.5.1 Determining vNICs Using the CIM Provider

1. Find all the `CIM_VLANEndpoint` instances whose `SystemName` property has a value equal to the name of the desired host. These instances are vNIC endpoints.
2. Find the `CIM_EndpointIdentity` instance associating the `CIM_VLANEndpoint` instance to a `CIM_LANEndpoint` instance. Follow the association to the `CIM_LANEndpoint`.
3. Find the `CIM_DeviceSAPImplementation` instance that associates the `CIM_LANEndpoint` to an instance of `CIM_EthernetPort`. Follow the associations to the `CIM_EthernetPort` instance representing the vNIC.

2.2.5.2 Determining vNICs Hosted by the Adapter in FCoE+NIC Mode in a Specific Host

1. Find the `CIM_Card` instance for the adapter in FCoE+NIC mode.
2. Find the vNICs in the host as detailed earlier, based on the version of the CIM provider.
3. For each `CIM_EthernetPort` representing a vNIC, perform these steps:
   a. Find the `CIM_ControlledBy` instance associating the `CIM_EthernetPort` to a `CIM_PortController` instance. Follow the association to the `CIM_PortController` instance.
   b. Find the `CIM_Realizes` instance associating the `CIM_PortController` to a `CIM_PhysicalConnector` instance. Follow the association to the `CIM_PhysicalConnector` instance.
   c. Find the `CIM_Container` instance associating the `CIM_PhysicalConnector` to a `CIM_Card` instance.

      Each `CIM_EthernetPort` with a CIM_Container instance referencing the CIM_Card instance found in step 1 is hosted by the targeted NIC+FCoE adapter.

2.2.5.3 Determining vNICs Operating Over the Same Physical Port

1. Find the `CIM_PhysicalConnector` instance representing the port in question. The key value for `CIM_PhysicalConnector` instances contains the serial number of the adapter hosting the port and the port number assigned to the physical port.
2. Find the vNICs in the host as detailed earlier, based on the version of the CIM provider. For each `CIM_EthernetPort` representing a vNIC:
   a. Find the `CIM_ControlledBy` instance associating the `CIM_EthernetPort` to an instance of `CIM_PortController`. Follow the association to the `CIM_PortController` instance.
   b. Find the `CIM_Realizes` instance associating the `CIM_PortController` to an instance of `CIM_PhysicalConnector`.

      Each `CIM_EthernetPort` with a `CIM_Realizes` instance referencing the `CIM_PhysicalConnector` instance is running over the targeted physical port.

Chapter 3: Troubleshooting

The following error message might appear if the CIM hosts are not properly added to the OneCommand Manager application:

```
"Unknown or invalid host specified"
```

There could be instances in which the drivers, the CIM provider, and the CIM Client on a Windows machine are all properly installed, but the CIM hosts are still not added to the OneCommand Manager application. The following table shows the most common reasons for this problem.

Table 2 Problems Adding a CIM Host

Situation	Resolution
The machine with the specified IP is not reachable.	Verify whether the machine is reachable from the CIM Client.
The specified protocol (HTTP or HTTPS) is not supported by the CIMOM.	Most often the CIMOM is configured to use HTTPS. Therefore, if you are trying to connect with HTTP, you might get an error. Try using HTTPS instead.
The namespace is invalid.	Verify that the namespace for the Emulex provider is `root/emulex`.
The username or password is invalid.	Verify that the username is correct, and retype the password.

The CIMOM is not running on the ESXi host.	You can check whether the CIMOM, a small footprint CIM broker (SFCB), is running by typing one of these commands.
	`/etc/init.d/sfcbd-watchdog status`
	or
	`ps -ef | grep sfcb`
	If the CIMOM is listening to a port other than 5988 or 5989, the connection might not take place. You can configure the SFCB CIMOM settings by editing `/etc/sfcb/sfcb.cfg`.
	By default, `sfcb` is enabled on ESXi 6.5. That is by default, in `/etc/sfcb/sfcb.cfg`:
	`enabled: true`

The CIM provider is running, but enumerations are not occurring properly.	Verify that the correct CIM provider for the ESXi host is installed.

If you still experience problems when adding the host, run the following commands on the ESXi host and send the output to the Broadcom Technical Support team.

```
vm-support
esxcfg-module -l
esxcfg-scsidevs -a
esxcfg-nics -l
lspci
esxcli software vib list | grep -i elx
esxcli software vib list | grep -i lpfc
esxcli software vib list | grep -i emu
esxcli software vib list | grep -i lpnic
```

Send the `/var/log/syslog.log` file for all of the above operations.


BROADCOM®

Emulex® OneCommand® Manager for VMware vCenter for LightPulse® Adapters

User Guide

Version 11.4
September 6, 2017

OCM-VM-LPe-UG114-100

Corporate Headquarters
San Jose, CA

Website
www.broadcom.com

Chapter 1: Introduction

Emulex® OneCommand® Manager for VMware vCenter is a comprehensive management utility for Emulex adapters that provides a powerful, centralized adapter management suite for the VMware vCenter management console.

This comprehensive solution builds upon standard Emulex Common Information Model (CIM) management providers and advanced functionality delivered with the OneCommand Manager application to present native configuration management, status monitoring, and online maintenance of Emulex adapters in VMware ESXi environments, using a graphical interface (GUI) or a command line interface (CLI).

1.1 Compatibility

For supported versions of operating systems, platforms, and adapters, go to http://www.broadcom.com.

OneCommand Manager for VMware vCenter can be hosted and registered on the following servers:

- vCenter Server 6.5
- vCenter Server 6.0
- vCenter Server 5.5

OneCommand Manager for VMware vCenter is supported for the following Windows operating systems:

- Windows 7 x64
- Windows 8 x64
- Windows 8.1 x64
- Windows Server 2012
- Windows Server 2016

See Table 1 to determine the support provided by the CIM Provider.

NOTE Illustrations in this guide are for illustrative purposes only. Your system information can vary.

This product supports the following Emulex LightPulse® host bus adapters (HBAs) and converged fabric adapters (CFAs):

- LPe12000-series adapters
- LPe15000-series adapters
- LPe16000-series adapters, including LPe16202/OCe15100 adapters
- LPe31000-series adapters
- LPe32000-series adapters

Table 1 lists the Emulex OneCommand Manager application support provided by the Emulex CIM Provider package and the individual Network Interface Card (NIC) and Fibre Channel/Fibre Channel over Ethernet (FC/FCoE) providers in

each package. The Emulex Common Information Model (CIM) Provider packages can be downloaded from http://www.broadcom.com.

Table 1 Support Provided by Emulex CIM Provider Versions

Emulex OneCommand Manager Application Features	ELX CIM Provider Package v11.4		OneCommand Manager for VMware vCenter v 11.4
	NIC	FC/FCoE	
Discover virtual ports connected to an FC or FCoE port	x	x	x
View virtual port information in a cluster (host-centric mode)	N/A	N/A	x
Discover hosts, adapters (FC, FCoE, and NIC), targets, and logical unit numbers (LUNs) for selected ESXi hosts	x	x	x
Discover hosts and adapters (FC and FCoE) for selected ESXi fabrics	x	x	x
View firmware version	x	x	x
View boot code version	x	x	x
Update firmware and boot code on a single adapter	x	x	x
Update firmware and boot code on a per-fabric basis	N/A	N/A	x
Change World Wide Port Name (WWPN) or World Wide Name (WWN) for FC and FCoE adapters	x	x	x
Locate adapters with beaconing	x	x	x
View PCI Express (PCIe) registers	x	x	x
D_Port (also referred to as ClearLink) test, for LPe16000, LPe31000, and LPe32000-series FC adapters connected to D_Port-enabled Brocade switches only	N/A	x	x
PCI loopback test	N/A	x	x
Internal and external loopback test	N/A	x	x
Echo test for LPe12000-series adapters	N/A	x	x
Power-on self-test (POST) test for LPe12000-series adapters	x	N/A	x
Batch update firmware and boot code to multiple adapters	N/A	N/A	x
Enable and disable ports	x	x	x
Get driver parameters (global and port)	x	x	x
Set global driver parameters to adapters	x	x	x
Set port driver parameters to adapters	x	x	x
Target and LUN information	x	x	x
ExpressLane ^{TM} priority queuing	x	x	x
ExpressLane LUN Level Frame Priority (FC Only)	N/A	x	x
Reset Port (FC and FCoE only)	x	x	x
View vital product data (VPD)	x	x	x
Enable and disable FCoE Initialization Protocol (FIP) data	x	x	x
Data Center Bridging (DCB) configuration for NIC + FCoE ports for LPe16202/OCe15100 adapters in NIC+FCoE mode	x	+	x
Display flash contents (wakeup parameters and flash load list) for FC ports	N/A	x	x
Export storage area network (SAN) configuration information at cluster level and at host level	N/A	N/A	x
Perform diagnostic dump at the adapter and port levels	x	x	x

1.2 Abbreviations

AL_PA	Arbitrated Loop Physical Address
BIOS	basic input/output system
BOFM	Blade Open Firmware Management Protocol
CA	certificate authority
CIM	Common Information Model
CIN	Cisco, Intel, Nuova (Data Center Bridging Exchange)
CLI	command line interface
CLP	Command Line Protocol
CNA	Converged Network Adapter
CRC	cyclic redundancy check
CSR	Certificate Signing Request
CSV	comma-separated values
DCB	Data Center Bridging
DCBX	Data Center Bridging Capabilities Exchange
DHCP	Dynamic Host Control Protocol
DNS	Domain Name System or Domain Name Server
DOCSIS	Data Over Cable Service Interface Specification
EC	engineering change
ETS	Enhanced Transmission Selection
FC	Fibre Channel
FCF	Fibre Channel over Ethernet Forwarder
FCoE	Fibre Channel over Ethernet
FCP	Fibre Channel Protocol
FIP	FCoE Initialization Protocol
FPMA	fabric-provided MAC address
FW	firmware
Gb	gigabit
Gb/s	gigabits per second
GFO	Get Fabric Object
GUI	graphical user interface
HBA	host bus adapter
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext Transfer Protocol Secure
IEEE	Institution of Electrical and Electronics Engineers
I/O	input/output
IP	Internet Protocol
IPL	initial program load
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
JEDEC ID	Joint Electron Device Engineering Council identification code

KB	Kilobyte (1024 bytes)
LAN	local area network
LIP	Loop Initialization Primitive
LLDP	Link Layer Discovery Protocol
LUN	logical unit number
MAC	Media Access Control
MB	megabyte
Mb	megabit
Mb/s	megabits per second
MN	manufacturer ID
MTU	maximum transmission unit
N/A	not applicable
NIC	network interface card
NOS	network operating system
NVRAM	non-volatile random-access memory
OAS	Optimized Access Storage
OS	operating system
OUI	organizationally unique identifier
PCI	Peripheral Component Interconnect (interface)
PCIe	PCI Express
PFC	priority flow control
PGID	priority group ID
POST	power-on self-test
PXE	Preboot Execution Environment
QoS	quality of service
RFC	request for comments
Rx	receive
SAN	storage area network
SCSI	Small Computer System Interface
SFCB	Small Footprint CIM Broker
SFP	small form-factor pluggable
SLI	Service Level Interface
SR-IOV	single root I/O virtualization
SSL	secure sockets layer
TCP	Transmission Control Protocol
TCP/IP	TCP over Internet Protocol
Tx	Transmit
ULP	Upper Layer Protocol
URL	Uniform Resource Locator
vCSA	VMware for vCenter Server Appliance
VF	virtual function
VLAN	virtual local area network
VLAN ID	VLAN identifier

VM	virtual machine
VMID	virtual machine ID
VPD	vital product data
vPort	virtual port
VM UUID	VM universal unique identifier
WLAN	wireless LAN
WWN	World Wide Name
WWNN	World Wide Node Name
WWPN	World Wide Port Name
XML	Extensible Markup Language

Chapter 2: Installing and Enabling OneCommand Manager for VMware vCenter

This section describes installing, uninstalling, reinstalling, and enabling OneCommand Manager for VMware vCenter. This section also describes registering and unregistering OneCommand Manager for VMware vCenter.

NOTE OneCommand Manager for VMware vCenter provides real-time management as a plug-in through the VMware vSphere console and VMware Web Client.

The web client allows you to run OneCommand Manager for VMware vCenter from the vSphere Web Vlient. System performance is directly influenced by the speed and efficiency of the underlying network infrastructure.

2.1 Best Installation Practices

- To avoid network latencies, ensure that all the managed ESXi hosts, the vCenter Server, and the VMware Infrastructure (VI) client are in the same network.
- To avoid configuration conflicts, install the vCenter Server and the VI infrastructure client on different hosts within the same network.

2.2 Hardware Requirements

- Physical or virtual (x86 or x86_64) servers with a minimum RAM of 2 GB and 250 GB of disk space.

2.3 Software Requirements

- VMware ESXi 5.5, 6.0, and 6.5 environments
- Operating system – Windows 7 (64 bit), Windows 8 (64 bit), Windows Server 2012 (64 bit), and Windows Server 2012 R2, and Windows Server 2016
- Adobe Flash Player 11.2 or later
- vCenter Server 5.5, 6.0, and 6.5
- vSphere Client 5.5, 6.0, and 6.5
- vSphere Web Client to manage the VMware vCenter Server instead of the vSphere client

  NOTE On the system where OneCommand Manager for VMware vCenter is installed, make sure the port numbers configured during the installation are open and dedicated to the OneCommand Manager for VMware vCenter Server only. No other service should be listening on this port.
- Emulex CIM Provider Package version 11.4

  NOTE Version 11.4 packages are not compatible with the 11.1 or earlier versions of Emulex software.

- Driver and firmware requirements

  Go to http://www.broadcom.com for the latest compatible driver and firmware versions.

2.4 Installing OneCommand Manager for VMware vCenter

The Emulex CIM Provider must be installed on your ESXi host before installing OneCommand Manager for VMware vCenter. For more information on installing the CIM Provider, refer to the *CIM Provider Package Installation Guide for LightPulse Adapters* available on http://www.broadcom.com.

NOTE As part of the OneCommand Manager for VMware vCenter installation, the OneCommand Manager for VMware vCenter Web Client components are installed and registered.

To install OneCommand Manager for VMware vCenter in Windows, perform these steps:

1. Go to http://www.broadcom.com to download the `ELXOCM-VMware-vCenter-<version>-Setup.exe` installation file to your system.
2. Navigate to the system directory to which you downloaded the file.
3. Double-click **ELXOCM-VMware-vCenter-<*version*>-Setup.exe**.

   The **OneCommand Manager for VMware vCenter** window appears.
4. Click **Next**.

   The **License agreement** window appears.
5. Read the agreement and select the **I agree to these terms and conditions** check box.
6. Click **Next**. The **Installation options** window with the default Installation folder appears (Figure 1).

Figure 1 Installation Options Window

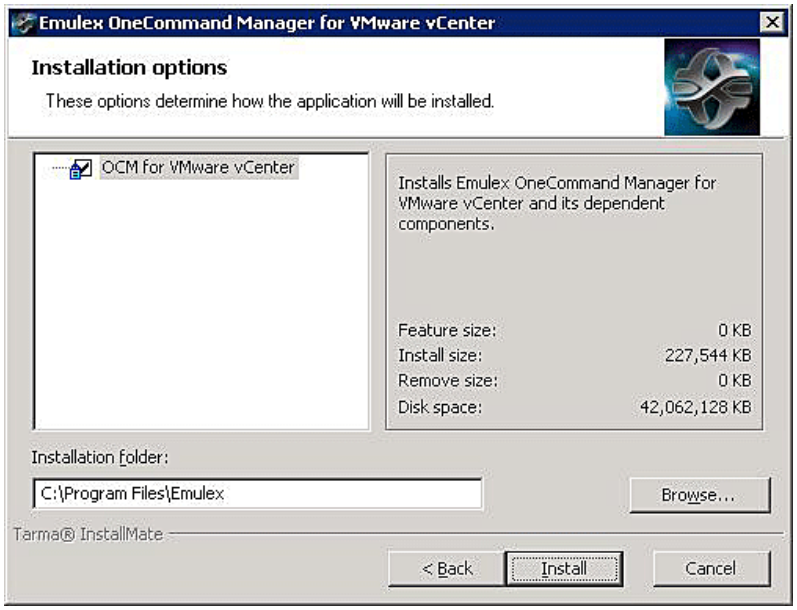

7. Ensure that the **OCM for VMware vCenter** check box is selected.
8. Program files install by default to `C:\Program Files\Emulex`. To change this location, click **Browse** and navigate to where you want the program files to reside.

9. Click **Install**. The **Operation in progress** screen appears. When the process is complete, the **OCM for VMware vCenter configuration** window appears (Figure 2).

Figure 2 OCM for VMware vCenter configuration Window

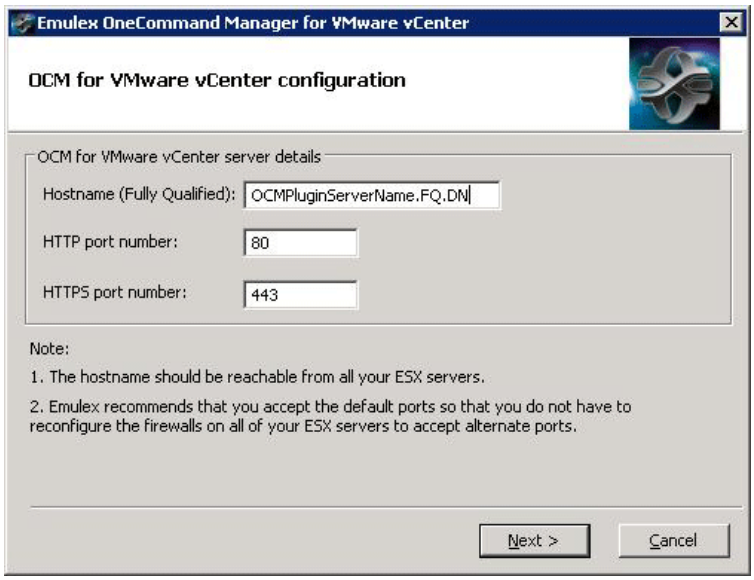

Default port numbers for OneCommand Manager for VMware vCenter Server are displayed.

NOTE The Windows firewall setting must allow incoming connections on the HTTP and HTTPS ports that you configure here.

If the port numbers are already in use, a warning appears next to the port number (Figure 3).

Figure 3 OCM for VMware vCenter Configuration Window with Port in use Warning

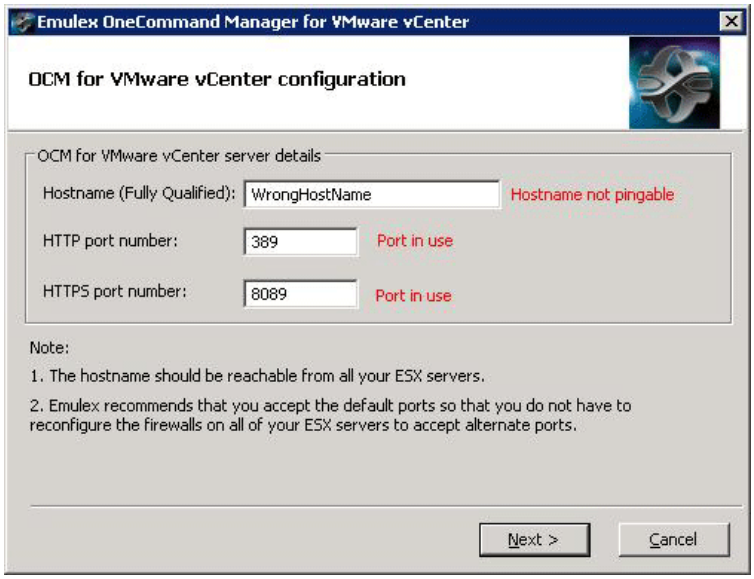

10. Click **Next**. The **Operation in progress** screen appears. When the installation process is complete, a message prompts you to launch the registration utility.
11. Click **Yes**. The **Register/Unregister** dialog appears in a new browser window (Figure 4).

Figure 4 Register/Unregister Dialog

12. Enter the following details of the vCenter Server:
    - vCenter Server Name – The IP address of the vCenter Server.
    - vCenter Server HTTPS Port – The HTTPS port number of the vCenter Server.

      NOTE The vCenter Server HTTPS port is 443 by default. You can change this value if you have configured a different HTTPS port while installing the vCenter.
    - Username – The user name with required privileges.
    - Password – The user password.
13. Click **Register** to register OneCommand Manager for VMware vCenter with a new vCenter Server.

    NOTES
    - You can unregister an existing OneCommand Manager for VMware vCenter by clicking **Unregister**.
    - If you change the host name of the machine hosting the vCenter Server, you must re-install the vCenter Server and re-register.
14. When the operation is successful, a message is displayed. Click **OK**.
15. Close the browser window. The **Installation completed** window appears.
16. Click **Finish**. The **OneCommand Manager for VMware vCenter Registration** icon is created on the desktop. You do not need to reboot the system.

2.4.1 Verifying the OneCommand Manager for VMware vCenter Installation

To verify the OneCommand Manager for VMware vCenter installation, perform these steps:

1. Log on to the vCenter Server through the VMware vSphere console.
2. Enter the IP address and credentials of the vCenter Server where OneCommand Manager for VMware vCenter is registered.

   The vSphere console appears (Figure 5).

Figure 5 vSphere Console, Getting Started Tab

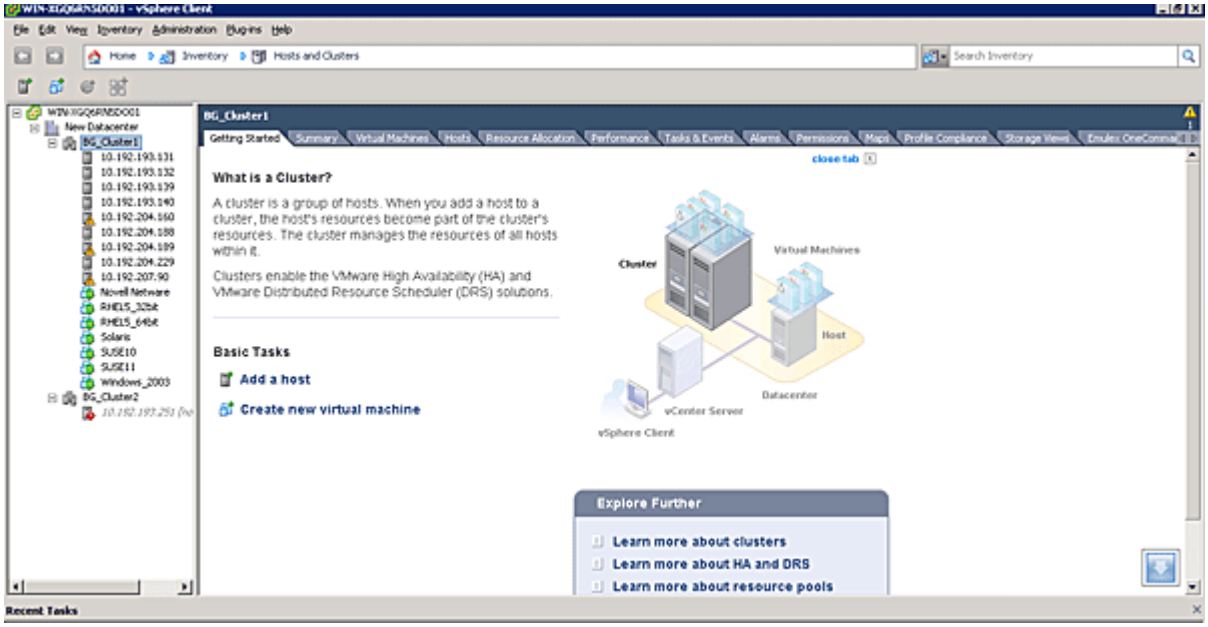

3. In the vSphere console, select **Plug-ins** from the menu bar and select **Manage Plug-ins**. The **Plug-in Manager** window appears (Figure 6).

Figure 6 Plug-in Manager Window

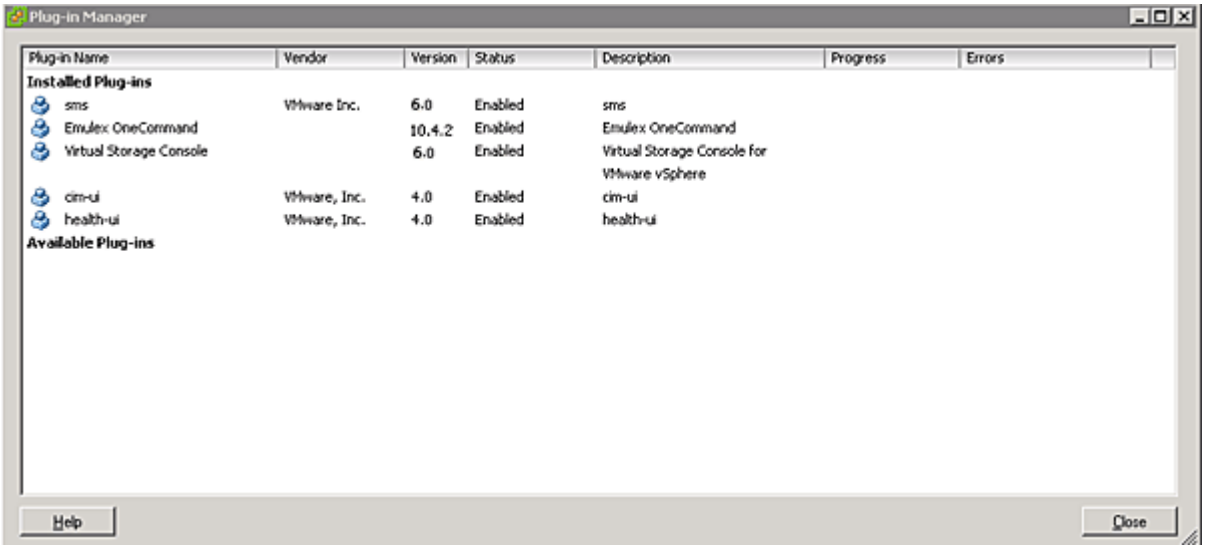

4. In the **Plug-in Manager** window, note the status of OneCommand Manager for VMware vCenter (Emulex OneCommand). If the OneCommand Manager for VMware vCenter installation is complete, the status of Emulex OneCommand is enabled by default.

2.5 Activating and Enabling ESXi Management

This section describes activating OneCommand Manager for VMware vCenter in the web client, enabling OneCommand Manager for vCSA, and enabling OneCommand Manager for VMware vCenter in the web client.

NOTE Refer to the vSphere guide on the VMware website for information on creating a user with required privileges and changing access permissions for a user in the Active Directory.

2.5.1 Requirements

Only a user with these specific privileges can read and perform active management in OneCommand Manager for VMware vCenter:

- **Extension.Register extension** to register OneCommand Manager for VMware vCenter using the registration utility.
- **Extension.Unregister extension** to unregister OneCommand Manager for VMware vCenter using the registration utility.
- **Host.CIM.CIM Interaction** to read and manage data through the OneCommand Manager for VMware vCenter

All other users, including the root user, of the ESXi host cannot perform any actions including reading data. If a user without the required privileges attempts to perform an action in OneCommand Manager for VMware vCenter, an error message is displayed.

NOTE To configure user roles and assign privileges, refer to the *VMware vCenter Server Guide* on the VMware website.

2.5.2 Lockdown Mode Feature

Refer to the vSphere guide on the VMware website for information on enabling and disabling lockdown mode.

If lockdown mode is enabled for an ESXi host, only a user with the required privileges can access the ESXi host and manage adapters using OneCommand Manager for VMware vCenter. All other users, including the root user, do not have access to the ESXi host.

2.5.3 Activating OneCommand Manager for VMware vCenter in the Web Client

To activate OneCommand Manager for VMware vCenter, which allows it to be accessed from the **Manage** tab, you must modify the `webclient.properties` file on the server where the vSphere web client is installed.

1. Stop the vSphere web client service.
2. Locate the `webclient.properties` file in the vSphere Web Client install directory, which is typically the following:

   ```
   %ProgramData%\VMware\vSphere Web Client
   ```

3. Add the following line:

   ```
   scriptPlugin.enabled = true
   ```

4. Save and close the file.
5. Restart the vSphere web client service.

2.5.4 Enabling OneCommand Manager for VMware for vCSA

vCSA is a preconfigured Linux-based virtual machine that runs the vCenter Server and associated services. vCSA is an alternative for centralized management of VMware vSphere instead of Windows vCenter Server.

To enable OneCommand Manager for VMware vCenter on vCSA, perform these steps:

1. Edit `webclient.properties` file at:

   `/var/lib/vmware/vsphere-client/webclient.properties`

   and add the statement:

   `scriptPlugin.enabled = true`
2. Save the file.
3. Restart the vSphere Web Client service on the web console.
4. Log on to vCSA again.

   The **Manage** tab for the OneCommand Manager for VMware vCenter is now displayed.

2.5.5 Enabling OneCommand Manager for VMware vCenter Web Client

OneCommand Manager for VMware vCenter can be enabled or disabled from the Plug-In Management section in the vSphere Web Client.

NOTE You must have sufficient privileges to access the Plug-In Management section. Refer to the VMware documentation for information on configuring users and privileges.

To enable OneCommand Manager for VMware vCenter in the vSphere Web Client, perform these steps:

1. From the navigation panel on the left-side of the vSphere Web Client Home page, click **Plugins > Manage Plug-ins**. The **Plug-in Manager** page is displayed (Figure 7).

Figure 7 vSphere Web Client Plug-In Management Page

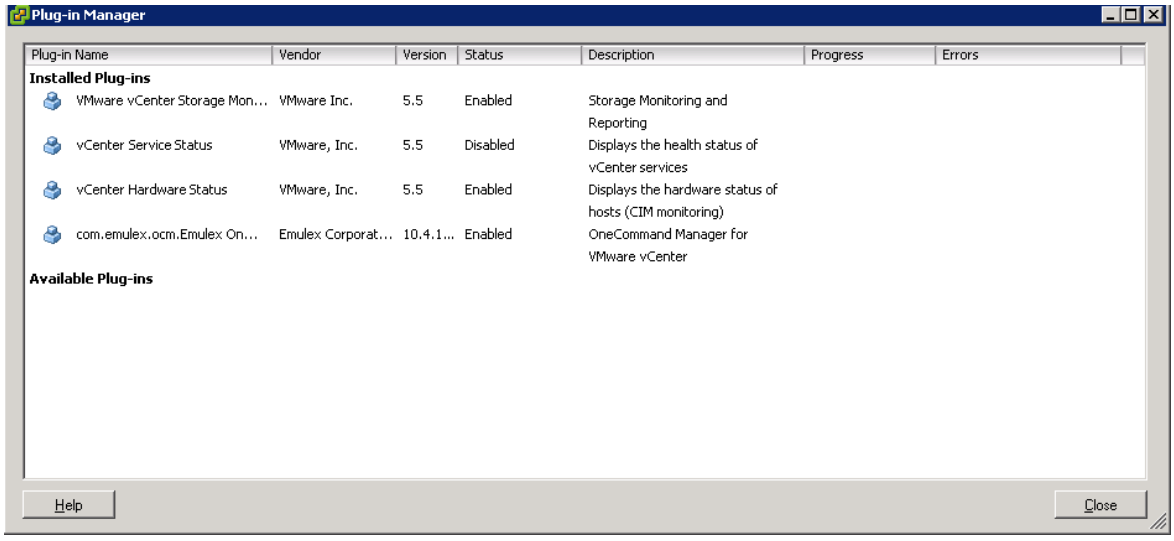

2. In the Plug-In Manager page, select **com.emulex.Emulex OneCommand** and right-click under the **Status** column.

   A context menu opens.

3. Either select **Enable** to enable OneCommand Manager for VMware vCenter, or select **Disable** to disable OneCommand Manager for VMware vCenter.

2.6 Enabling and Disabling OneCommand Manager for VMware vCenter with the Plug-in Manager

If you enable OneCommand Manager for VMware vCenter, the OneCommand Manager for VMware vCenter functionality immediately appears on the vSphere console. If two or more plug-ins are registered, the **Emulex OneCommand** sub-tab (under the **Classic Solutions** tab) immediately appears on the vSphere console. If only one plug-in is registered, the OneCommand Manager for VMware vCenter functionality immediately appears on the **Classic Solutions** tab.

If you disable OneCommand Manager for VMware vCenter, the OneCommand Manager for VMware vCenter functionality immediately disappears. If two or more plug-ins are registered, the **Emulex OneCommand** sub-tab immediately disappears from the vSphere console. If only one plug-in is registered, the OneCommand Manager for VMware vCenter functionality immediately disappears from the **Classic Solutions** tab.

To change the status of OneCommand Manager for VMware vCenter, perform these steps:

1. Log on to the vCenter Server through the VMware vSphere console.
2. Enter the IP address and credentials of the vCenter Server where OneCommand Manager for VMware vCenter is registered.

   The vSphere console appears.
3. In the vSphere console, select **Plug-ins** from the menu bar, and select **Manage Plug-ins**. The **Plug-in Manager** window appears.
4. Click the **Emulex OneCommand** row. A context menu appears (Figure 8).

Figure 8 Plug-in Manager with Selected Row

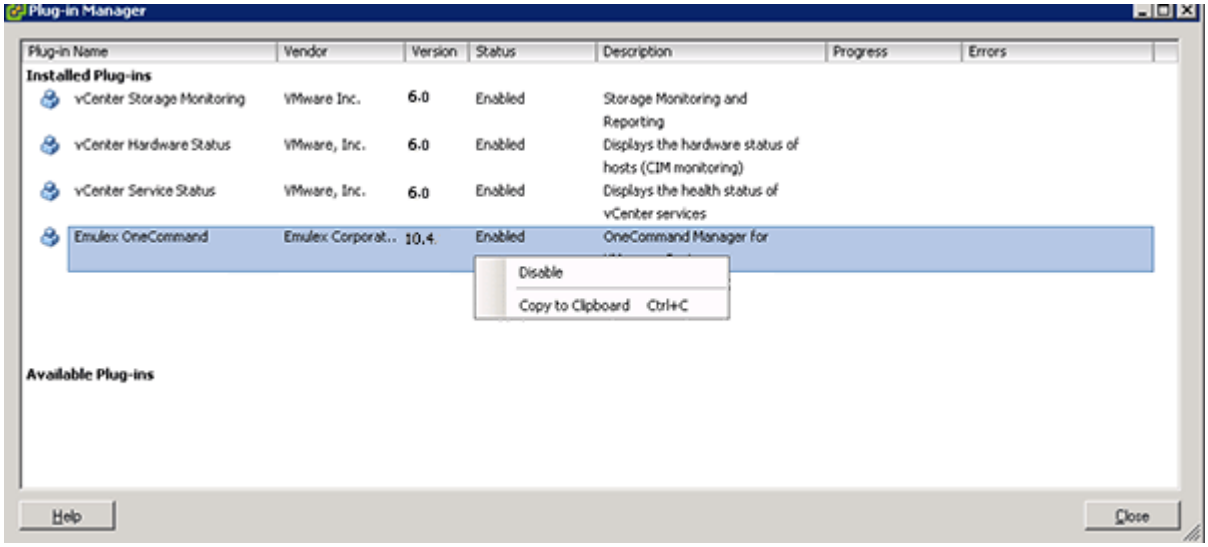

5. Select **Enable** or **Disable**.

2.7 Registering and Unregistering OneCommand Manager for VMware vCenter

OneCommand Manager for VMware vCenter can be registered with more than one vCenter server. If you are using the vSphere Web Client, you must verify that the vSphere Web Client and the vCenter server are registered to the same vCenter Single Sign On server to ensure that the vSphere Web Client can access the vCenter server inventory.

To register or unregister OneCommand Manager for VMware vCenter with a new vCenter server, perform these steps:

1. Double-click the **OCM for VMware vCenter Registration** icon on the desktop. This icon is created when OneCommand Manager for VMware vCenter is successfully installed. The **register/unregister** dialog is displayed (Figure 4).
2. Enter the following details of the vCenter server:
   - vCenter Server Name – The IP address of the vCenter server.
   - vCenter Server HTTPS Port – The HTTPS port number of the vCenter server.
   - The vCenter Server HTTPS port is 443 by default. You can change this value if you have configured a different HTTPS port while installing the vCenter.
   - Username – The user name with required privileges.
   - Password – The user password.
3. Do one of the following:
   - Click **Register** to register OneCommand Manager for VMware vCenter with a new vCenter server.

   or

   - Click **Unregister** to unregister an existing OneCommand Manager for VMware vCenter with a vCenter server.

   NOTES
   - If you change the host name of the machine hosting the vCenter server, you must reinstall the vCenter server and re-register.
   - If the vCenter server is already registered with another instance of OneCommand Manager for VMware server, it is replaced with this server instance.
4. When the operation is successful, a message is displayed. Click **OK**.
5. Close the window.

2.8 Uninstalling OneCommand Manager for VMware vCenter

Before you uninstall OneCommand Manager for VMware vCenter, you must unregister it from the vCenter Server. For more information, see Section 3, Using OneCommand Manager for VMware vCenter.

CAUTION When you uninstall OneCommand Manager for VMware vCenter, ensure that you do not delete the default configuration and log files that are stored in the `%TEMP%\Emulex\OCM for VMware` directory. If these files are deleted, all historical information of active management performed from the host is permanently lost.

To uninstall OneCommand Manager for VMware vCenter in a Windows system, perform these steps:

1. Navigate to the system directory to which you downloaded the `ELXOCM-VMware-vCenter-<version>-Setup.exe` file.
2. Double-click the `ELXOCM-VMware-vCenter-<version>-Setup.exe` file. The OneCommand Manager for VMware vCenter window prompts you to reinstal or uninstall the application. Select **Uninstall the application completely** and click **Next**. A progress window is displayed. The screen indicating the detection of OneCommand Manager for VMware vCenter appears.

3. Click **OK**.
   When uninstallation is complete, the **Uninstallation Completed** window is displayed.
4. Click **Finish**. You do not need to reboot the system.

NOTE You can also uninstall the OneCommand Manager application from the **Programs and Features** window.

2.9 Reinstalling OneCommand Manager for VMware vCenter

To reinstall OneCommand Manager for VMware vCenter in a Windows system, perform these steps:

1. Navigate to the system directory to which you downloaded the `ELXOCM-VMware-vCenter-<version>-Setup.exe` file.
2. Double-click `ELXOCM-VMware-vCenter-<version>-Setup.exe`.
   The **OneCommand Manager for VMware vCenter** window prompts you to reinstall or uninstall the application.
3. Select **Re-install the application** and click **Next**.
   A progress window is displayed. The **OCM for VMware vCenter** window appears.
4. Click **Next**.
   The **License agreement** screen appears.
5. Read the agreement and select the **I agree to these terms and conditions** check box.
6. Click **Next**.
   The **Installation Options** window with the previous installation folder location appears (Figure 1).
7. Ensure that the **OCM for VMware vCenter** check box is selected.
8. To change the installation folder location, click **Browse** and navigate to where you want the program files to reside.
9. Click **Install** on the Installation Options window.
   The **operation in progress** window appears. When the installation process is complete, the **OneCommand Manager for VMware vCenter configuration** window appears (Figure 9).

Figure 9 Reinstallation Configuration Dialog

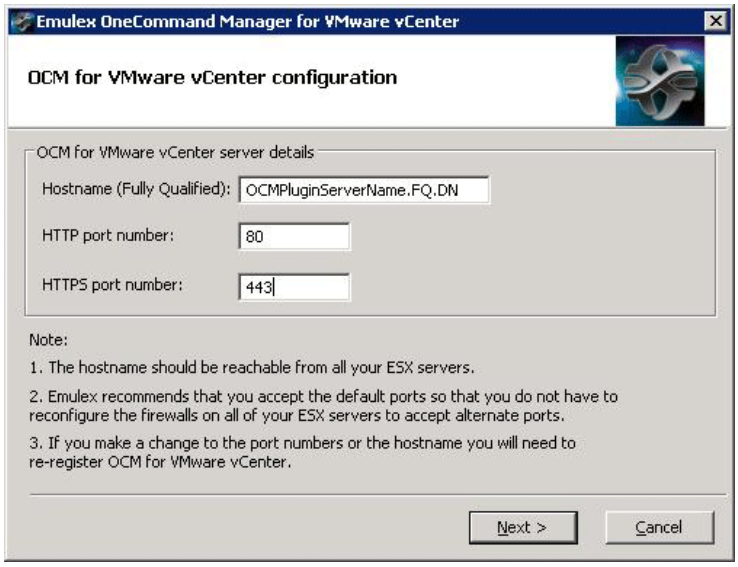

If OneCommand Manager for VMware vCenter was installed earlier with port numbers other than the defaults provided, those configured ports are shown. If the port numbers are already in use, a warning appears next to the port number.

NOTE The Windows firewall setting must allow incoming connections on the HTTP and HTTPS ports that you configure here.

10. Follow the on-screen instructions and complete the installation with steps 10 to 15 of Section 2.4, Installing OneCommand Manager for VMware vCenter.

Chapter 3: Using OneCommand Manager for VMware vCenter

OneCommand Manager for VMware vCenter is available at the host level and the cluster level in the vSphere inventory list.

NOTE To increase the size of the OneCommand Manager for VMware screen, the Recent Tasks panel on the right can be unpinned and collapsed.

3.1 Accessing with VMware Web Client 6.0 and 6.5

After you are logged into the VMware Web Client, the OneCommand Manager for VMware vCenter version 6.5 or version 6.0 is under the **Manage** tab for a particular host or cluster that you select in the vSphere client. When the **Emulex OneCommand** tab is loaded, the functionality and features available are the same as those for the OneCommand Manager for VMware vCenter available for vSphere console.

To launch the OneCommand Manager for VMware vCenter Web Client, perform these steps:

1. Log on to the vSphere Web Client. The vSphere Web Client home page is displayed.
2. Navigate to an ESXi host or cluster in the Navigation pane.
3. Perform one of the following actions:

   From the Host level view, select the host you want to display.

   From the Cluster level view, select the cluster you want to display.
4. Go to the **Manage** tab to access the OneCommand Manager for VMware vCenter. In Web Client 6.5, go to the **Configure** tab and under **more,** see **Emulex OneCommand**.

3.2 Accessing with VMware Web Client 5.5

After you are logged into the vSphere client, the OneCommand Manager for VMware vCenter version 5.5 is under the **Classic Solutions** tab for a particular host or cluster that you select in the vSphere client. When the **Classic Solutions** tab is loaded, the functionality and features available are the same as those for the OneCommand Manager for VMware vCenter available for vSphere console.

To launch the OneCommand Manager for VMware vCenter Web Client, perform these steps:

1. Log on to the vSphere Web Client. The vSphere Web Client home page is displayed.
2. Navigate to an ESXi host or cluster in the Navigation pane.
3. Perform one of the following actions:

   From the Host level view, select the host you want to display.

   From the Cluster level view, select the cluster you want to display.
4. Go to the **Classic Solutions** tab to access the OneCommand Manager for VMware vCenter.

NOTE If OneCommand Manager for VMware vCenter is the only plug-in registered on the vCenter server, the tab title **Emulex OneCommand** does not appear under the **Classic Solutions** tab. OneCommand Manager for VMware vCenter directly loads under the **Classic Solutions** tab with no title. The **Emulex OneCommand** plug-in is shown as a sub-tab only if two or more plug-ins are registered.

3.3 Accessing the vSphere Console

To access the vSphere console, perform these steps:

1. Log on to the VMware vSphere Client:
   a. Double-click the **VMware vSphere Client** icon.
   b. Enter the credentials and click **Login**.
   c. Select a cluster or host in the vSphere console tree-view, and if applicable, click the **Emulex OneCommand** tab to perform active management on a cluster or a host.

   NOTE If OneCommand Manager for VMware vCenter is the only plug-in registered on the vCenter server, the tab title **Emulex OneCommand** does not appear under the **Classic Solutions** tab. OneCommand Manager for VMware vCenter directly loads under the **Classic Solutions** tab with no title. The **Emulex OneCommand** plug-in is shown as a sub-tab only if two or more plug-ins are registered.
2. A Security Alert is displayed. Click **Yes**. The vSphere console is displayed.

3.3.1 vSphere Console Views

The vSphere console supports two views, cluster view and host view.

The difference between the two views is the Emulex Device Management area, which may contain options or a tree-view.

In a cluster view, you can view fabric-centric or host-centric information, about all the hosts in a cluster. Figure 10 shows the Emulex Device Management options available in a host-centric cluster view. You can perform active management simultaneously on a fabric or a group of hosts belonging to a cluster.

In a host view, you can only view information about a single host. Figure 11 shows the Emulex Device Management tree-view elements for a selected host. In the host view, in addition to selecting an element in the tree-view, you must also select a tab in the information pane to view the information for the element selected in the tree-view.

NOTES
- A host can be managed by only one vCenter Server at a time.
- The OneCommand Manager for VMware vCenter user interface is auto-refreshed every five minutes.

To access the cluster view, select a cluster in the vSphere console tree-view, and if applicable, select the **Emulex OneCommand** tab. The cluster view appears (Figure 10).

To access the host view, select a host in the vSphere console tree-view, and if applicable, select the **Emulex OneCommand** tab. The host view appears (Figure 11).

In both cluster and host views, the OneCommand Manager for VMware vCenter window consists of three major parts:

- The Emulex Device Management area
- The Information pane corresponding to the option or element selected in the Emulex Device Management area
- The About area (shows version information about the Emulex OneCommand Manager for VMware vCenter)

3.3.2 OneCommand Manager for VMware vCenter Window Elements

The OneCommand Manager for VMware vCenter window (Figure 10 and Figure 11) contains four basic components:

- The Emulex Device Management area
- The Information pane
- The console buttons
- The Filter options menu (if applicable)

Figure 10 Cluster View with Callouts

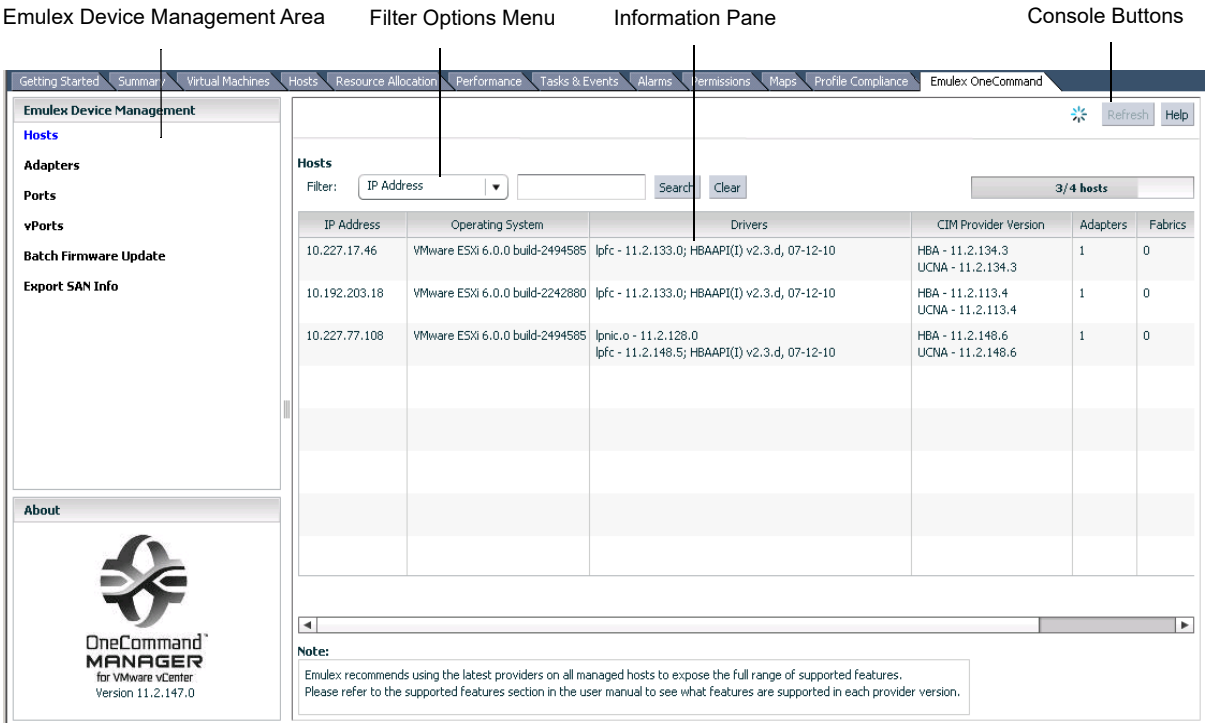

Figure 11 Host View with Callouts

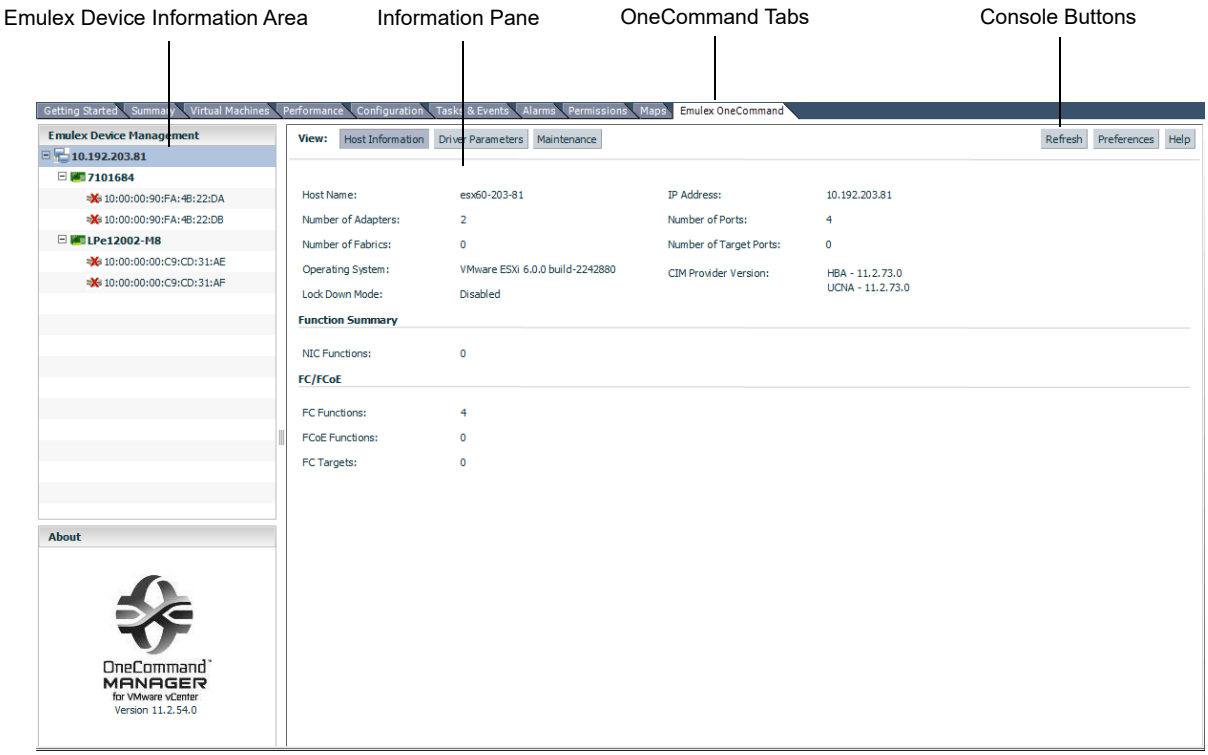

3.3.2.1 Emulex Device Management Area

In a cluster view, the Emulex Device Management area contains links that determine what is displayed in the Information pane.

In a host view, the Emulex Device Management area is a discovery-tree with icons that represent discovered hosts, adapters, ports, virtual ports, fabrics, targets, and LUNs.

3.3.2.2 OneCommand Tabs

In a host view, the **OneCommand** tabs display configuration, statistical, and status information for network elements.

3.3.2.3 Information Pane

In a cluster view, the information pane shows information based upon what is selected in the Emulex Device Management area.

In a host view, the information pane shows information based upon the **OneCommand** tab that is selected.

3.3.2.4 Filter Options Menu

In a cluster view, selecting an item from the Filter Options menu shows information that is sorted and displayed based upon by the item selected.

3.3.2.5 Console Buttons

- **Help** – Click to load the complete indexed online help for OneCommand Manager for VMware vCenter. You can search for information for all OneCommand Manager for VMware vCenter tabs and functions.
- **Preferences** – Click to access the **User Preferences** window. The **User Preferences** window is available in host view only. In the **User Preferences** window, select the **Event Logging** check box to show up or down events for FC and FCoE ports logged into the vSphere console. NIC port events are not logged, as VMware already does this

action. FC and FCoE port events are limited to the active vCenter client. If the same user logon is used from another vCenter client, the **User Preferences** window does not show these events.

- **Refresh** – Click to refresh cluster or host information.

NOTE The **User Preferences** window logs only up and down events for the port. Other events, such as temperature, are not posted.

Chapter 4: Managing Clusters and Hosts

This section pertains to viewing cluster and host information.

4.1 Managing Clusters

From within a cluster, you can view information about:

- Hosts in a cluster
- Adapters belonging to hosts in a cluster
- Physical ports in host-centric mode
- Virtual ports in host-centric mode
- Physical ports in fabric-centric mode

Figure 10 shows the main elements of the cluster view.

4.1.1 Viewing Hosts in a Cluster

NOTE Hosts in a cluster with different provider versions support features as listed in Table 1.

To view information about hosts in a cluster:

1. Select a cluster in the vSphere console tree-view. The **Getting Started** tab is selected by default.
2. If applicable, select the **Emulex OneCommand** tab. The Hosts information pane is displayed (Figure 12).

To filter information for a host within a cluster field:

1. Use the Filter Options list to filter information.
   a. Search by any column title by selecting the title from the Filter Options list.
   b. To further narrow the search, enter a value in the field to the right of the Filter Options list. You can also enter a wildcard for this field.
   c. Click **Search**.
2. Click **Clear** to clear the search criteria.
3. Click **Refresh** to refresh the information. If a host is added to the currently selected cluster, the host information is refreshed.

Figure 12 Hosts within a Cluster

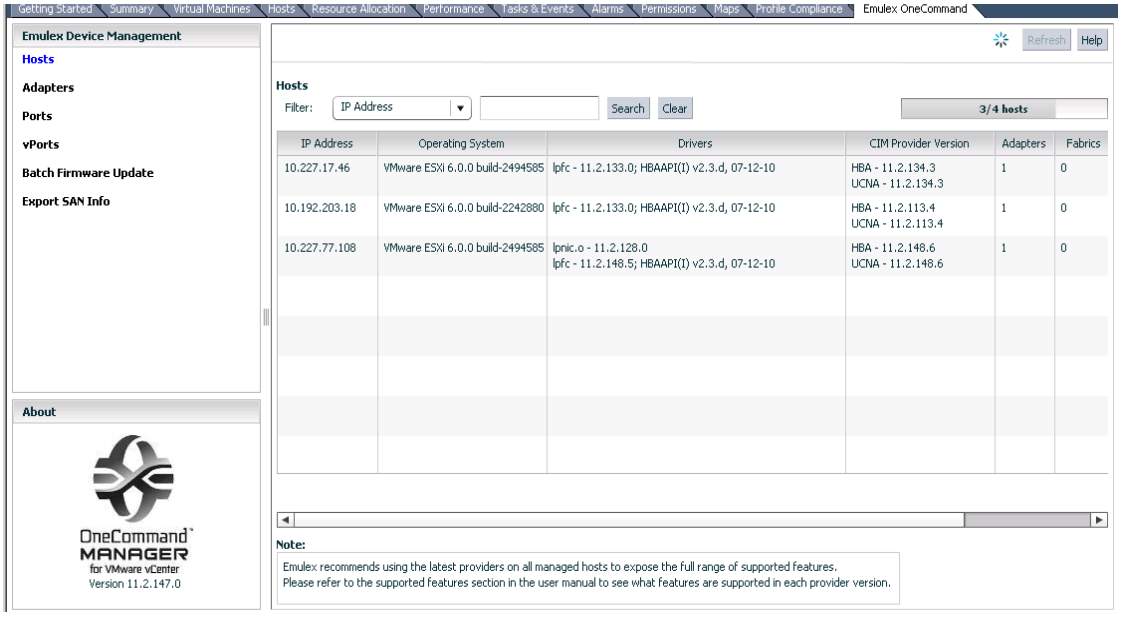

The following Hosts within a Cluster fields are displayed:

- **IP Address** – The IP address of the host in the cluster.
- **Operating System** – The operating system and version installed on the host.
- **Drivers** – The drivers and their versions installed on the host.
- **CIM Provider Version** – The version of the Emulex CIM Provider that is running on the ESXi host.
- **Adapters** – The number of adapters installed in the host.
- **Fabrics** – The number of fabrics to which the host is connected.
- **Ports** – The number of discovered physical ports that can be managed by the host.
- **Lock Down Mode** – Specifies if lockdown mode is enabled or disabled.

4.1.2 Viewing Adapters in a Cluster

To view information about adapters belonging to a host in a cluster, perform these steps:

1. Select a cluster in the vSphere console tree-view, and if applicable, select the **Emulex OneCommand** tab.
2. From the Emulex Device Management options, select **Adapters**. Adapters information is displayed (Figure 13).

To filter adapter information within a cluster, perform these steps:

1. Use the Filter Options list to filter the adapter information. Choose any of the available adapter information fields from the list, enter a wildcard for the field, and click **Search**.
2. Click **Clear** to clear the search criteria and the corresponding adapter information.
3. Click **Refresh** to refresh the information. If an adapter is added to any of the hosts, the adapter information is refreshed.

Figure 13 Adapters within a Cluster

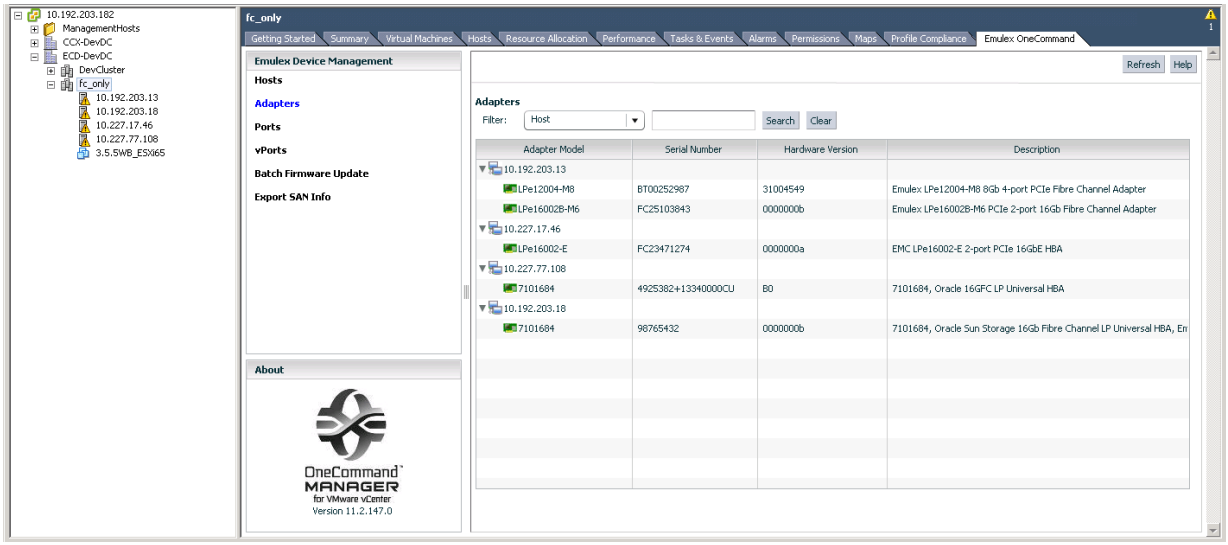

The following Adapters within a Cluster field are displayed:

- **Adapter Model** – The model of the adapter.
- **Serial Number** – The serial number of the adapter.
- **Hardware Version** – For LPe16202/OCe15100 adapters in NIC+FCoE mode, this field shows the board type and code version. For all other adapters, this field shows the JEDEC ID.
- **Description** – The type of adapter.

4.1.3 Viewing Physical Ports in a Cluster (Host-centric Mode)

To view information about a physical port that is in a cluster, in host-centric mode, perform these steps:

1. Select a cluster in the vSphere console tree-view.
2. Select the **Emulex OneCommand** tab. Hosts information is displayed.
3. From the Emulex Device Management options, select **Ports**. The host-centric information pane is displayed (Figure 14 and Figure 15).

NOTE Make sure the **Group by Fabric** check box is not selected.

Figure 14 Physical Ports within a Cluster, Host-centric View

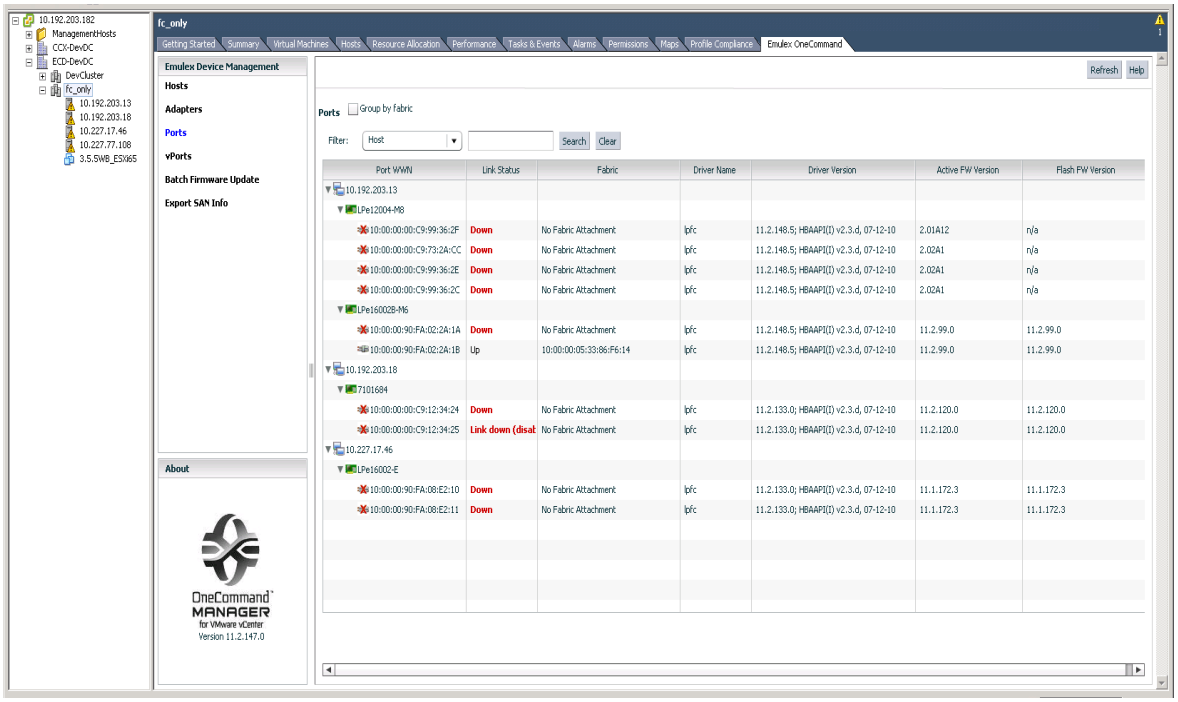

Figure 15 Close-Up of Physical Ports within a Cluster, Host-centric View

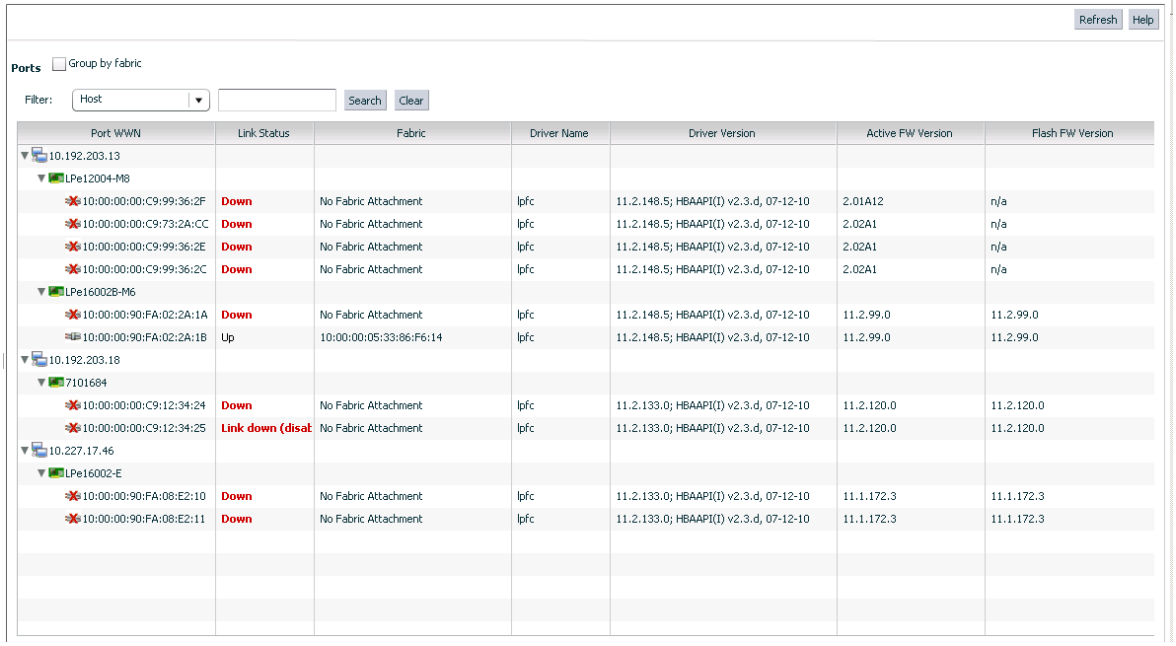

The following Physical Ports within a Cluster, Host-centric fields are displayed:

- **Port WWN** – The port World Wide Name.
- **Link Status** – The status of the link on the selected port.
- **Fabric** – The 64-bit worldwide unique identifier assigned to the fabric.

- **Driver Name** – The executable file image name for the driver as it appears in the Emulex driver download package.
- **Driver Version** – The version of the installed driver.
- **Active FW Version** – The active FW version.
- **Flash FW Version** – The FW version that becomes active after the system is rebooted.

4.1.4 Viewing Virtual Ports in a Cluster (Host-centric Mode)

To view information about a virtual port that is in a cluster, in host-centric mode, perform these steps:

1. Select a cluster in the vSphere console tree-view.
2. If applicable, select the **Emulex OneCommand** tab. Port host information is displayed.
3. From the Emulex Device Management options, select **vPorts**. The virtual port information is displayed (Figure 16).

To filter virtual port information in a cluster, perform these steps:

1. Use the Filter Options list to filter the port information.
2. Choose any of the available host information fields from the list, enter a wildcard for the field, and click **Search**.
3. Click **Clear** to clear the search criteria.

Figure 16 Virtual Ports in a Cluster, Host-centric View

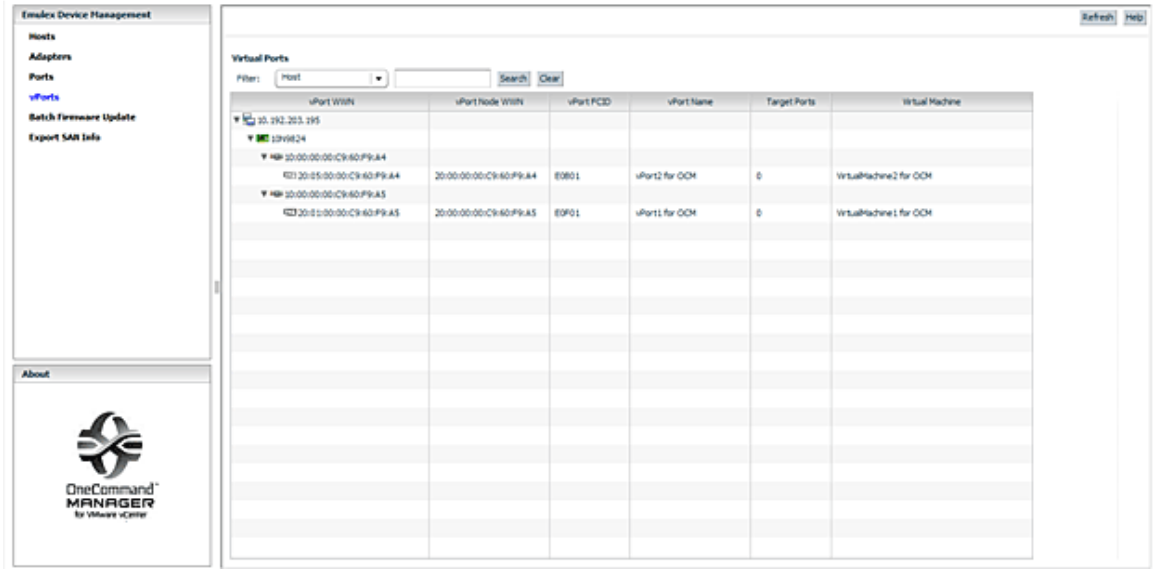

The following Virtual Ports in a Cluster, Host-centric fields are displayed:

- **vPort WWN** – The virtual port World Wide Name.
- **vPort Node WWN** – The virtual port node World Wide Name.
- **vPort FCID** – The virtual port FC ID.
- **vPort Name** – The virtual port name.
- **Target Ports** – The number of target ports.
- **Virtual Machine** – Virtual machine information.

4.1.5 Viewing Physical Port Information in a Cluster (Fabric- centric Mode)

To view physical port information in a cluster, in fabric-centric mode, perform these steps:

1. Select a cluster in the vSphere console tree-view.
2. Select the **Emulex OneCommand** tab, host information is displayed.
3. From the Emulex Device Management options, select **Ports**. The host-centric information pane is displayed.
4. Select the **Group by Fabric** check box. Fabric information is displayed (Figure 17 and Figure 18).

Figure 17 Information for a Physical Port in a Cluster, Fabric-Centric View

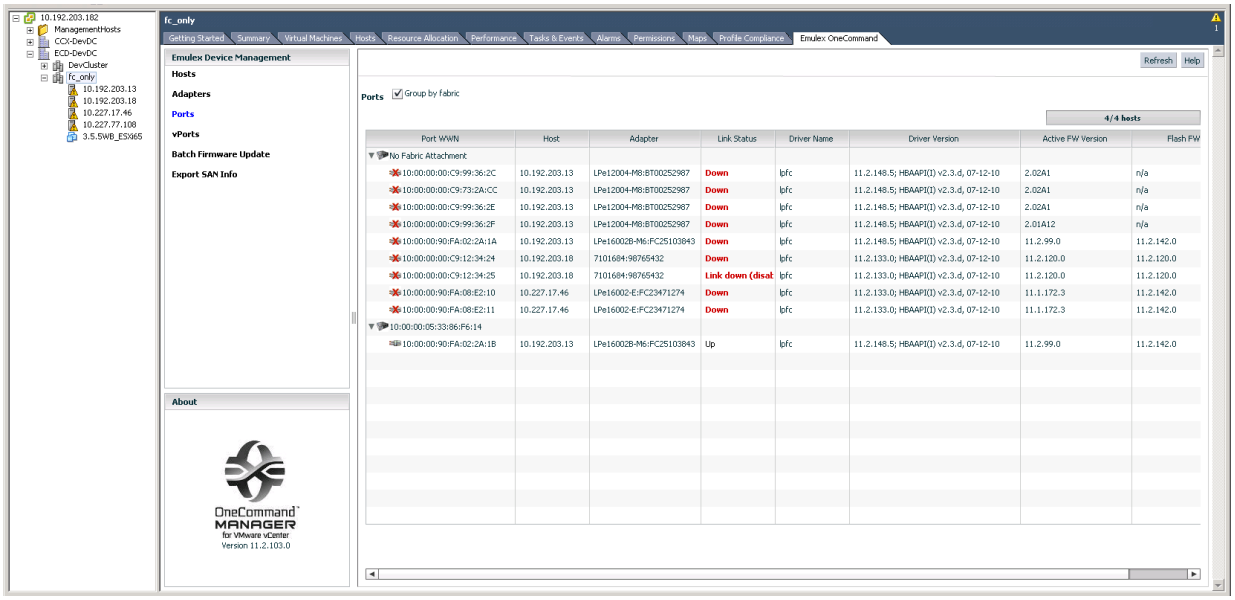

NOTE NIC ports are not attached to any fabric, so they are listed under **No Fabric Attachment**.

Figure 18 Close-up of Information for a Physical Port in a Cluster, Fabric-centric View

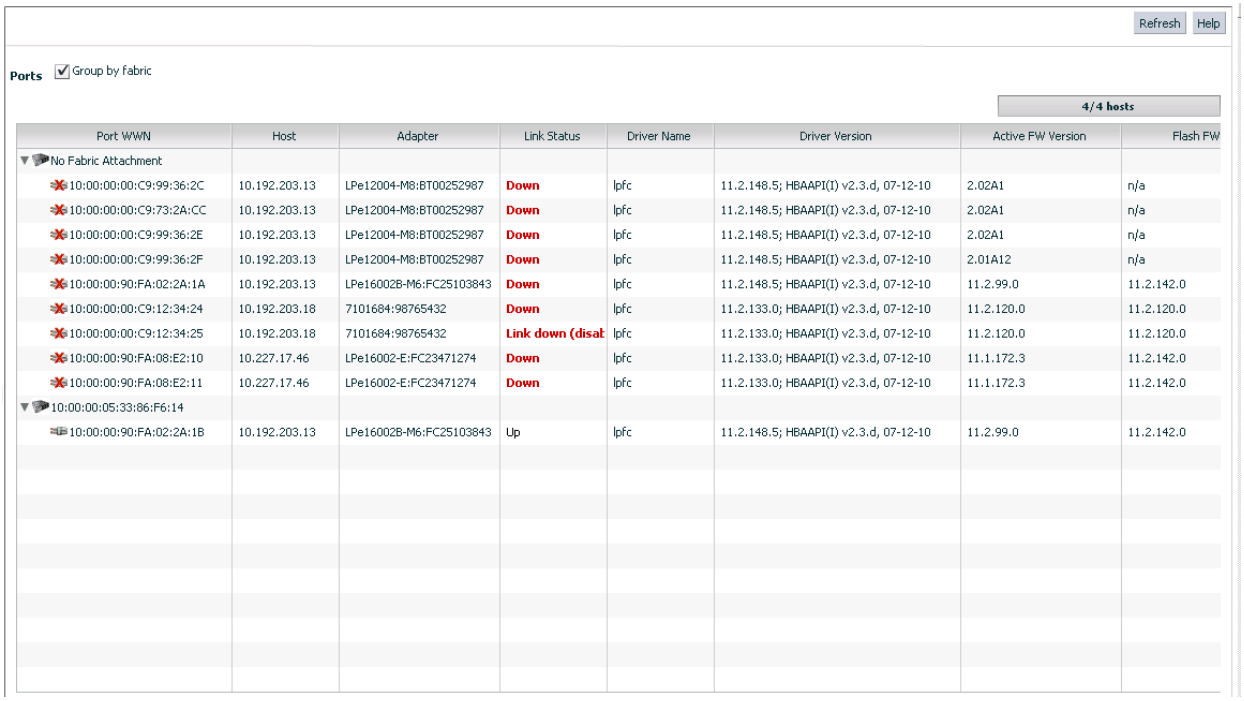

Port WWN	Host	Adapter	Link Status	Driver Name	Driver Version	Active FW Version	Flash FW
No Fabric Attachment							
10:00:00:00:C9:99:36:2C	10.192.203.13	LPe12004-M8:BT00252987	Down	lpfc	11.2.148.5; HBAAPI(I) v2.3.d, 07-12-10	2.02A1	n/a
10:00:00:00:C9:73:2A:CC	10.192.203.13	LPe12004-M8:BT00252987	Down	lpfc	11.2.148.5; HBAAPI(I) v2.3.d, 07-12-10	2.02A1	n/a
10:00:00:00:C9:99:36:2E	10.192.203.13	LPe12004-M8:BT00252987	Down	lpfc	11.2.148.5; HBAAPI(I) v2.3.d, 07-12-10	2.02A1	n/a
10:00:00:00:C9:99:36:2F	10.192.203.13	LPe12004-M8:BT00252987	Down	lpfc	11.2.148.5; HBAAPI(I) v2.3.d, 07-12-10	2.01A12	n/a
10:00:00:90:FA:02:2A:1A	10.192.203.13	LPe16002B-M6:FC25103843	Down	lpfc	11.2.148.5; HBAAPI(I) v2.3.d, 07-12-10	11.2.99.0	11.2.142.0
10:00:00:00:C9:12:34:24	10.192.203.18	7101684:98765432	Down	lpfc	11.2.133.0; HBAAPI(I) v2.3.d, 07-12-10	11.2.120.0	11.2.120.0
10:00:00:00:C9:12:34:25	10.192.203.18	7101684:98765432	Link down (disat	lpfc	11.2.133.0; HBAAPI(I) v2.3.d, 07-12-10	11.2.120.0	11.2.120.0
10:00:00:90:FA:08:E2:10	10.227.17.46	LPe16002-E:FC23471274	Down	lpfc	11.2.133.0; HBAAPI(I) v2.3.d, 07-12-10	11.1.172.3	11.2.142.0
10:00:00:90:FA:08:E2:11	10.227.17.46	LPe16002-E:FC23471274	Down	lpfc	11.2.133.0; HBAAPI(I) v2.3.d, 07-12-10	11.1.172.3	11.2.142.0
10:00:00:05:33:86:F6:14							
10:00:00:90:FA:02:2A:1B	10.192.203.13	LPe16002B-M6:FC25103843	Up	lpfc	11.2.148.5; HBAAPI(I) v2.3.d, 07-12-10	11.2.99.0	11.2.142.0

The following Fabric-centric Information fields are displayed:

- **Port WWN** – The port World Wide Name.
- **Host** – The host IP address.
- **Adapter** – The adapter model.
- **Link Status** – The status of the link on the selected port.
- **Driver Name** – The executable file image name for the driver as it appears in the Emulex driver download package.
- **Driver Version** – The version of the installed driver.
- **Active FW Version** – The active firmware version.
- **Flash FW Version** – The firmware version that becomes active after the system is rebooted.

4.2 Managing Hosts

Host information includes:

- Information for a single host
- Driver parameters for all adapters in a host
- Firmware information for all adapters in a host

Figure 11 shows the main elements of the host view.

4.2.1 Viewing Host Information for a Single Host

To view host information for a single host, select a host in the vSphere console tree-view, and if applicable, select the **Emulex OneCommand** tab. The **Host Information** tab is selected by default and the information of the selected host appears (Figure 19).

To filter the host information, perform these steps:

1. Use the Filter Options list to filter the fabric information.
2. Choose any of the available fabric information fields from the list, enter a wildcard for the field, and click **Search**.
3. Click **Clear** to clear the search criteria.

Figure 19 Information for a Single Host

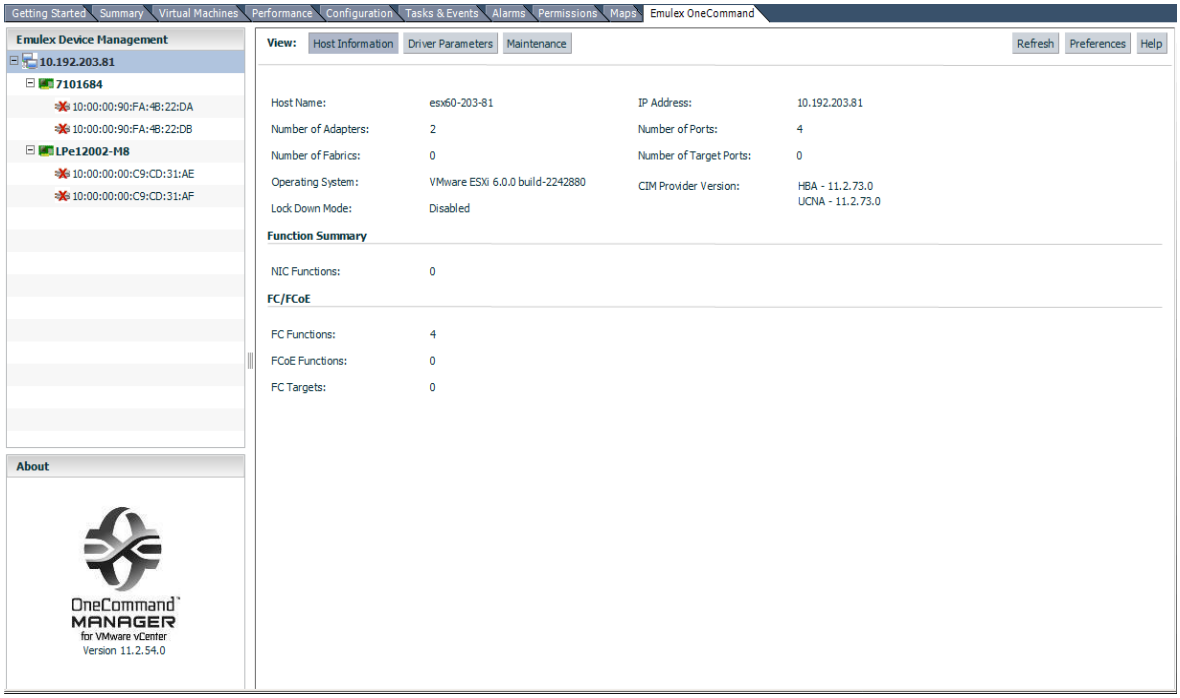

The following Host Information fields are displayed:

- **Host Name** – The host identifier.
- **Number of Adapters** – The number of adapters installed in the host.
- **Number of Fabrics** – The number of fabrics to which the host is connected.
- **Operating System** – The operating system and version installed on the selected host.
- **Lock Down Mode** – Indicates whether lockdown mode is enabled or disabled.
- **IP Address** – The IP address of the host.
- **Number of Ports** – The number of discovered physical ports that can be managed by this host.
- **Number of Target Ports** – The number of targets discovered across the ports.
- **CIM Provider Version** – The versions of the Emulex CIM Providers that are running on the ESXi host.
- Function Summary area (NIC, LPe16202/OCe15100 adapters in NIC+FCoE mode):
  - **NIC Functions** – The number of NIC functions running on the discovered adapters on this host.

- Function Summary area (FC and FCoE):
  - **FC Functions** – The number of FC functions running on the discovered adapters on this host.
  - **FCoE Functions** – The number of FCoE functions running on the discovered adapters on this host.
  - **FC Targets** – The number of FC targets discovered on the FC or FCoE functions on this host.

4.2.2 Viewing Driver Parameters of All Adapters in a Host

The host **Driver Parameters** tab enables you to view and edit the adapter driver parameter settings for a specific host. The host driver parameters are global values and apply to all adapters in that host unless they are overridden by parameters assigned to a specific adapter using the port **Driver Parameters** tab. For each parameter, the information pane shows the current value, the range of acceptable values, the default value, and whether the parameter is dynamic. A dynamic parameter allows the change to take effect without resetting the adapter or rebooting the system. For information on changing parameters for the host, see Section 6.2.5, Configuring FC and FCoE Port Driver Parameters.

NOTE If there are no discovered FC or FCoE adapters, the driver parameters table is empty. This event occurs because there are no adapters to which the host driver parameters apply.

To view driver parameters for all adapters in a host, perform these steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Emulex OneCommand** tab.
2. Select the **Driver Parameters** tab (Figure 20).

Figure 20 Host Driver Parameters

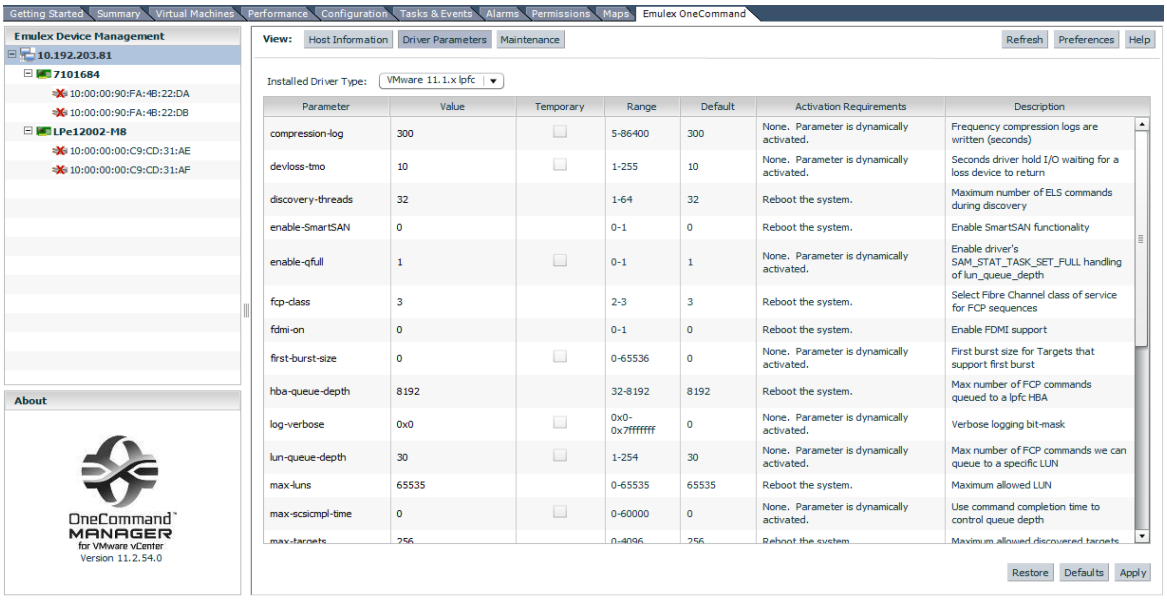

The following host **Driver Parameters** tab fields are displayed:

- **Installed Driver Type** – The current driver installed on this host.
- **Driver Parameter table** – A list of adapter driver parameters and their current values.

  Driver-parameter-specific information is displayed in this area. This information includes value, range, default, activation requirements and description.
  - **Parameter** – The name of the driver parameter.

- **Value** – The value of the driver parameter.
- **Temporary** – Indicates if the value can be set temporarily at port level.
- **Range** – The range of acceptable values for the driver parameter.
- **Default** – The default value of the parameter.
- **Activation Requirements** – The steps required to activate the changed value of the driver parameter.
- **Description** – The description of the driver parameter.

To change the driver parameters for all adapters in a host, perform these steps:

1. From the vSphere console tree-view, select the host whose adapter driver parameters you want to change. If applicable, select the **Emulex OneCommand** tab.
2. Select the host **Driver Parameters** tab (Figure 20). If there are adapters with different driver types installed, the **Installed Driver Types** menu shows a list of all driver types and driver versions that are installed. Select the driver whose parameters you want to change. This menu does not appear if all the adapters are using the same driver.
3. In the driver parameter table, click the **Value** field of a parameter that you want to change. The range for the value is displayed. Enter a value in decimal or hexadecimal format, depending on how the current value is presented. If the value is in hexadecimal format, it is prefixed by 0x (for example, -0x2d).
4. Parameters that can be changed temporarily can only be changed from the corresponding port. Such parameters are represented with a check box next to them.
5. Click **Apply**.

If you changed some parameters and did not click **Apply**, you can restore the parameters back to the value they had before you made the changes. To restore the parameters, click **Restore**.

To reset all parameters back to their default values, click **Defaults**.

4.2.3 Viewing Firmware Information for All Adapters in a Host

The host **Maintenance** tab enables you to view and update firmware on multiple adapters in a specific host. To update firmware on FC adapters in a host, see Section 7.1.1, Updating Firmware on an LPe12000-Series Adapter in a Host.

To view firmware for all adapters in a host, perform these steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Emulex OneCommand** tab.
2. Select the **Maintenance** tab (Figure 21).

Figure 21 Host Maintenance Tab

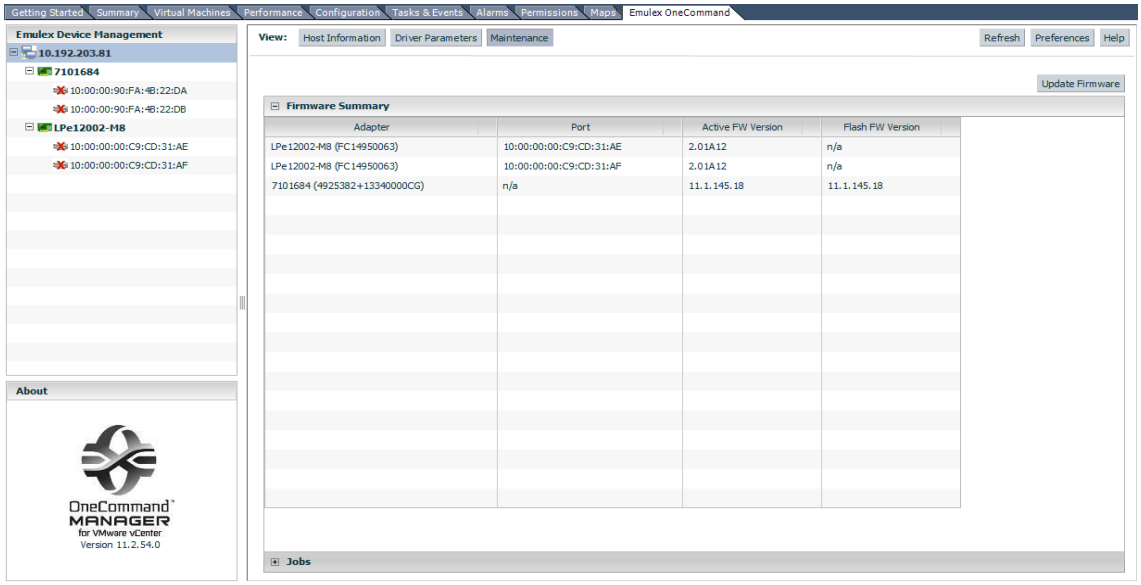

The following host **Maintenance** tab fields are displayed:

- **Adapter** – The model of the adapter.
- **Port** – The port WWN.
- **Active FW Version** – The current firmware on the adapter.
- **Flash FW Version** – The flashed firmware on the adapter. Shows n/a for FC ports that are not available.

Chapter 5: Managing Adapters and Ports

This section describes the various adapter and port management functions you can perform using OneCommand Manager for VMware vCenter.

5.1 FC Adapters

This section pertains to viewing FC adapter and port information.

5.1.1 Viewing FC Adapter Information

When you select an FC adapter from the Emulex Device Management tree-view, the Adapter Information pane shows general attributes associated with the selected adapter.

To view information for an FC adapter, perform these steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Emulex OneCommand** tab.
2. In the Emulex Device Management tree-view, select an FC adapter. The **Adapter Information** tab is displayed (Figure 22).

Figure 22 FC Adapter Information Tab

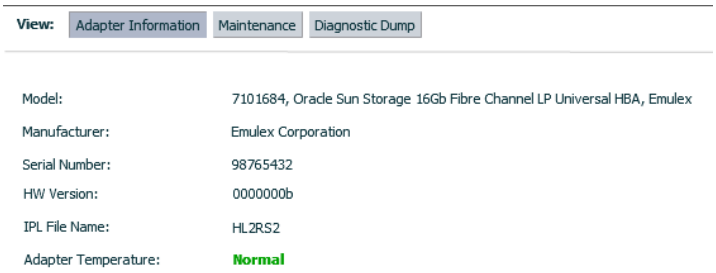

The following FC **Adapter Information** tab fields are displayed:

- **Model** – The complete model name of the adapter.
- **Manufacturer** – The manufacturer of the adapter.
- **Serial Number** – The manufacturer's serial number for the selected adapter.
- **HW Version** – For Emulex LightPulse adapters, this field shows the JEDEC ID.
- **IPL File Name** – For Emulex LightPulse adapters, this field shows the initial program load file name.
- Adapter Temperature:
  - **Normal**: The adapter's temperature is within normal operational range.
  - **Exceeded operational range – Critical:** The adapter's temperature is beyond normal operational range. If the temperature continues to increase, the adapter will shut down. You must determine the cause of the temperature issue and fix it immediately. Check for system cooling issues. Common causes of system cooling issues include clogged air filters, inoperative fans, and air conditioning issues that cause high ambient air temperatures.
  - **Exceeded operational range – Adapter stopped:** The temperature has reached the critical limit, forcing the adapter to shut down. You must determine the cause of the temperature issue and fix it before resuming

operation. Check for system cooling issues. Common causes of system cooling issues include clogged air filters, inoperative fans, and air conditioning issues that cause high ambient air temperatures.

— **Not Supported** - The adapter temperature is not available.

After the system overheating issue is resolved and the adapter has cooled down, reboot the system or, if the system supports hot swapping, cycle the power of the adapter slot.

5.1.2 Viewing FC Port Details

When you select an FC port from the Emulex Device Management tree-view, the **Port Details** tab contains general attributes associated with the selected FC port.

You can also configure virtual machine ID (VMID) when it is supported by the switch.

To view details for an FC port, perform these steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Emulex OneCommand** tab.
2. In the Emulex Device Management tree-view, select the FC port whose information you want to view.

   The **Port Details** tab is displayed (Figure 23).

Figure 23 FC Port Details Tab (Priority Tagging supported)

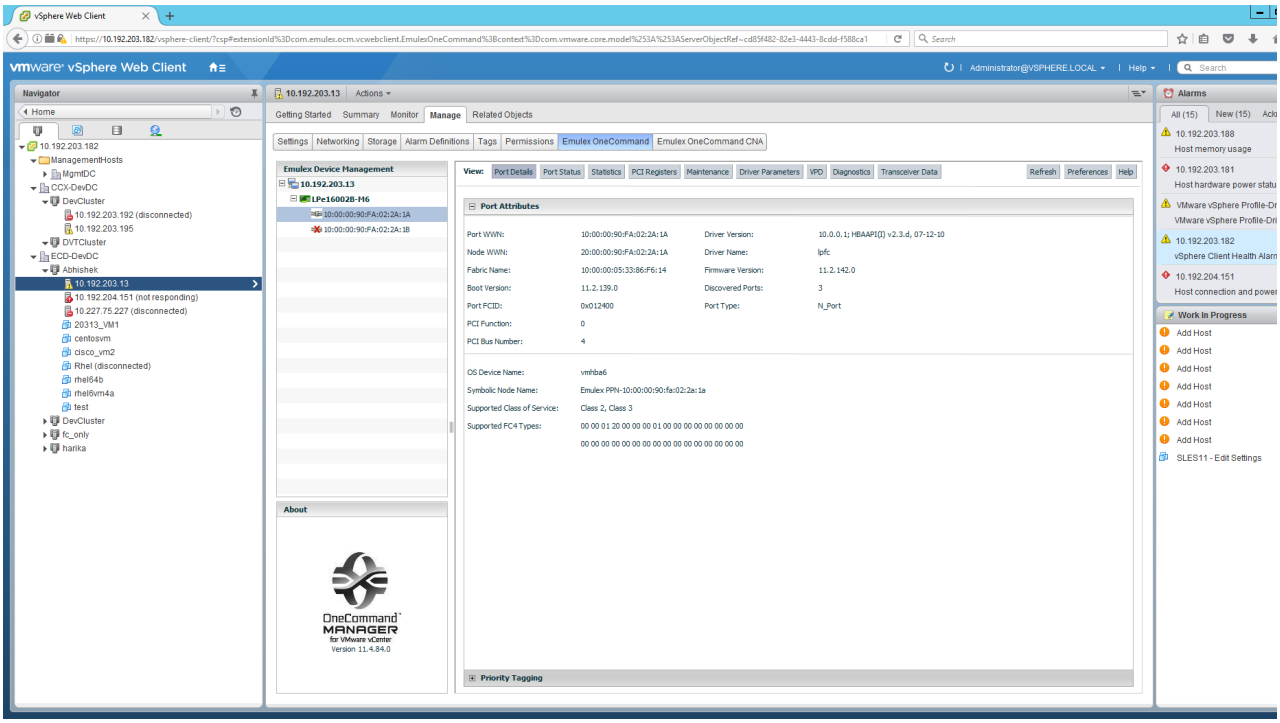

The following FC **Port Details** tab fields are displayed:

- Port Attributes area:
  - **Port WWN** – The port World Wide Name of the selected adapter.
  - **Node WWN** – The node World Wide Name of the selected adapter.
  - **Fabric Name** – The 64-bit worldwide unique identifier assigned to the fabric.
  - **Boot Version** – The version of boot code installed on the selected adapter port. If the boot code is disabled, the field shows **Disabled**.

- **Port FCID** – The FC ID of the selected adapter port.
- **PCI Function** – The PCI function number of the selected port.
- **PCI Bus Number** – The PCI bus number.
- **Driver Version** – The version of the driver installed for the adapter.
- **Driver Name** – The executable file image name for the driver as it appears in the Emulex driver download package.
- **Firmware Version** – The version of Emulex firmware currently active on the adapter port.
- **OS Device Name** – The platform-specific name by which the selected adapter is known to the operating system.
- **Symbolic Node Name** – The FC name used to register the driver with the name server.
- **Supported Class of Service** – A frame delivery scheme exhibiting a set of delivery characteristics and attributes. There are three classes of service.
  - **Class 1**– Provides a dedicated connection between a pair of ports with confirmed delivery or notification of non-delivery.
  - **Class 2** – Provides a frame switched service with confirmed delivery or notification of non-delivery.
  - **Class 3** – Provides a frame switched service similar to Class 2, but without notification of frame delivery or non-delivery.
- **Supported FC4 Types** – A 256-bit (8-word) map of the FC-4 protocol types supported by the port containing the selected adapter.
- **Priority Tagging** – Click this button to configure priority tagging. (This feature is not available if it is not supported by the switch and the adapter.)

5.1.2.1 Configuring Priority Tagging

The **Priority Tagging** page displays all the VMs on an ESXi host and the corresponding priority mapping for each VM. Each entry contains the VM name, the VM universal unique identifier (VM UUID), the assigned priority group, the time the VM was last accessed, and 'Delete If Inactive' indicating whether the mapping will be deleted when there is no I/O from a VM.

Using the **Priority Tagging** page, you can view the priority group assigned to a VM.

NOTE Priority tagging configuration is supported only when the `vmid-priority-tagging` driver parameter is enabled. See Section 6.2.5, Configuring FC and FCoE Port Driver Parameters, for additional information.

Figure 24 Priority Tagging Page

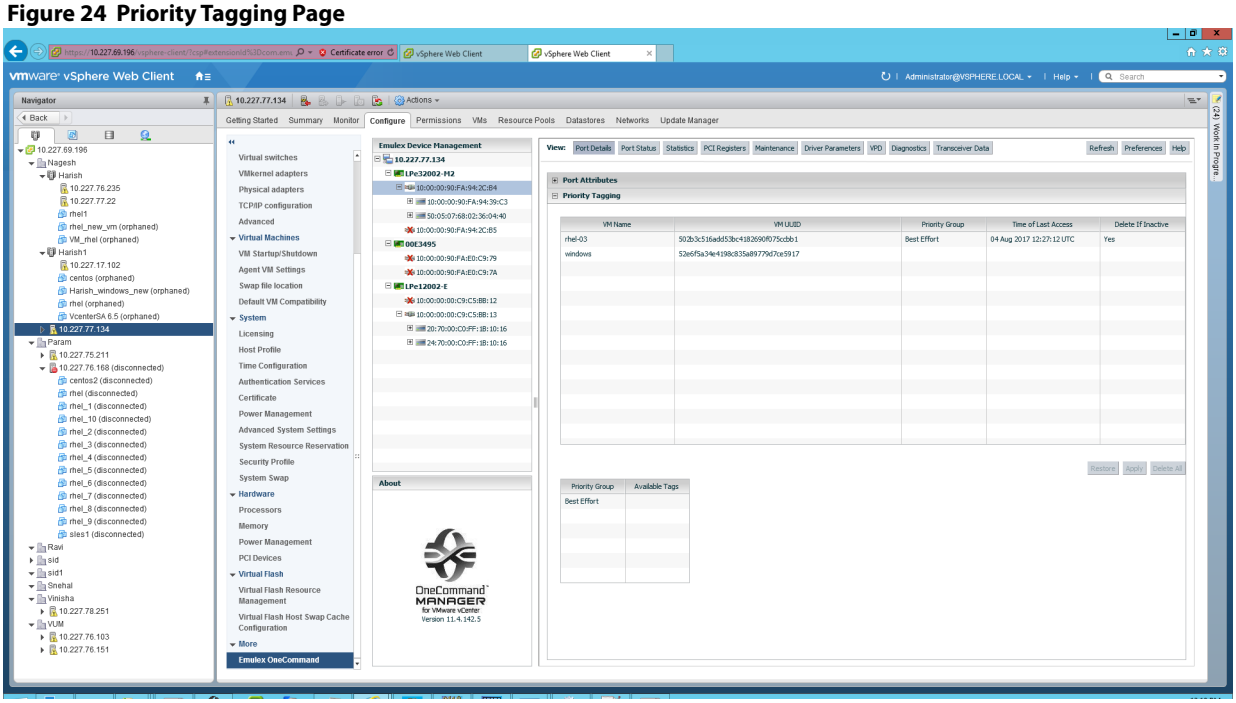

5.1.3 Viewing FC Port Status Information

The FC **Port Status** tab displays the link status and port speed for the selected port. You can also enable or disable the port and configure its link speed.

Figure 25 FC Port Status Tab

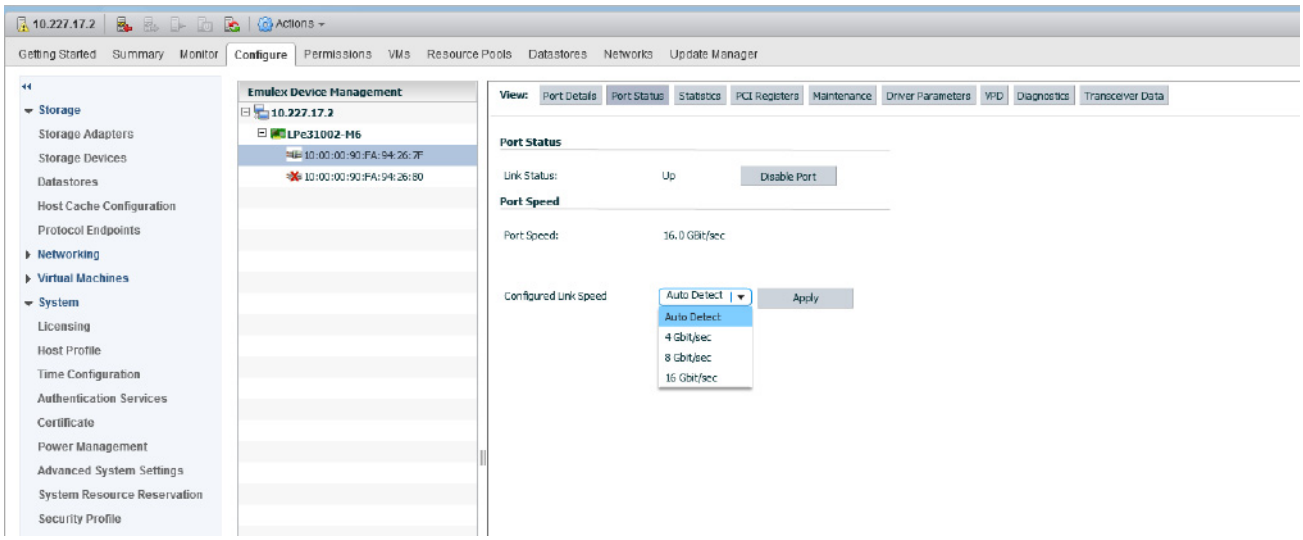

The following FC **Port Status** tab fields are displayed:

- Port Status area:
  - **Link Status** – This field indicates the status of the link on the selected adapter port.
- Port Speed area:
  - **Port Speed** – This field indicates the current port speed of the selected adapter port.

— **Configured Link Speed** – Link speeds that are supported on the port. The list varies depending on the adapter type. The list also includes an Auto Detect option, which indicates that the link speed should be auto-negotiated.

NOTE If an installed adapter does not support forced link speeds, the Configured Link Speed settings and the **Apply** button are not shown.

5.1.3.1 Enabling and Disabling an FC Port

When you disable an FC port, you disable all functions for the port. Disabled ports appear in the Emulex Device Management tree-view with an x over the port icon.

CAUTION Do not disable a boot port; doing so could result in data loss or corruption.

NOTE Ensure that there is no I/O traffic on the port before disabling it.

To enable or disable an FC port, perform these steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Emulex OneCommand** tab.
2. In the Emulex Device Management tree-view, select the FC port that you want to enable or disable.
3. Select the **Port Status** tab (Figure 25).
4. Click either **Enable Port** or **Disable Port**.

5.1.3.2 Configuring the Link Speed on an FC Port

To configure link speed on an FC port, perform these steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Emulex OneCommand** tab.
2. In the Emulex Device Management tree-view, select the FC port that you want to enable or disable.
3. Select the **Port Status** tab (Figure 25).
4. Select a link speed from the **Configured Link Speed** pull-down menu.
5. Click **Apply** to set the new link speed. The **Apply** button is only enabled if the currently selected link speed does not match the currently configured speed.

   If the speed has been set successfully, the following message is displayed stating that the port must be reset to activated the new speed setting.

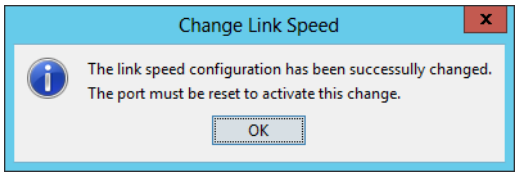

6. Click **OK** to reset the port.

In some situations, the currently configured link speed is not in the supported speed list for the port. This situation can occur if a new SFP is installed that supports a different set of link speeds than the previously installed SFP. If the currently configured link speed is not in the supported speed list, the following message is displayed:

```
Warning: The currently configured port speed is not a valid supported speed.
Please select a link speed and click Apply.
```

The **Apply** button remains enabled until you select a valid port speed.

If the installed SFP is not supported by the adapter, you cannot configure a link speed. If this is attempted, the following message is displayed:

```
Unsupported optics installed..
```

Chapter 6: Managing Protocols

This section pertains to managing NIC, FC, and FCoE protocols.

6.1 NIC

This section pertains to viewing NIC adapter and port information. This section applies only to LPe16202/OCe15100 adapters in NIC+FCoE mode.

6.1.1 Viewing NIC Port Information

When you select a NIC port from the Emulex Device Management tree-view, the NIC **Port Information** tab shows general attributes associated with the selected NIC port.

If SR-IOV is enabled, multiple VFs can be created on a NIC PF for an adapter port. These VFs are assigned by the ESXi host to the VMs. Each VM can be assigned one or more VFs by the guest operating system running on the VM. To the guest operating system, the VF is an independent NIC function with its own MAC address, and it is available for network I/O.

The discovery-tree displays the VFs running on the discovered PFs. It also shows the selected VFs, MAC Address, VLAN ID, Link Status, and Transmit Rate.

NOTE To use SR-IOV, you must enable SR-IOV in the OCM for VMware vCenter application and the system BIOS. Refer to the *Emulex Drivers for VMware ESXi User Guide* for information on enabling SR-IOV in the system BIOS.

To view general NIC port information, perform these steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Emulex OneCommand** tab.
2. In the Emulex Device Management tree-view, select the NIC port for which you want to view information.
3. Select the **Port Information** tab.

Figure 26 NIC Port Information Tab

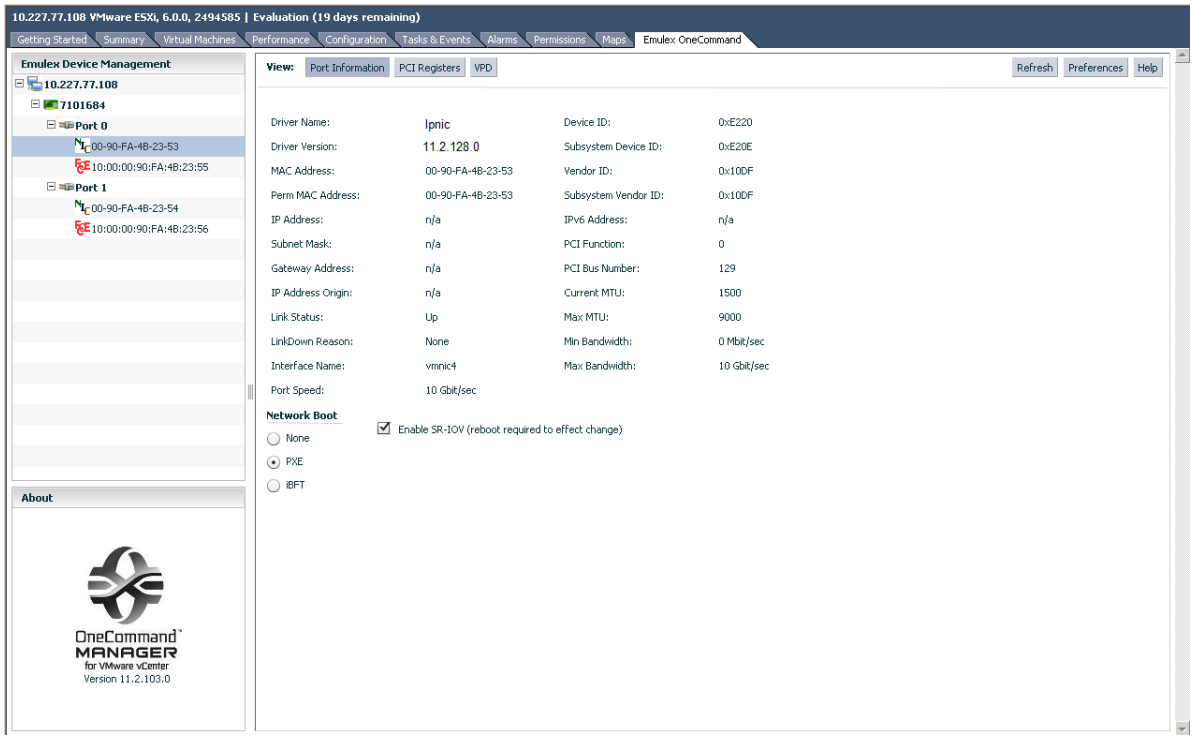

The following NIC **Port Information** tab fields are displayed:

- **Driver Name** – The NIC driver file name.
- **Driver Version** – The NIC driver version.
- **MAC Address** – The NIC MAC address currently assigned to the port.
- **Perm MAC Address** – The original factory-assigned NIC MAC address.
- **IP Address** – The IP address for the NIC port.
- **Subnet Mask** – The subnet mask for the NIC port.
- **Gateway Address** – The gateway address for the NIC port.
- **IP Address Origin** – The origin of the IP address (DHCP or static).
- **Link Status** – The link status is always displayed as n/a (not available).
- **LinkDown Reason** – The reason the link is down on the selected port.
- **Interface Name** – The interface assigned to this port by the host operating system.
- **Port Speed** – The current port speed of the NIC port.
- **Device ID** – The PCI device ID assigned to the NIC function.
- **Subsys Device ID** – The PCI subsystem ID assigned to the NIC function.
- **Vendor ID** – The PCI vendor ID assigned to the NIC function.
- **Subsys Vendor ID** – The PCI subsystem vendor ID assigned to the NIC function.
- **IPv6 Address** – The IPv6 address, if applicable.
- **PCI Function** – The PCI function number assigned to the NIC function.
- **PCI Bus Number** – The PCI bus number to which the port is mapped.
- **Current MTU** – The current transmission unit for Ethernet traffic.
- **Max MTU** – The maximum transmission unit for Ethernet traffic.
- **Min Bandwidth** – The minimum SoftCap bandwidth assigned.
- **Max Bandwidth** – The maximum SoftCap bandwidth assigned.

- **Network Boot** radio buttons – For additional information on network boot, see Section 6.1.1.1, Network Boot. The **Network Boot** radio buttons do not appear if SR-IOV is enabled on the adapter.
  - **None** – Select to disable the network boot functions.
  - **PXE** – Select to boot from a network interface using a preboot execution environment.

6.1.1.1 Network Boot

NIC functions include PXE boot configuration, which is displayed under the Network Boot section of the NIC **Port Information** tab. If the currently selected NIC port does not support network boot, the Network Boot section is not displayed.

NOTE Network boot is supported by CIM Provider 3.6.12 and later.

Network boot options are only available on ports with a PCI function of 0 or 1.

To select a network boot configuration on the selected port, click the desired option from the Network Boot section:

- **PXE** – To boot from a network interface using a preboot execution environment.
- **None** – To disable the network boot functions.

6.1.2 Viewing PCI Registers for a NIC Port

The **PCI Registers** tab contains PCI information about the selected NIC function.

To view a NIC port's PCI registers, perform these steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Emulex OneCommand** tab.
2. In the Emulex Device Management tree-view, select the NIC port for which you want to view PCI information.
3. Select the **PCI Registers** tab (Figure 27).

Figure 27 NIC Port PCI Registers Tab

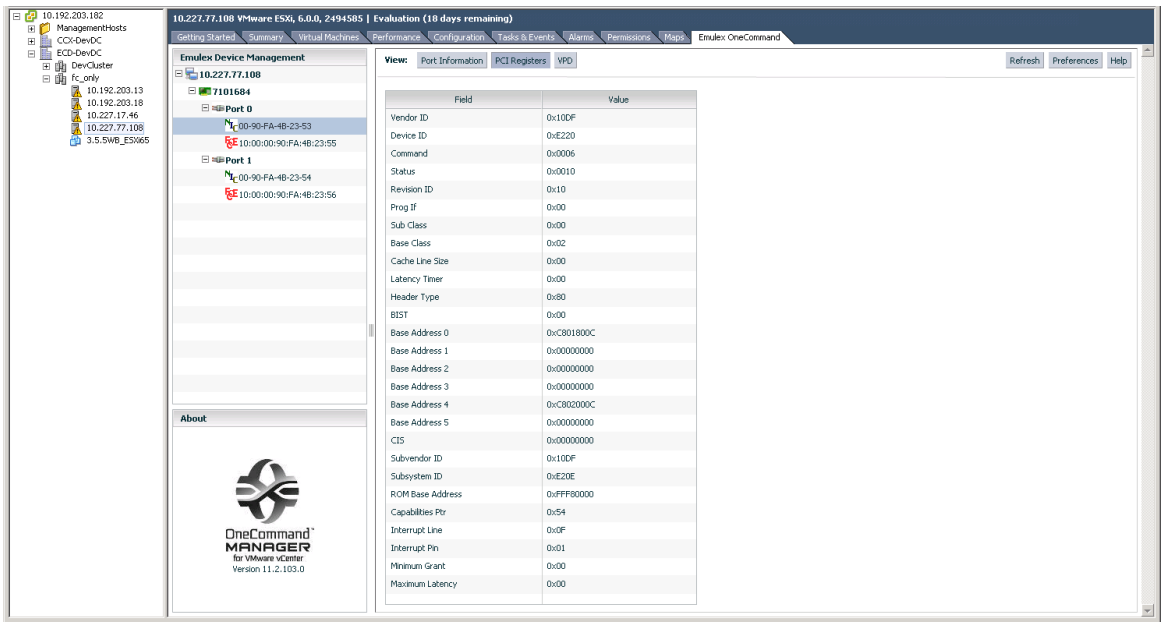

6.1.3 Viewing NIC Port Vital Product Data (VPD)

The **VPD** tab shows vital product data (if available) for the selected NIC port such as the product name, part number, and serial number for example.

To view VPD information of an NIC port, perform these steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Emulex OneCommand** tab.
2. In the Emulex Device Management tree-view, select the NIC port for which you want to view VPD information.
3. Select the **VPD** tab (Figure 28).

Figure 28 NIC Port VPD Tab

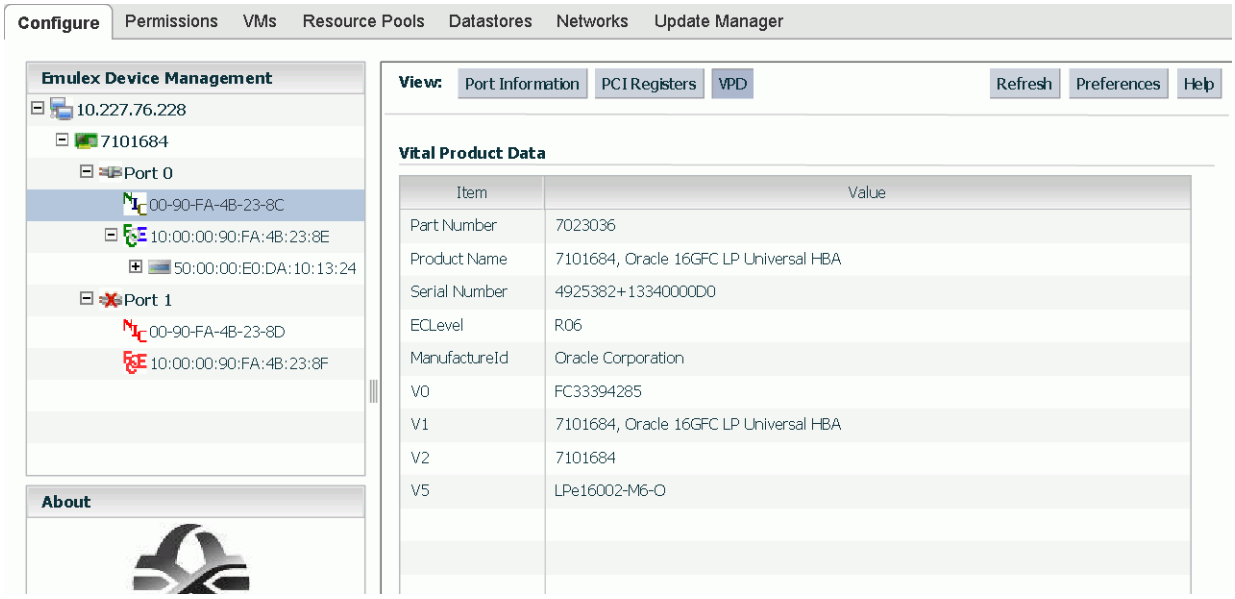

The following **VPD** tab fields are displayed:

- **Part Number** – The adapter's part number.
- **Product Name** – The product information about a selected adapter port.
- **Serial Number** – The adapter's serial number.
- **ECLevel Manufacturer ID** – The ECLevel manufacturer ID.
- **VO** – Vendor-unique data. V indicates a vendor-specific field. An adapter can have none, one, or more than one of these fields defined. Valid values for this field are VO (the letter O, not the number zero) and V*x* (where *x* is a number).

NOTE Some adapters may show additional VPD information, such as EC and MN.

6.2 FC and FCoE

This section pertains to managing FC and FCoE adapter and port information.

6.2.1 Viewing FCoE Port Details

Use the **Port Details** tab to view general attributes associated with the selected FCoE port.

NOTE FCoE-specific information applies only to LPe16202/OCe15100 adapters in NIC+FCoE mode.

To view port details for an FCoE port, perform these steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Emulex OneCommand** tab.
2. In the Emulex Device Management tree-view, select the FCoE port whose details you want to view.
3. Select the **Port Details** tab (Figure 29).

Figure 29 FCoE Port Details Tab

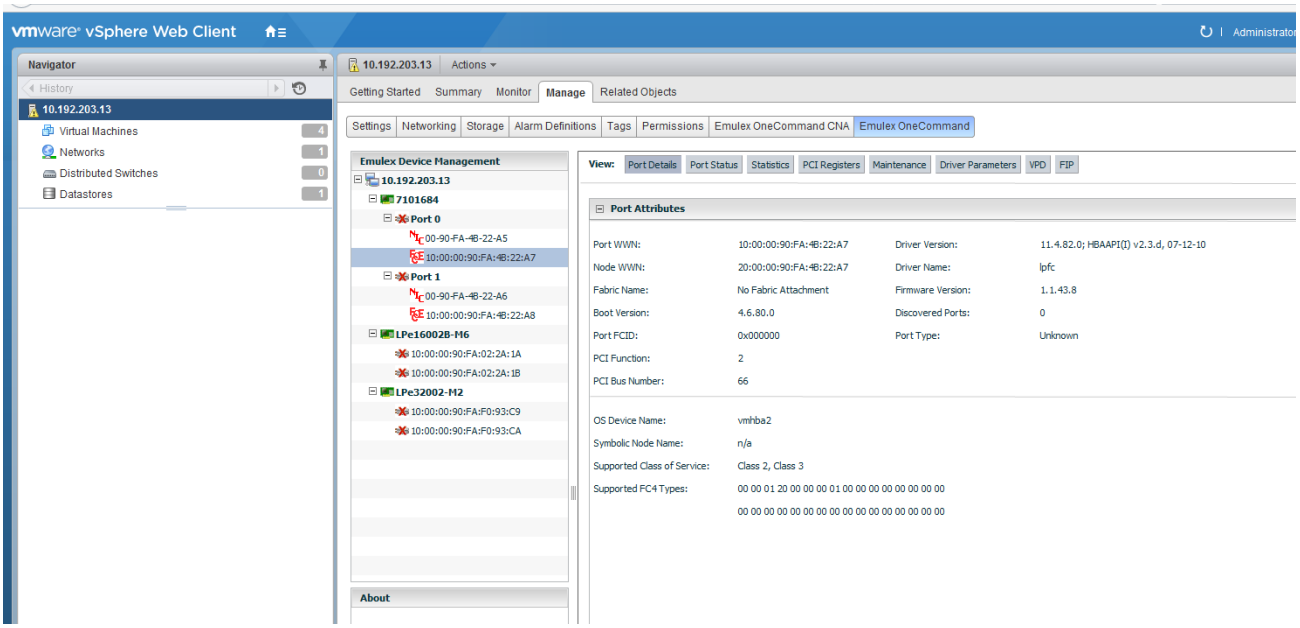

The following **FCoE Port Details** tab fields are displayed:

- Port Attributes Area
  - **Port WWN** – The port World Wide Name of the adapter.
  - **Node WWN** – The node World Wide Name of the selected adapter.
  - **Fabric Name** – The Fabric Name field is displayed in Host view. This is a 64-bit worldwide unique identifier is assigned to the fabric. The host name is the name of the host containing the adapter.
  - **Boot Version** – The version of boot code installed on the selected adapter port. If the boot code is disabled, the field shows **Disabled**.
  - **Port FCID** – The FCoE ID for the selected adapter port.
  - **PCI Function** – The PCI function number assigned by the system.
  - **PCI Bus Number** – The PCI bus number.
  - **Driver Version** – The version of the driver installed for the adapter.
  - **Driver Name** – The executable file image name for the driver as it appears in the Emulex driver download package.
  - **Firmware Version** – The version of Emulex firmware currently active on the adapter port.

- **Discovered Ports** – The number of mapped and unmapped ports found during discovery by the Emulex adapter driver. The mapped ports are targets and the unmapped ports are non-targets, such as switches or adapters.
- **Port Type** – The current operational mode of the selected adapter's port.
- **OS Device Name** – The platform-specific name by which the selected adapter is known to the operating system.
- **Symbolic Node Name** – The FC name used to register the driver with the name server.
- **Supported Class of Service** – A frame delivery scheme exhibiting a set of delivery characteristics and attributes. There are three classes of service:
  - **Class 1** – Provides a dedicated connection between a pair of ports with confirmed delivery or notification of non-delivery.
  - **Class 2** – Provides a frame switched service with confirmed delivery or notification of non-delivery.
  - **Class 3** – Provides a frame switched service similar to Class 2 but without notification of frame delivery or non-delivery.
- **Supported FC4 Types** – A 256-bit (8-word) map of the FC-4 protocol types supported by the port containing the selected adapter.

6.2.2 Viewing FC and FCoE Port Statistics

When you select an FC or FCoE port from the discovery-tree, the **Port Statistics** tab shows cumulative totals for error events and statistics on the port. Some statistics are cleared when the adapter is reset.

To view statistics for a port, perform these steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Emulex OneCommand** tab.
2. In the Emulex Device Management tree-view, select the port whose statistics you want to view.
3. Select the **Statistics** tab (Figure 30).

Figure 30 Port Statistics Tab

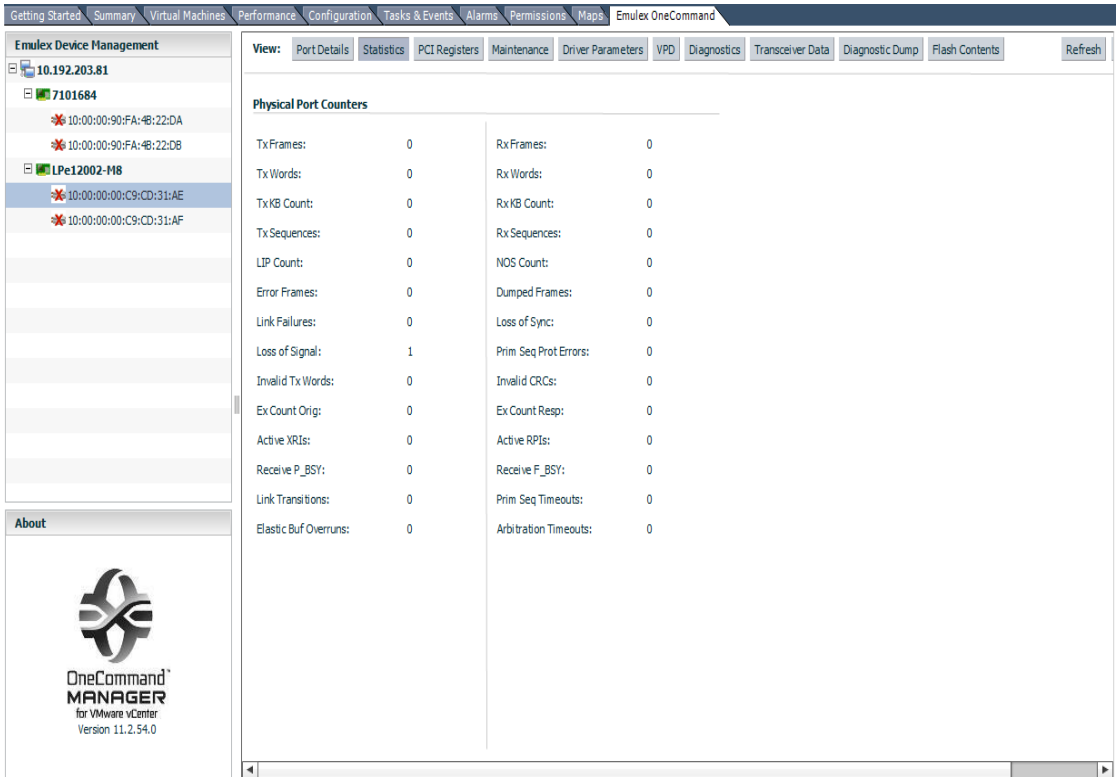

The following **Port Statistics** tab fields are displayed:

- **Tx Frames** – The FC frames transmitted by this adapter port.
- **Tx Words** – The FC words transmitted by this adapter port.
- **Tx KB Count** – The FC kilobytes transmitted by this adapter port.
- **Tx Sequences** – The FC sequences transmitted by this adapter port.
- **LIP count** – The number of LIP events that have occurred for the port. This field is supported only if the topology is arbitrated loop.

  Loop initialization consists of the following:

  — Temporarily suspending loop operations.
  — Determining whether loop capable ports are connected to the loop.
  — Assigning AL_PA IDs.
  — Providing notification of configuration changes and loop failures.
  — Placing loop ports in the monitoring state.
- **Error Frames** – The number of frames received with CRC errors.
- **Link Failures** – The number of times the link has failed. A link failure can cause a timeout.
- **Loss of Signal** – The number of times the signal was lost.
- **Invalid Tx Words** – The total number of invalid words transmitted by this adapter port.
- **Ex Count Orig** – The number of FC exchanges originating on this port.
- **Active XRIs** – The number of active exchange resource indicators.
- **Received P_BSY** – The number of FC port-busy link response frames received.
- **Link Transitions** – The number of times the SLI port sent a link attention condition.
- **Elastic Buf Overruns** – The number of times the link interface has had its elastic buffer overrun.
- **Rx Frames** – The number of FC frames received by this adapter port.

- **Rx Words** – The number of FC words received by this adapter port.
- **Rx KB Count** – The received kilobyte count by this adapter port.
- **Rx Sequences** – The number of FC sequences received by this adapter port.
- **NOS Count** – The number of NOS events that have occurred on the switched fabric (not supported for an arbitrated loop).
- **Dumped Frames** – The number of frames that were lost due to a lack of host buffers available.
- **Loss of Sync** – The number of times loss of synchronization has occurred.
- **Prim Seq Prot Errs** – The primitive sequence protocol error count. This counter is incremented whenever there is any type of protocol error.
- **Invalid CRCs** – The number of frames received that contain CRC failures.
- **Ex Count Resp** – The number of FC exchange responses made by this port.
- **Active RPIs** – The number of remote port indicators.
- **Receive F_BSY** – The number of FC port-busy link response frames received.
- **Primitive Seq Timeouts** – The number of times a primitive sequence event timed out.
- **Arbitration Timeouts** – The number of times the arbitration loop has timed out. Large counts could indicate a malfunction somewhere in the loop or heavy usage of the loop.

6.2.3 Viewing PCI Registers

The **PCI Registers** tab shows PCI information, including PCIe details, about the selected port.

NOTE The PCI fields may vary with the type of adapter installed.

To view PCI registers for an FC or FCoE port, perform these steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Emulex OneCommand** tab.
2. In the Emulex Device Management tree-view, select the port whose PCI information you want to view.
3. Select the **PCI Registers** tab (Figure 31).

Figure 31 Port PCI Registers Tab

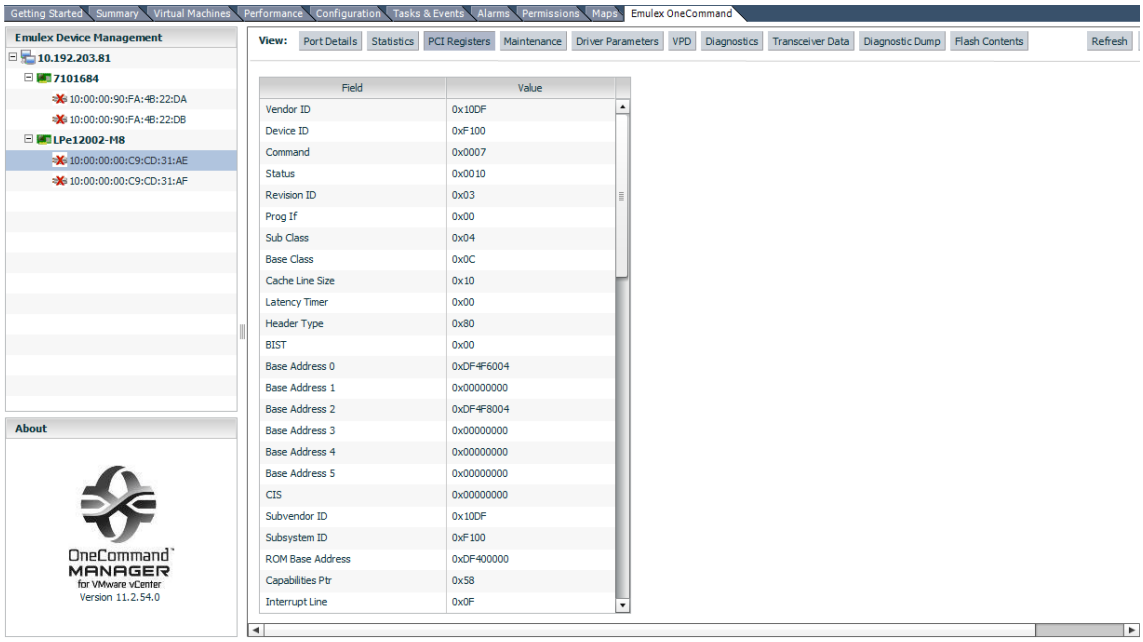

6.2.4 Viewing FC and FCoE Port Maintenance and Firmware Information

The **Maintenance** tab shows firmware information for an FC or FCoE port.

To view firmware information for a port, perform these steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Emulex OneCommand** tab.
2. In the Emulex Device Management tree-view, select the port whose firmware information you want to view.
3. Select the **Maintenance** tab for a port on an LPe12000-series adapter (Figure 32) or a port on any other LightPulse adapter port (Figure 33).

Figure 32 Maintenance Tab for a Port on an LPe12000-Series Adapter

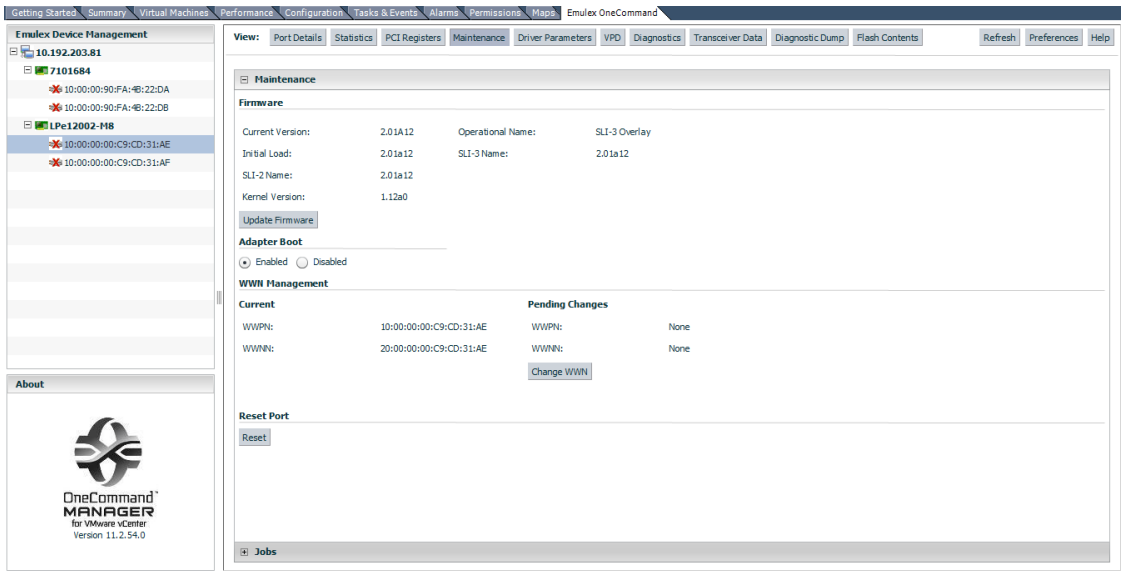

Figure 33 Port Maintenance Tab for a Port on Other LightPulse Adapters

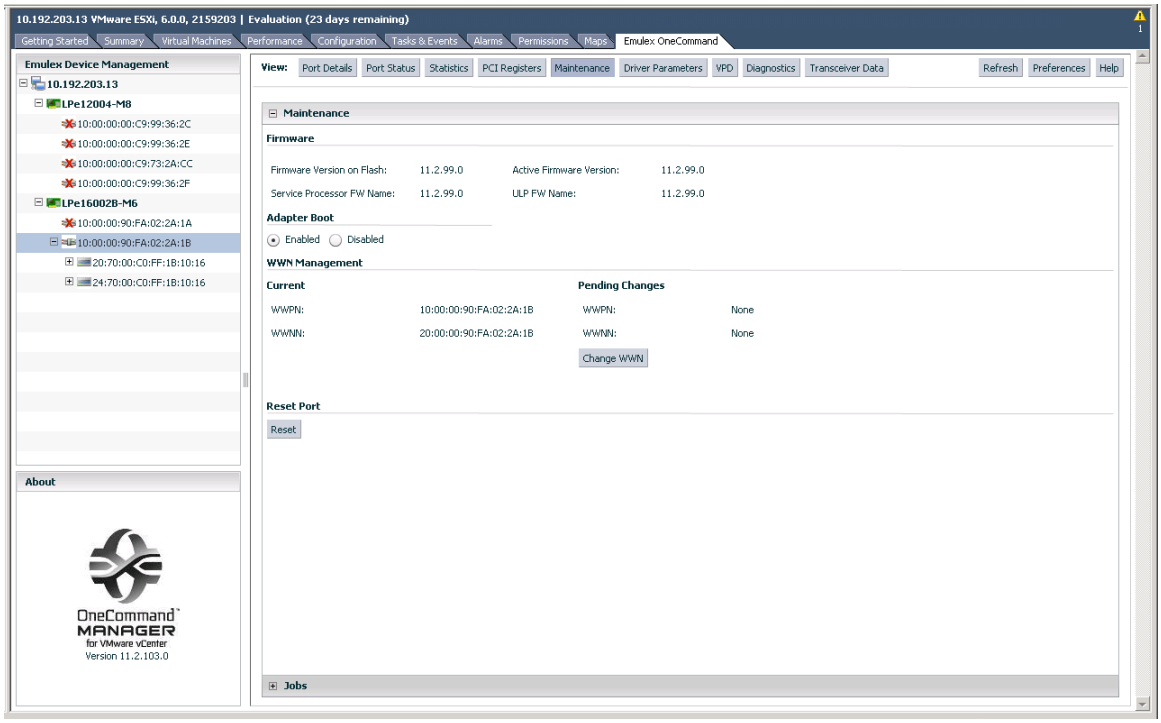

The **Maintenance** tab fields displayed depend on the adapter selected.

- FC Port Firmware area:
  - **Current Version** – The Emulex firmware version number for this adapter.
  - **Initial Load** – The firmware version stub responsible for installing SLI code into its proper slot.
  - **SLI-2 Name** – The name of the SLI-2 firmware overlay.
  - **Kernel Version** – The version of the firmware responsible for starting the driver.
  - **Operational Name** – The name of the operational firmware for the adapter.
  - **SLI-1 Name** – The name of the SLI-1 firmware overlay.
  - **SLI-3 Name** – The name of the SLI-3 firmware overlay.

  For information on updating firmware on an FC port, see Section 7.1, Updating Firmware for an LPe12000-Series Adapter.
- FCoE Port Firmware area:
  - **Firmware Version on Flash** – The firmware version stored on the adapter's non volatile storage. When the system restarts, this version becomes the active firmware version.
  - **Service Processor FW Version** – The operational firmware version on the adapter.
  - **Active Firmware Version** – The firmware version running on the selected adapter.
  - **ULP FW Name** – The firmware version running on the ULP processors within the ASIC.
- WWN Management area (FC and FCoE):
  - Current
    - **WWPN** – The WWPN for the selected adapter port.
    - **WWNN** – The WWNN for the selected adapter port.
  - Pending Changes
    - **WWPN** – If the WWPN has been changed, the new WWPN is displayed in this list. After rebooting, the new WWPN is displayed in the Current list.

- **WWNN** – If the WWNN has been changed, the new WWNN is displayed in this list. After rebooting, the new WWNN is displayed in the Current list.

See Section 6.2.4.1, Changing WWN Configuration for an FC or FCoE Port, for more information about changing the WWN configuration.

For instructions on resetting a port, see Section 6.2.4.2, Resetting an FC or FCoE Port, for more information.

6.2.4.1 Changing WWN Configuration for an FC or FCoE Port

The **Maintenance** tab enables you to change the WWPN and the WWNN of a selected adapter port. For example, you might want to use an installed adapter as a standby if another installed adapter fails. By changing the standby adapter's WWPN or WWNN, it can assume the identity and configuration (such as driver parameters and persistent binding settings) of the failed adapter.

There are three options for referencing WWNs:

- Factory default WWN – The WWN as shipped from the factory.
- Non-volatile WWN – The values that are saved in the non-volatile adapter's flash memory that survive a reboot or a power outage.
- Volatile WWN – A temporary value that is saved in the volatile memory on the flash. If volatile WWNs are set, they are used instead of the non-volatile WWNs. Volatile WWN changes require a warm system reboot to take effect. Volatile WWN changes are lost on systems that power cycle the adapters during the reboot.

CAUTION Changing volatile WWNs result in taking the selected adapter offline. Ensure that this adapter is not controlling a boot device and that all I/O activity on this adapter has stopped before proceeding. This change could result in data loss or corruption.

6.2.4.1.1 Considerations When Changing WWN Configuration for an FC or FCoE Port

- To avoid address conflicts, do not assign a WWNN or WWPN with OneCommand Manager for VMware vCenter if you also use another address management tool.
- The WWPN and WWNN in the Pending Changes list can show **n/a** instead of **None**. This display occurs when the remote host is busy processing some critical task and WWN Management cannot obtain the current state of WWN management.
- In an environment where preboot management exists, a WWPN or WWNN modified by OneCommand Manager for VMware vCenter can be overridden by preboot management, such as Lenovo System X BOFM and industry-standard CLP.

  For example:

  In an environment with CLP/BOFM, OneCommand Manager for VMware vCenter modifies the WWNN or WWPN. OneCommand Manager for VMware vCenter requires a reboot to complete the change. After a reboot, the CLP string is sent during the system boot and rewrites the WWNN or WWPN, or EFIBoot finds the BOFM protocol and uses the default WWNN or WWPN per the BOFM's command.

  In an environment without CLP/BOFM, OneCommand Manager for VMware vCenter modifies the WWNN or WWPN. OneCommand Manager for VMware vCenter requires a reboot to complete the change. The system boots, and the OneCommand Manager for VMware vCenter-modified WWNN or WWPN is used.
- On a system where OneCommand Manager for VMware vCenter is installed, make sure the port numbers configured during the installation are open and dedicated to the OneCommand Manager for VMware vCenter server only. No other service should be listening on this port.

To change a port's WWPN or WWNN, perform these steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Emulex OneCommand** tab.
2. In the Emulex Device Management tree-view, select the port for which you want to change the WWN information.

3. Select the **Maintenance** tab (Figure 32).
4. Click **Change WWN**. The **Change WWN Configuration** dialog is displayed (Figure 34).

Figure 34 Change WWN Configuration Dialog

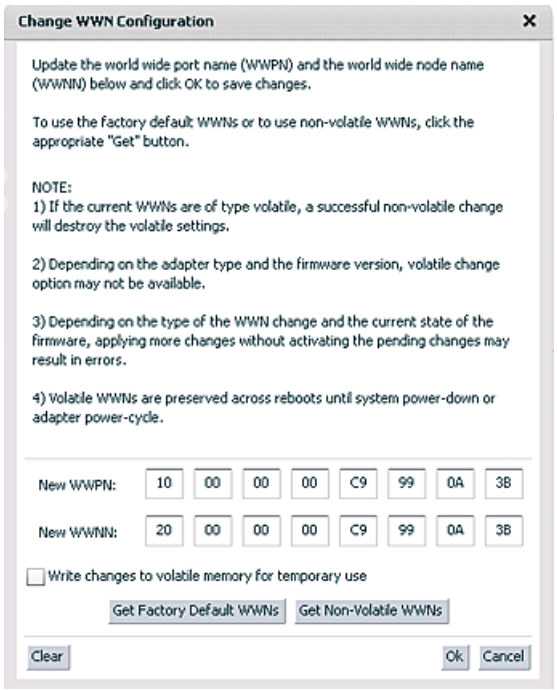

5. Do one of the following:
   - Enter a new WWPN or WWNN.
   - Click **Get Factory Default WWNs** to load the settings that were assigned when the adapter was manufactured. These values can then be modified and saved as volatile or non volatile WWNs.
   - Click **Get Non-Volatile WWNs** to load the current non-volatile WWN. These values can be modified and saved to volatile or non-volatile memory. You can edit the data returned from the button.
6. Check **Write changes to volatile memory for temporary use** to save the **New WWPN** and **New WWNN** settings as volatile WWNs. If cleared, the **New WWPN** and **New WWNN** settings are saved as non-volatile WWNs.

   NOTE If the adapter or firmware does not support volatile WWNs, the **Write changes to volatile memory for temporary use** check box is not selected. This type of change is supported locally and using TCP/IP connections. This check box is not available for remote in-band adapters, regardless of adapter models and firmware version.
7. Click **OK**. The following warning is displayed (Figure 35).

Figure 35 Change WWN Configuration Warning

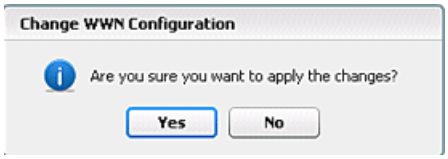

8. Click **Yes**. The New WWPN and new WWNN values are saved. The new WWPN and WWNN appear in the Pending Changes list in the WWN Management area of the **Maintenance** tab.
9. Reboot the system for the changes to take effect (the new WWPN and WWNN appear in the Pending Changes list of the **Maintenance** tab until the system is rebooted). After rebooting, the changes are applied and appear in the Current section of the **Maintenance** tab.

NOTE After changing the WWN of an adapter, you must reboot the system before trying to access the adapter on that system.

6.2.4.2 Resetting an FC or FCoE Port

For FCoE ports, a reset is only necessary to activate changed driver parameters or FIP settings. The reset does not actually perform an adapter level reset of the port.

CAUTION Do not reset an adapter port while copying or writing files. This action could result in data loss or corruption.

NOTE When you reset a port or change WWN configuration on OneCommand Manager for VMware vCenter, do not perform any active management operations on the ESXi host.

To reset an FC or FCoE port, perform these steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Emulex OneCommand** tab.
2. In the Emulex Device Management tree-view, select the FC port you want to reset.
3. Select the **Maintenance** tab (Figure 32 for LPe12000-series adapter ports or Figure 33 for other LightPulse adapter ports).
4. Click **Reset**. The following warning is displayed (Figure 36).

Figure 36 Reset Port Warning

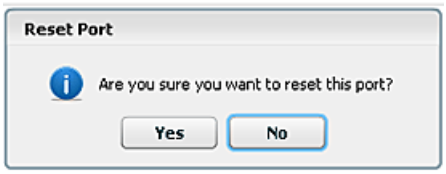

5. Click **Yes**. The adapter port resets. The reset can require several seconds to complete. While the adapter port is resetting, the message `Operation is in progress` is displayed. When the reset is finished, the message `Reset Port Completed` is displayed.

6.2.5 Configuring FC and FCoE Port Driver Parameters

The **Driver Parameters** tab shows driver parameters for a port.

To view driver parameters for a port, perform these steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Emulex OneCommand** tab.
2. In the Emulex Device Management tree-view, select the port whose driver parameters you want to view.
3. Select the **Driver Parameters** tab (Figure 37).

Figure 37 FC and FCoE Port Driver Parameters Tab

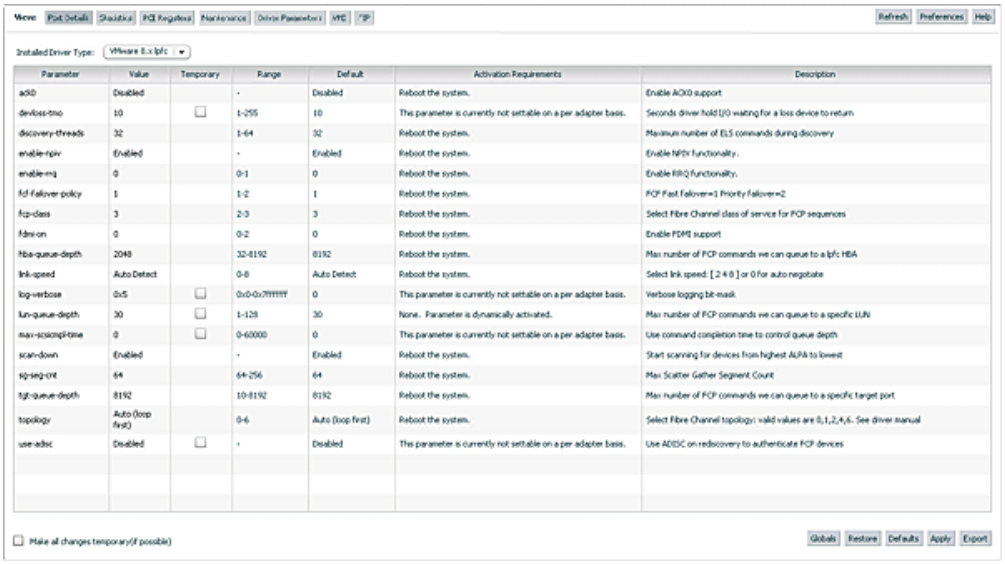

The following **Driver Parameters** tab fields are displayed:

- **Installed Driver Type** – The current driver installed on this host.
- **Port Parameter table** – A list of port driver parameters and their current values.
  - **Parameter** – The driver parameter's name.
  - **Value** – The driver parameter's value.
  - **Temporary** – An indication the value is temporary.
  - **Range** – The range of acceptable values for the driver parameter.
  - **Default** – The driver parameter's default value.
  - **Activation Requirements** – The steps required to activate the changed value of the driver parameter.
  - **Description** – The driver parameter's description.

To change the driver parameters for a port using the **Value** field, perform these steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Emulex OneCommand** tab.
2. In the Emulex Device Management tree-view, select the port for which you want to change the driver parameters.
3. Select the **Driver Parameters** tab (Figure 37).
4. In the driver parameter table, click the **Value** field of a parameter that you want to change. The range for the value is displayed. Enter a value in decimal or hexadecimal format, depending on how the current value is presented. If the value is in hexadecimal format, it is prefaced by 0x (for example, 0x2d). You can enter a hexadecimal value without the 0x. For example, if you enter `ff10`, this value is interpreted and displayed as 0xff10.
5. If you want the change to be temporary (causing the parameter to revert to its last permanent setting when the system is rebooted), select the **Temporary** check box. This option is available only for dynamic parameters.
6. If you are making changes to multiple parameters, and you want all the changes to be temporary, select the **Make all changes temporary** check box. This setting overrides the setting of the **Make change temporary** check box. Only dynamic parameters can be made temporary.
7. Click **Apply**.

To set a port parameter value to the corresponding host parameter value, click **Globals**. All parameter values are set the same as the global, or host, values.

To apply the global values, click **Apply**.

If you changed parameters but did not click **Apply,** and you want to restore the parameters to their last saved values, click **Restore**.

To reset all parameter values to their default (factory) values, click **Defaults**.

To save driver parameters to a file, perform these steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Emulex OneCommand** tab.
2. In the Emulex Device Management tree-view, select the FC port for which you want to change the driver parameters.
3. Select the **Driver Parameters** tab (Figure 37).
4. Click **Export** to create and save a desired port parameter configuration. Each definition is saved in a comma-delimited file with the following format:

   <parameter-name>=<parameter-value>
5. Click **Apply** to apply your configuration changes.

6.2.6 Viewing FC or FCoE Port Vital Product Data (VPD)

The **VPD** tab shows vital product data (if available) for the selected FC or FCoE port, such as the product name, part number, serial number, and so on.

To view VPD information for a port, perform these steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Emulex OneCommand** tab.
2. In the Emulex Device Management tree-view, select the port whose VPD information you want to view.
3. Select the **VPD** tab (Figure 38).

Figure 38 Port VPD Tab

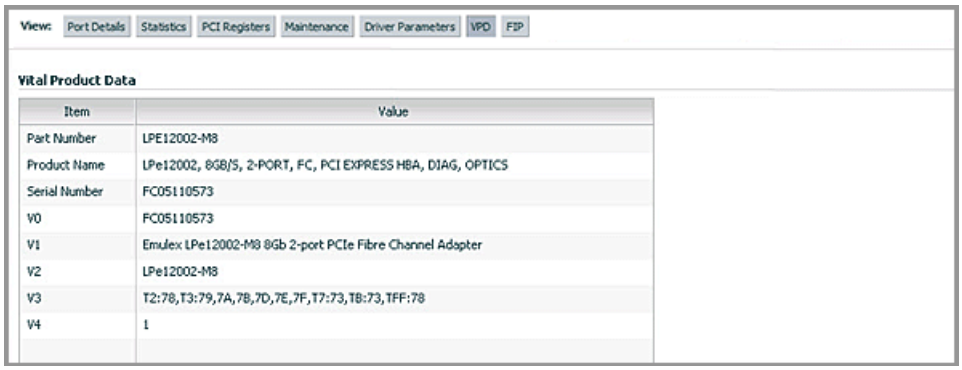

The following **VPD** tab fields are displayed:

- **Part Number** – The adapter's part number.
- **Product Name** – The product information about the selected adapter port.
- **Serial Number** – The adapter's serial number.
- **VO** – Vendor-unique data. V indicates a vendor-specific field. An adapter may have none, one, or more of these fields defined. Valid values for this field are VO (the letter O, not the number zero) and V*x* (where *x* is a number).

NOTE Some adapters may show additional VPD information such as EC and MN.

6.2.7 Configuring FIP

The **FIP** tab allows you to view configure FIP for an FCoE adapter port.

To configure FIP, perform these steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Emulex OneCommand** tab.
2. In the Emulex Device Management tree-view, select the FCoE adapter port whose FIP properties you want to configure.
3. Select the **FIP** tab (Figure 39).
4. Set the parameters you want, and click **Apply Changes**.

Figure 39 FCoE Port FIP Tab

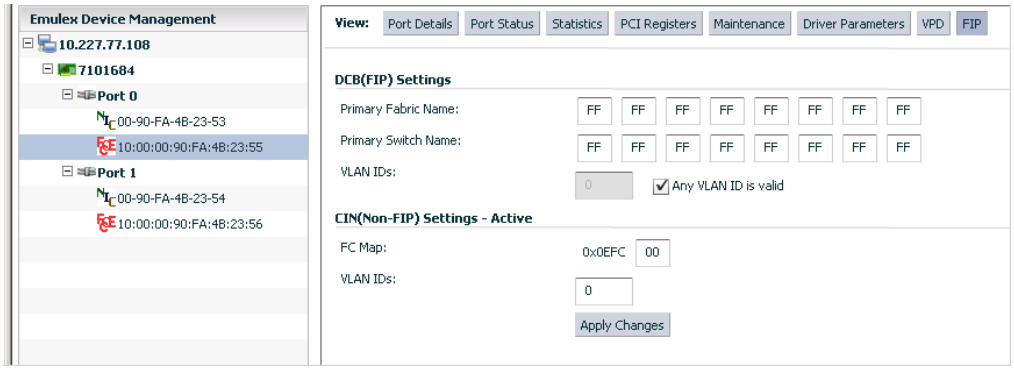

The following **FIP** tab fields are displayed:

- DCB (FIP) Settings area: DCB settings are only applicable when the DCBX Mode in the **DCB** tab is set to **DCB**. See Chapter 8 for more information.

  NOTE If the DCB (FIP) Settings area has Active in its title, this indicates that DCB (FIP) is the current DCBX mode (DCBX mode is set from the **DCB** tab). If CIN (Non-FIP) is the current DCBX mode, Active is in the CIN (Non-FIP) Settings area.

  - **Primary Fabric Name** – Indicates the FC fabric's WWN to which to connect. If the Primary Fabric Name is wild (all 0xFFs), then connection to any fabric name is allowed.
  - **Primary Switch Name** – Indicates the FC switch's WWN to which to connect. If the Primary Switch Name is wild (all 0xFFs) then connection to any switch name is allowed.
  - **VLAN ID** – Determines the VLAN where the adapter FCoE services are available. It can have a value from 0 to 4095 and supports wildcard values if the **Any VLAN ID is valid** check box is selected.
  - **Any VLAN ID is valid** check box – When checked, the **VLAN ID** field of the FCoE forwarder can be any valid value.
- CIN Settings area:
  - **FC Map** – Enter the bit value that completes the FPMA.
  - **VLAN ID** – Determines the VLAN where the adapter FCoE services are available. It can have a value from 0 to 4095.

5. Click **Apply Changes** to set the new FIP parameters.

6.2.8 Viewing FC and FCoE Port Transceiver Information

The **Transceiver Data** tab shows transceiver information such as vendor name, serial number, and part number. If the adapter or transceiver does not support some or all of the transceiver data, the fields display N/A.

To view transceiver information for an FC or FCoE port, perform these steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Emulex OneCommand** tab.
2. In the Emulex Device Management tree-view, select the FC or FCoE port whose transceiver information you want to view.
3. Select the **Transceiver Data** tab (Figure 40).

Figure 40 Port Transceiver Data Tab

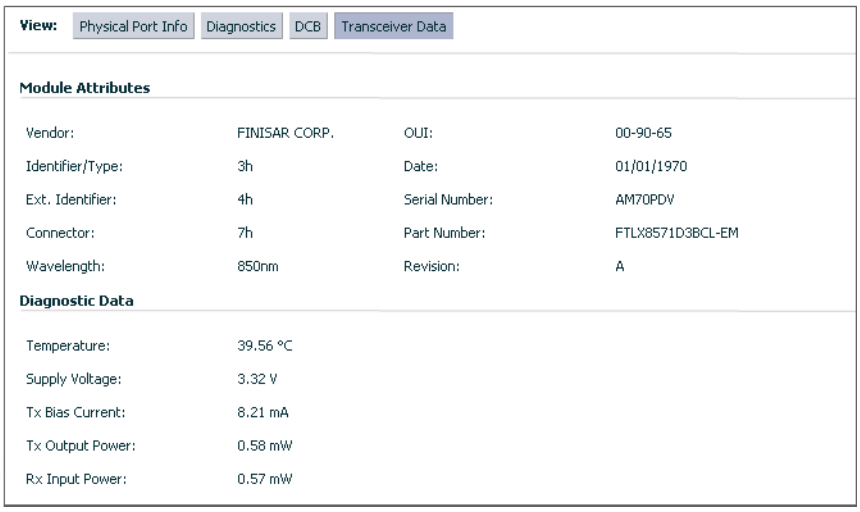

The following **Transceiver Data** tab fields are displayed:

- Module Attributes area:
  - **Vendor** – The name of the vendor.
  - **Identifier/Type** – A value that specifies the physical device described by the serial information.
  - **Ext. Identifier** – Additional information about the transceiver.
  - **Connector** – The external optical or electrical cable connector provided as the media interface.
  - **Wavelength** – The nominal transmitter output wavelength at room temperature.
  - **OUI** – The vendor's Organizationally Unique Identifier. It is also known as the IEEE Company Identifier for the vendor.
  - **Date** – The vendor's date code in the MM/DD/YY format.
  - **Serial Number** – The serial number provided by the vendor.
  - **Part Number** – The part number provided by the SFP vendor.
  - **Revision** – The vendor revision level.
- Diagnostic Data area:
  - **Temperature** – The internally measured module temperature.
  - **Supply Voltage** – The internally measured supply voltage in the transceiver.
  - **Tx Bias Current** – The internally measured Tx bias current.
  - **Tx Output Power** – The measured Tx output power.
  - **Rx Input Power** – The measured Rx input power.

6.2.9 Viewing Flash Contents for an FC Port

To view the flash contents for an FC port, perform these steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Emulex OneCommand** tab.
2. In the Emulex Device Management tree-view, select the FC port whose Flash contents you want to view.
3. Select the **Flash Contents** tab. The Flash Contents information is displayed (Figure 41).

Figure 41 FC Port Flash Contents

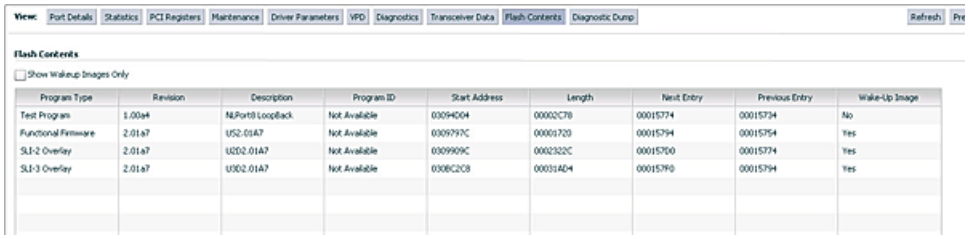

4. Select the **Show Wakeup Images Only** check box if you want to see only the flash contents with the wake-up images.

6.2.10 Viewing FC and FCoE Target Information

When you select a port target associated with an FC or FCoE adapter from the Emulex Device Management tree-view, the **Target Information** tab shows information associated with that target.

To view FC or FCoE target information, perform these steps:

1. Select a host in the vSphere console tree-view, and select the **Emulex OneCommand** tab.
2. In the Emulex Device Management tree-view, select the FC or FCoE target port whose information you want to view. The **Target Information** tab appears (Figure 42).

Figure 42 Target Information Tab

The following **Target Information** tab fields are displayed:

- **FC ID** – The FC ID for the target; assigned automatically in the firmware.
- **SCSI Bus Number** – The SCSI bus number to which the target is mapped.
- **SCSI Target Number** – The target's identifier on the SCSI bus.
- **Node WWN** – A unique 64-bit number, in hexadecimal, for the target (N_PORT or NL_PORT).

- **Port WWN** – A unique 64-bit number, in hexadecimal, for the fabric (F_PORT or Switched Fabric Loop Port [FL_PORT]).
- **OS Device Name** – The operating system device name.

6.2.11 Viewing FC and FCoE LUN Information

When you select a LUN associated with an FC or FCoE adapter from the Emulex Device Management tree-view, the LUN Info pane shows information associated with that LUN.

To view LUN information, perform these steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Emulex OneCommand** tab.
2. In the Emulex Device Management tree-view, select the FC or FCoE LUN whose information you want to view. The corresponding LUN Info pane is displayed (Figure 43).

Figure 43 LUN Info Tab (ExpressLane LUN with Frame Priority Selection Supported)

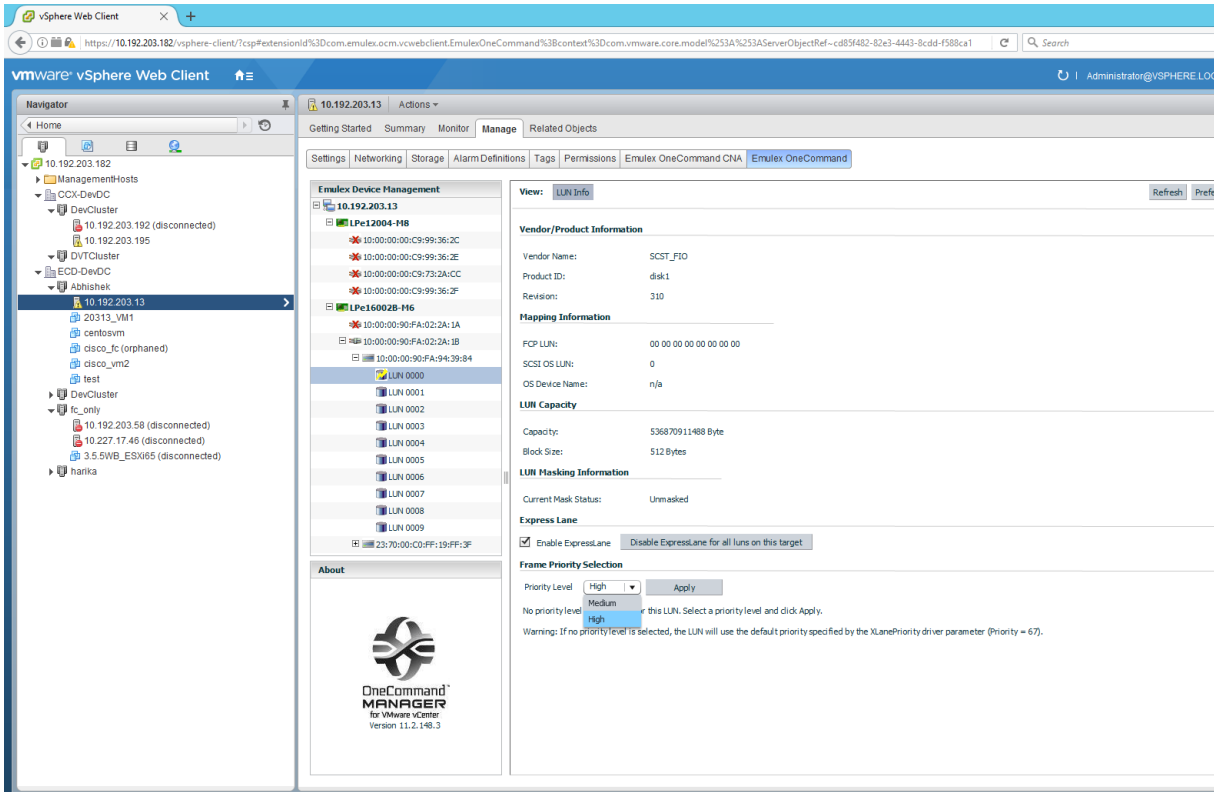

The following **LUN Information** tab fields are displayed:

- Vendor Product Information area:
  - **Vendor Name** – The name of the vendor of the LUN.
  - **Product ID** – The vendor-specific ID for the LUN.
  - **Revision** – The vendor-specific revision number for the LUN.
- Mapping Information area:
  - **FCP LUN** – The FC identifier used by the adapter to map to the operating system LUN.
  - **SCSI OS LUN** – The SCSI identifier used by the operating system to map to the specific LUN.
  - **OS Device Name** – The name assigned by the operating system to the LUN.

- LUN Capacity area: LUN capacity information is only provided when the LUN is a mass-storage (disk) device. Other devices, such as tapes and scanners, do not display capacity.
  - **Capacity** – The capacity of the LUN, in megabytes.
  - **Block Size** – The length of a logical unit block in bytes.
- LUN Masking Information area:
  - **Current Mask Status** – The current status is masked or unmasked.
- ExpressLane area (Only appears if supported by the adapter):
  - **Enable ExpressLane** – Select this check box to create an ExpressLane LUN.
  - Priority Level – When ExpressLane is enabled, you can assign the frame priority to the ExpressLane LUN.
  - **Disable ExpressLane on all LUNs** button – Clicking this button disables ExpressLane on all ExpressLane LUNs.
  - **Apply** button - Click this button to save your ExpressLane changes

The adapter information that is displayed depends upon the type of adapter you select.

6.2.11.1 Enabling an Optimized Access Storage (ExpressLane) LUN

Enabling an ExpressLane LUN is supported only for the LPe16000-series, LPe31000-series, and LPe32000-series adapters, and, is supported only on ESXi 5.5 and later versions.

NOTES

- An ExpressLane LUN must be enabled in the driver before it can be used by the OneCommand Manager for VMware vCenter. Additionally, the priority value to mark each of the ExpressLane frames must be specified to the FC driver through new driver parameters.
- ExpressLane is not supported for LUNs attached to virtual ports.
- Masked LUNs cannot be ExpressLane enabled because they are not presented to the host. Conversely, ExpressLane LUNs cannot be masked.

In host view, the **ExpressLane LUN** icon  is displayed in the Emulex Device Management area and represents a LUN with ExpressLane priority queuing enabled.

If your system contains flash and rotational storage I/O on the same initiator port, the improved performance of the flash storage can be hindered by the slower rotational storage I/O performance. Because both are sharing the same queue, rotational storage I/O requests queued up in front of flash storage I/O requests can prevent the performance gains of flash storage I/O (that is, head of line blocking).

To prevent this situation, you can create a separate priority queue in the driver to process the requests ahead of the normal queue. Additionally, the I/O frames placed on this queue can be marked with a priority so that the switch and target also prioritizes the flash I/O requests. By indicating to the FC driver which LUNs receive this special priority queuing (for example, ExpressLane LUNs), performance on these LUNs can be greatly improved. Even if only FC flash storage is available on your system, there can be LUNs that you want to run I/O with a higher priority than other LUNs.

To enable an ExpressLane LUN, perform these steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Emulex OneCommand** tab.
2. From the discovery-tree, select a LUN under the adapter on which you want to enable ExpressLane. The **LUN Information** tab appears (Figure 43).
3. Select the **Enable ExpressLane** check box.

   The LUN icon in the discovery-tree changes to the **ExpressLane LUN** icon  .

NOTE The ExpressLane state of a LUN is enabled at the driver (active) level and the configured level (persists across reboots).

To disable an ExpressLane LUN, perform these steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Emulex OneCommand** tab.
2. From the discovery-tree, select a LUN under the adapter on which you want to disable ExpressLane. The **LUN Information** tab appears (Figure 43).
3. Clear the **Enable ExpressLane** check box to disable the selected LUN.

To disable all ExpressLane LUNs, perform these steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Emulex OneCommand** tab.
2. From the discovery-tree, select a LUN under the adapter on which you want to disable ExpressLane. The **LUN Information** tab appears (Figure 43).
3. Click **Disable ExpressLane for all LUNs on this target**.
4. A dialog appears warning you that you are about to delete all ExpressLane LUNS on this target. Click **OK**.

   All **ExpressLane LUN** icons in the discovery-tree (for the selected adapter port) will change to the regular **LUN** icon and any assigned frame priority is removed.

6.2.11.2 Selecting a Frame Priority

If the adapter and switch support it, you can assign a frame priority to the ExpressLane LUN. Switches can provide up to three priority levels; **Low**, **Medium**, and **High,** but they might provide fewer options.

To select a frame priority, perform these steps:

NOTE The `EnableXLane` driver parameter must be enabled on the **Driver Parameters** tab to set frame priorities.

1. Select a host in the vSphere console tree-view. If applicable, select the **Emulex OneCommand** tab.
2. From the discovery-tree, select a LUN under the adapter on which you want to enable ExpressLane. The **LUN Information** tab appears (Figure 43).
3. Select the **Enable ExpressLane** check box if it is not already checked. The **LUN** icon in the discovery-tree changes to the **ExpressLane LUN** icon.
4. Select a frame priority from the **Priority Level** pull-down menu.

NOTE If the switch connected to the FC or FCoE initiator does not support LUN specific frame priority levels using the Get Fabric Object (GFO), you must manually enter the frame priority values in the range of 0–127 for all ExpressLane enabled LUNs as depicted in Figure 44.

Figure 44 LUN Information Tab (Frame Priority is not Supported by the Switch)

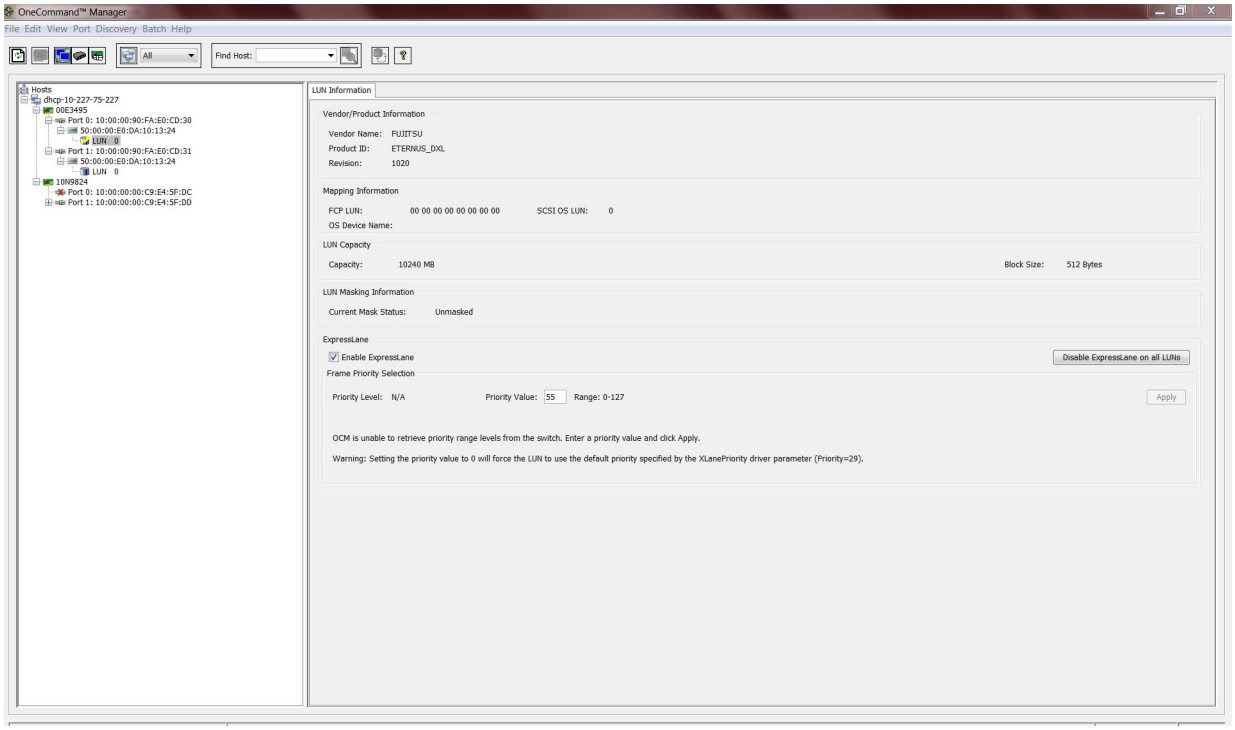

5. Click **Apply**.

If there were problems assigning the frame priority, the **LUN Information** tab displays a message with a suggested solution.

Chapter 7: Updating Firmware

OneCommand Manager for VMware vCenter enables you to update firmware for a single adapter or simultaneously for multiple adapters.

The submitted firmware update job can be tracked in the VMware tasks window.

Multiple firmware update jobs can be submitted for different adapters and ports on the same or different ESXi hosts simultaneously to OneCommand Manager for VMware. However, only a single job is processed on a given ESXi host. The remaining jobs on that host will be queued and processed sequentially.

The firmware update progresses in the background until all the jobs are completed. During this period, you can still browse through the other tabs. However, if you start a firmware update and log out from the vSphere console before the firmware update is completed, all pending jobs fail.

NOTE If a secure version of firmware (version 11.0 or later) is installed on an LPe16000-series adapter and you want to update to an earlier unsecured version of firmware, you must remove the secure firmware jumper block before performing the update. Refer to the installation guide for the adapter for more information.

If you attempt to update unauthenticated firmware for a secure LPe16000-series, LPe31000-series, or LPe32000-series adapter, the following error message will be displayed.

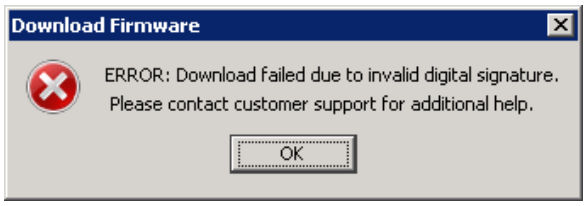

If you attempt to update unsecured firmware for a secure LPe16000-series, LPe31000-series, or LPe32000-series adapter, the following error message will be displayed.

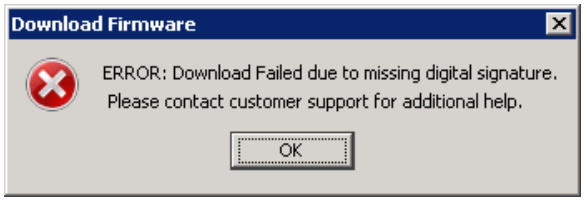

Contact customer support for more information.

7.1 Updating Firmware for an LPe12000-Series Adapter

CAUTION Updating firmware or boot code on an LPe12000-series adapter that is being used to boot from SAN is not advisable. After the update has completed, an adapter reset is issued that may cause a loss of connectivity to the SAN and possible loss of data. To update firmware on an LPe12000-series adapter, you must make sure that the adapter is not currently being used to boot from SAN.

Do one of the following:

- Move the adapter to be updated to a non-boot from SAN host, and perform the update from that location.
- If the host with the target adapter is also hosting other boot from SAN adapters, perform a boot from SAN using one of the other boot from SAN adapters. The target adapter can now be updated.

To update firmware for a port on an LPe12000-series adapter, perform these steps:

1. Select a host in the vSphere console tree-view; and, if applicable, select the **Emulex OneCommand** tab (Figure 5).
2. In the Emulex Device Management tree-view, select LPe12000-series adapter port for which you want to update firmware.
3. Select the **Maintenance** tab (Figure 45).

Figure 45 Maintenance Tab

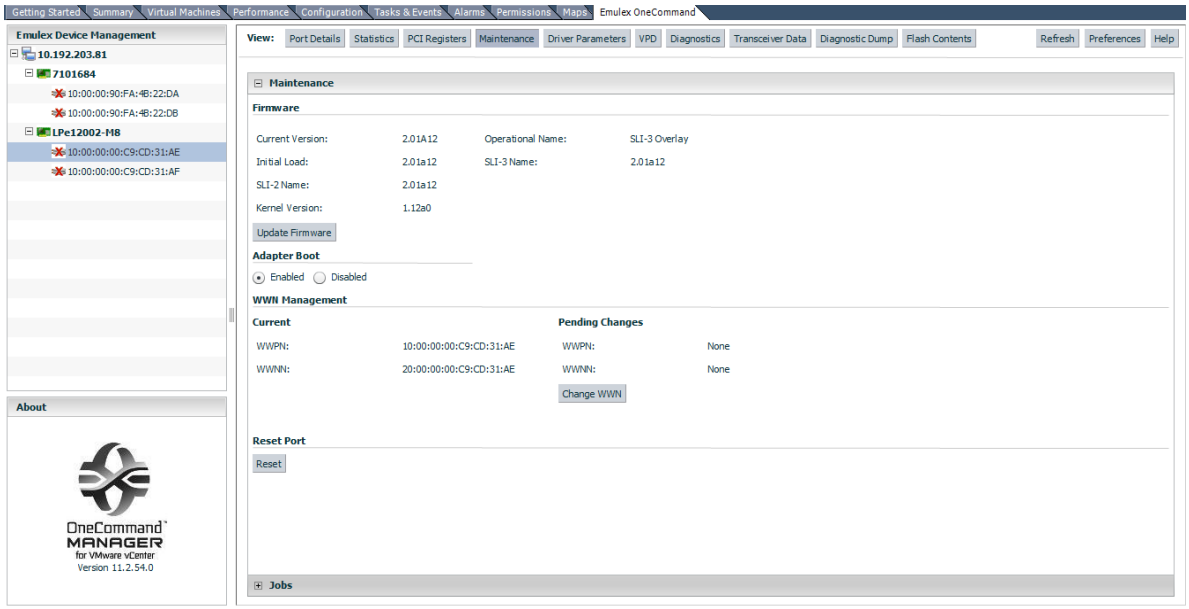

4. Click **Update Firmware**. If a warning screen appears, click **Yes**. The **Firmware Download** dialog is displayed (Figure 46).

Figure 46 Firmware Download Dialog

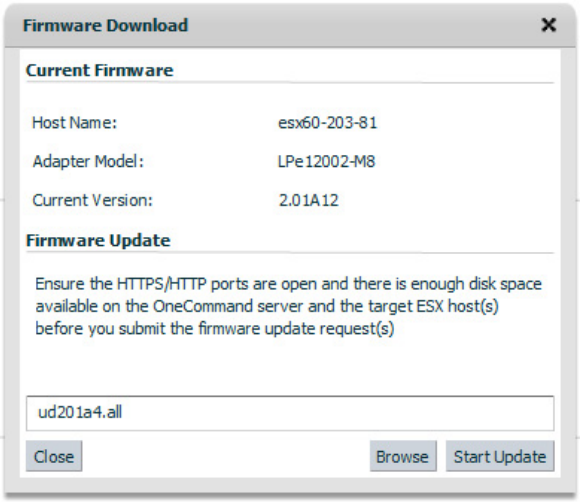

5. Click **Browse** and navigate to the unzipped, extracted image file you want to download.
6. On the browse window, select the file and click **OK**. The **Firmware Download** dialog appears.
7. Click **Start Update**. A message prompting you to confirm the firmware update appears.
8. Click **Yes**. When the update begins, the **Jobs** window is displayed.(Figure 47)

Figure 47 Jobs Window

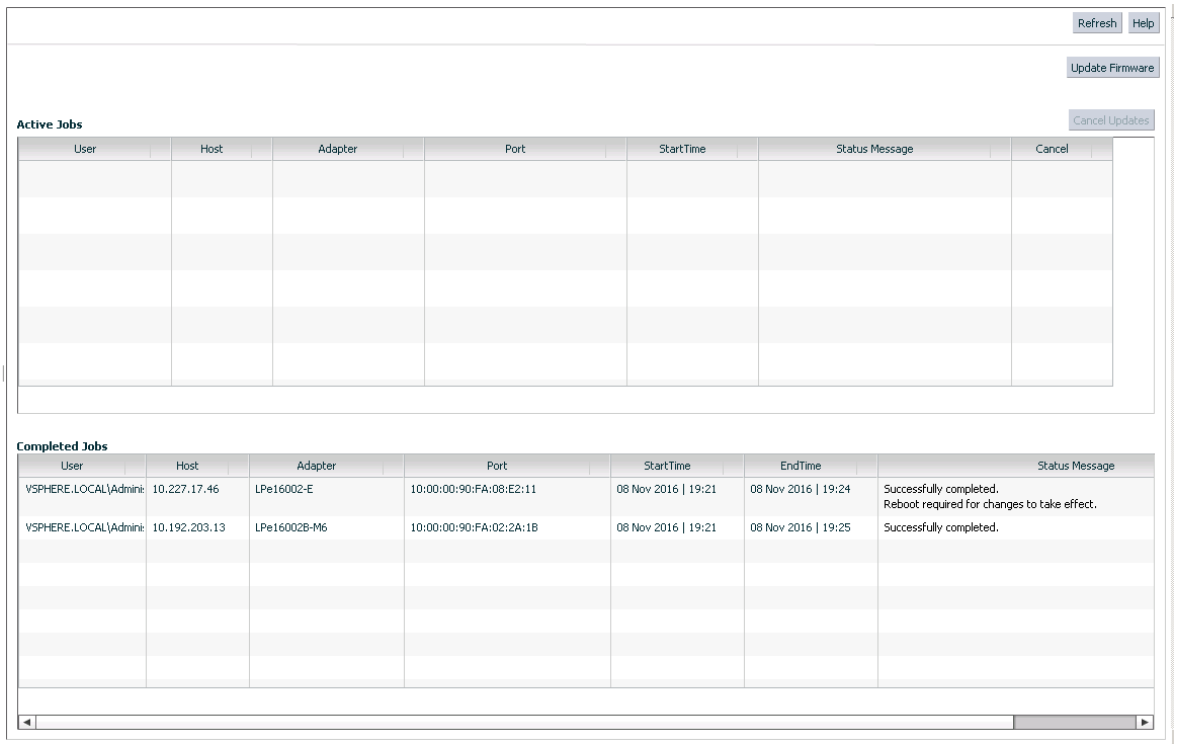

A status message in the Active Job list shows the progress of the download. The ports on which the firmware is being downloaded have the status **Job is in progress**; the others have **Waiting in queue to start**. When the download is completed, the entry moves to the Completed Jobs list. The **status message** field in the Completed Jobs list shows the status of the completed job.

NOTE The firmware update progresses in the background until all the jobs are completed. During this period, you can still browse through the other tabs.

If you start a firmware update and logout from the vSphere console before the firmware update is completed, all pending jobs fail.

7.1.1 Updating Firmware on an LPe12000-Series Adapter in a Host

On the system where OneCommand Manager for VMware vCenter is installed, make sure the port numbers configured during the installation are open and dedicated to the OneCommand Manager for VMware vCenter server only. No other service should be listening on this port.

Before you can perform a batch update, the firmware file must be downloaded from http://www.broadcom.com and extracted to a directory on your local drive.

To update firmware for compatible adapters, perform these steps:

1. Select a host in the vSphere console tree-view, and, if applicable, select the **Emulex OneCommand** tab.
2. On the **Batch Firmware Update Information** pane, click **Update Firmware**.

   NOTE Do not select a particular tree element for this operation.
3. Select the **Maintenance** tab and click **Update Firmware**. If a warning message appears, click **Yes**. The **Batch Firmware Download** dialog is displayed (Figure 48).

Figure 48 Batch Firmware Download Dialog

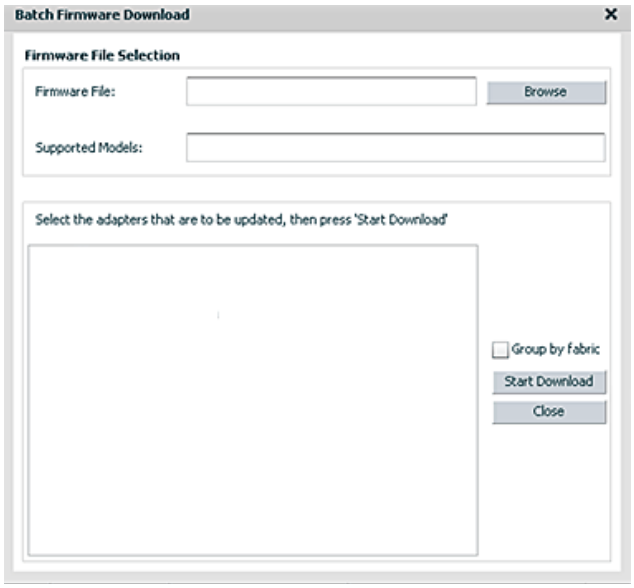

4. Click **Browse**. The **Firmware File Selection** dialog appears. Select the file you want to use and click **OK**. A status message appears indicating that OneCommand Manager for VMware vCenter is searching for compatible adapters.

   After compatible adapters are found, the **Firmware File** field of the **Batch Firmware Download** dialog shows the selected image file name. The **Supported Models** field shows a list of all adapter models that are compatible with the selected image file. The set of compatible adapters appears in the dialog's tree-view (Figure 49).

Figure 49 Populated Batch Firmware Download Dialog

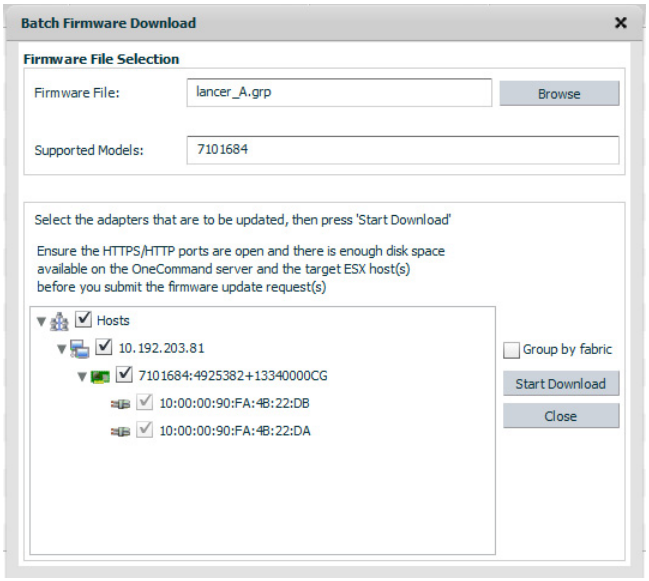

The tree-view shows all adapters and their corresponding hosts for which the selected firmware file is compatible. Use the check boxes next to the host and adapter entries to select or deselect an entry. Selecting or clearing an adapter selects or removes that adapter, respectively; selecting a host removes or selects all eligible adapters for that host.

To view the compatible adapters in a fabric-centric mode, select the **Group by Fabric** check box.

For adapters where each individual port can have new firmware installed, you can select the ports on the adapter to which you want to download firmware.

5. Make your selections and click **Start Download**. When a message prompting you to confirm the firmware update appears, click **Yes**.

   When the update begins, a status message in the Active Job list shows the progress as either **Job is in progress** or **Waiting in queue to start**. The ports on which the firmware is being downloaded have the status **Job is in progress**; the others have **Waiting in queue to start**. You can select the check box to cancel the jobs with status **Waiting in queue to start**. When download is completed, the entry moves to the Completed Jobs list. The **status message** field in the Completed Jobs list shows the status of the completed job.

   NOTE If you start a firmware update and log out from the vSphere console before the firmware update is completed, all pending jobs fail.

   The firmware update progresses in the background until all the jobs are completed. During this period, you can still browse through the other tabs.

7.2 Updating Firmware for LPe16000-Series, LPe31000-Series, and LPe32000-Series Adapters

For all LightPulse adapters except LPe12000-series adapters, you update the firmware for the entire adapter and not for individual ports.

To update firmware for an adapter, perform these steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Emulex OneCommand** tab.

2. In the Emulex Device Management tree-view, select the adapter for which you want to update firmware information.
3. Select the **Maintenance** tab (Figure 50).

Figure 50 Maintenance Tab

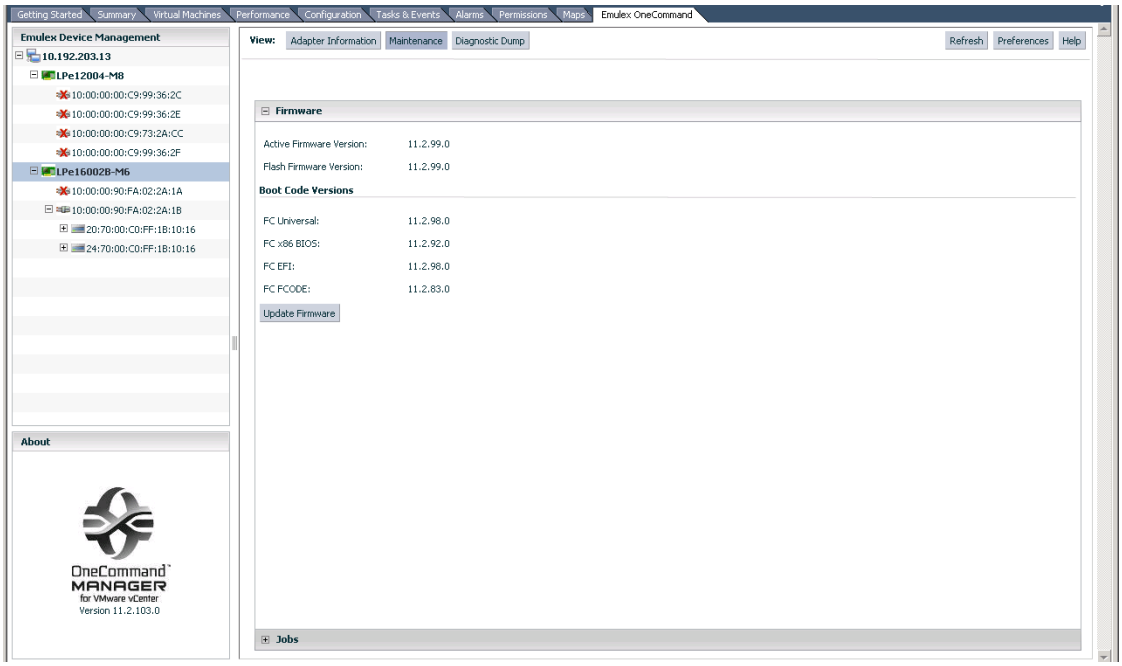

4. Click **Update Firmware**. If a warning message appears, click **Yes**. The **Firmware Download dialog** is displayed (Figure 51).

Figure 51 Firmware Download Dialog

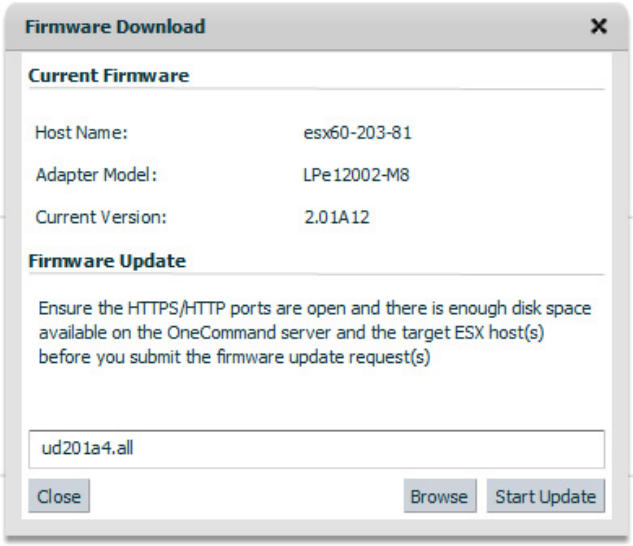

5. Click **Browse** and navigate to the unzipped, extracted image file you want to download.
6. On the browse window, select the file and click **OK**.

7. Click **Start Update**.

   When the update begins, the **Jobs** window is displayed (Figure 56). A status message in the Active Job list shows the progress of the download. The ports on which the firmware is being downloaded have the status **Job is in progress**, the other ports display the **Waiting in queue to start** status.

   When the download has completed, the entry moves to the Completed Jobs list. The **status message** field in the Completed Jobs list shows the status of the completed job.

   You can monitor jobs submitted through the OneCommand Manager for VMware vCenter command line interface. If you start a firmware update and log out from the vSphere console before the firmware update is completed, all pending jobs fail. The firmware update progresses in the background until all the jobs are completed. During this period, you can still browse through the other tabs. The firmware update job submitted can also be tracked in the VMware tasks window.

8. Open **Firmware Summary**, and the updated firmware information for the selected adapter is displayed.

7.2.1 Performing a Batch Firmware Update in Cluster View

Before you can perform a batch update, you must download the firmware file from http://www.broadcom.com and extract it into a directory on your local drive.

NOTE On the system where OneCommand Manager for VMware vCenter is installed, make sure the port numbers configured during the installation are open and dedicated to the OneCommand Manager for VMware vCenter server only. No other service should be listening on this port.

To perform a batch firmware update on a host, perform these steps:

1. Select a cluster in the vSphere console tree-view, and if applicable, select the **Emulex OneCommand** tab.
2. From the Emulex Device Management options, select **Batch Firmware Update** (Figure 52). The **Batch Firmware Download** dialog appears (Figure 48).

Figure 52 Batch Firmware Update Information

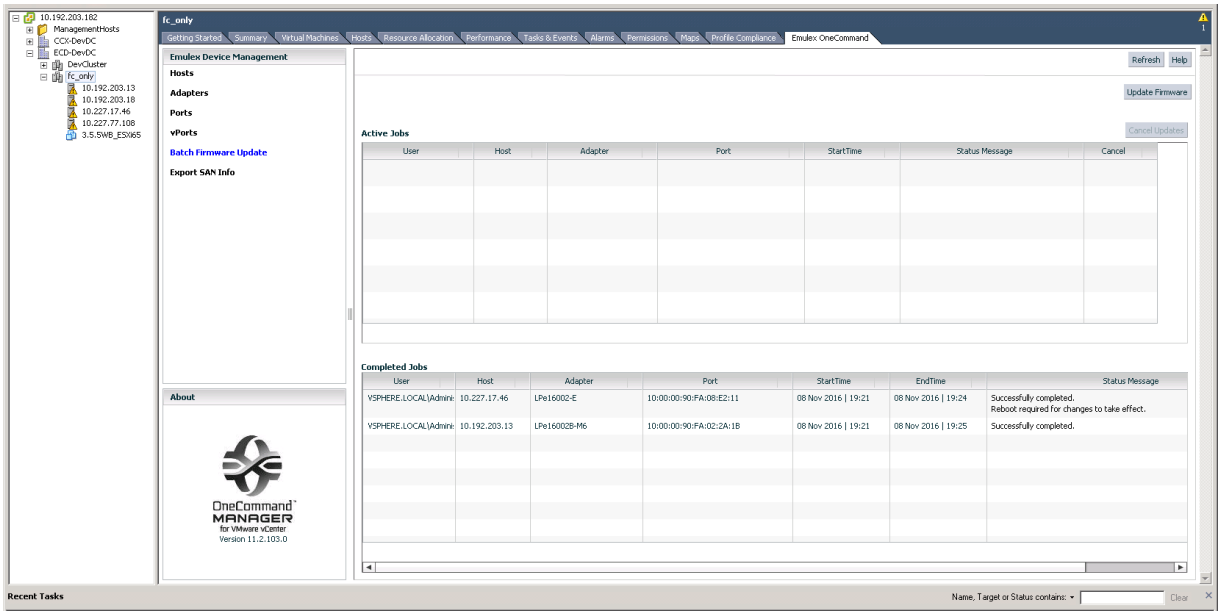

3. Click **Start Download** to install firmware on multiple adapters in a single step. Batch firmware loading is restricted to a single firmware file and to all accessible adapters for which that file is compatible.

7.2.2 Updating Firmware on Multiple Adapters in a Host

To update firmware for multiple adapters on a single host, perform these steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Emulex OneCommand** tab.
2. Select the **Maintenance** tab and click **Update Firmware**. If a warning message appears, click **Yes**. The **Batch Firmware Download** dialog is displayed (Figure 53).

Figure 53 Batch Firmware Update Dialog

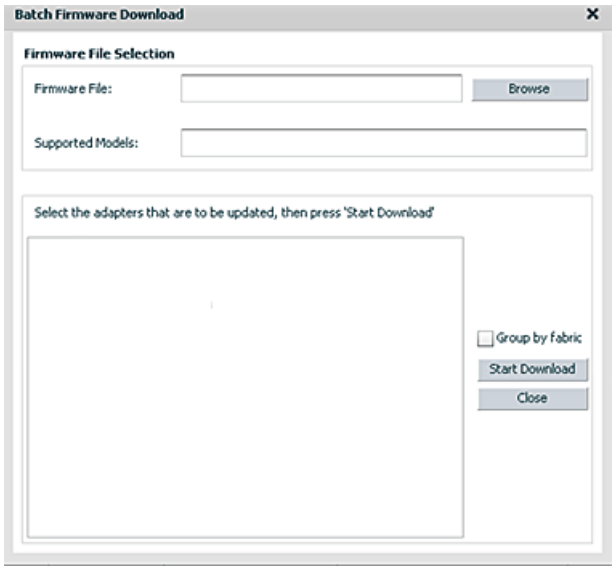

NOTE Do not select a particular tree element for this operation.

3. Click **Browse,** and a search dialog appears. On the search dialog, select the file you want to use and click **OK**. A status message appears indicating that OneCommand Manager for VMware vCenter is searching for compatible adapters. After compatible adapters are found, the following is displayed in the **Batch Firmware Download** dialog (Figure 54):
   - **Firmware File** – This field shows the selected image file name.
   - **Supported Models** – This field shows a list of all adapter models that are compatible with the selected image file.
   - The set of compatible adapters appears in the dialog's tree-view.
4. To view the compatible adapters in host-centric mode, make sure the **Group by Fabric** check box is not selected.

Figure 54 Batch Firmware Download Dialog, Host-Centric View

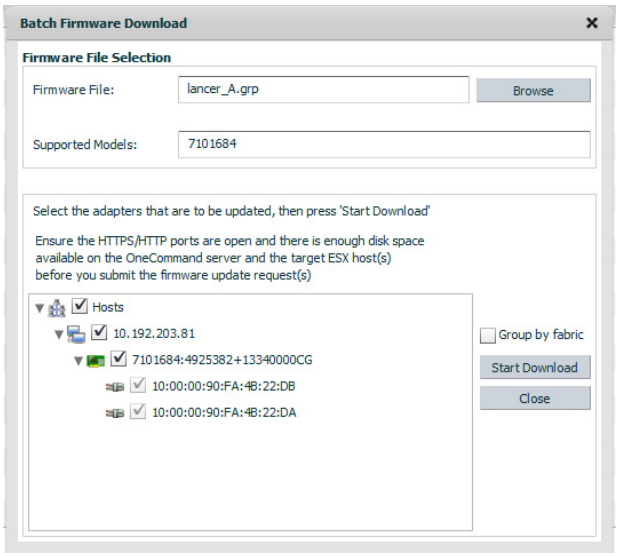

5. To view the compatible adapters in a fabric-centric mode, select the **Group by Fabric** check box.

Figure 55 Batch Firmware Download Dialog, Fabric-Centric View

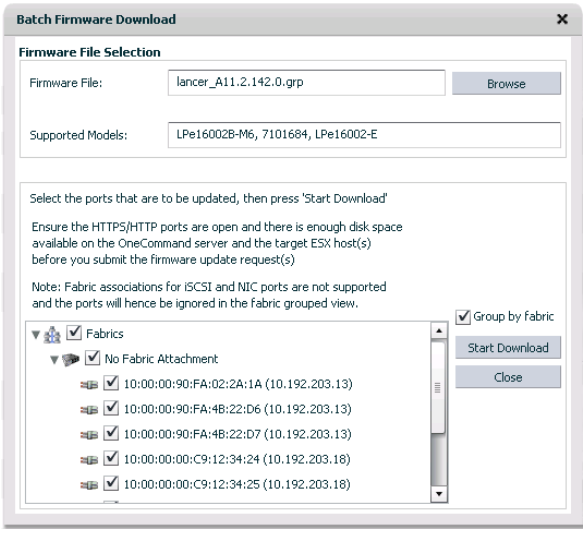

The tree-view shows all adapters and their corresponding hosts or fabrics for which the selected firmware file is compatible. Use the check boxes next to the entries to select or deselect a host, fabric, adapter, or port (if the adapters where the individual port can have new firmware downloaded).

6. Make your selections and click **Start Download**. A message prompting you to confirm the firmware update appears, click **Yes**.

7.2.3 Jobs Window

When the download begins, the **Jobs** window is displayed. A status message in the Active Job list shows the progress of the firmware download. The ports on which the firmware is being downloaded have the status **Job is in progress**,

the other ports have **Waiting in queue to start**. When the download is completed, the entry moves to the Completed Jobs list. The **status message** field in the Completed Jobs list shows the status of the completed job (Figure 56).

Figure 56 Completed Jobs Window

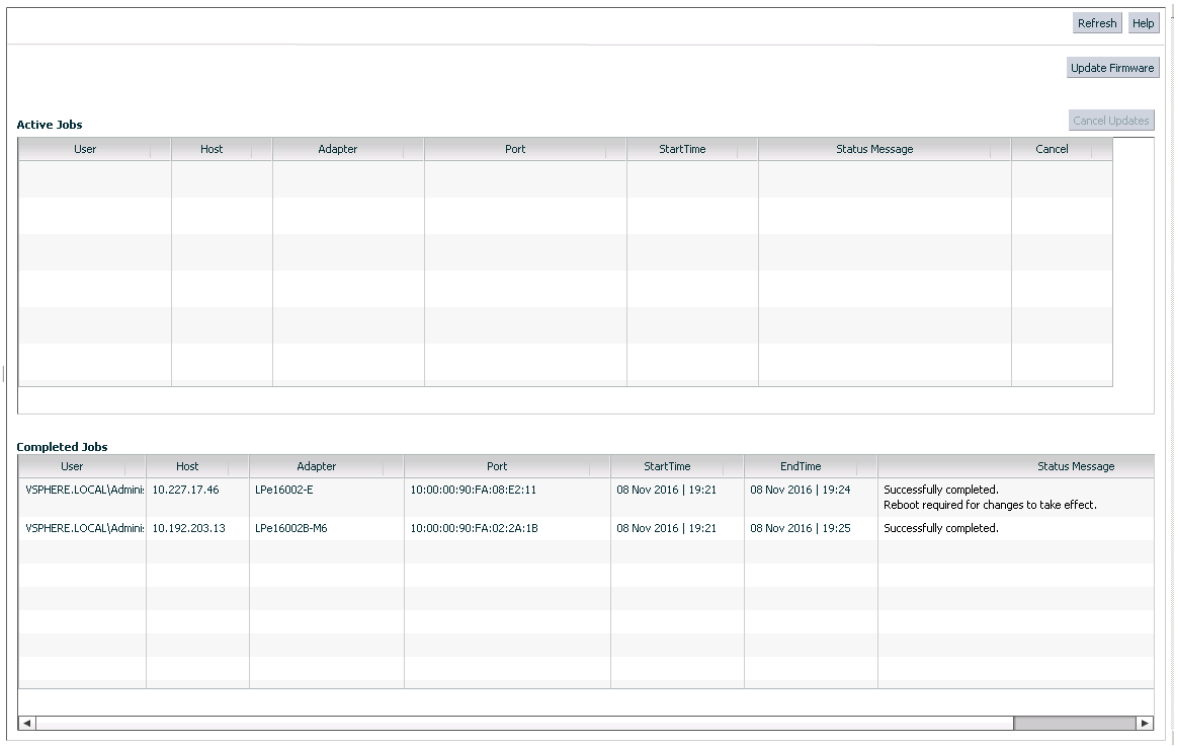

The following **Jobs** window fields are displayed:

- Active Jobs list:
  - **User** – The user who updated the firmware.
  - **Host** – The IP address of the host to which the adapter or port belongs.
  - **Adapter** – The model of the adapter selected for the firmware update request.
  - **Port** – The Port WWN for an FC adapter. This field is blank for a CNA.
  - **Start Time** – The time the job is submitted.
  - **Status Message** – The status of the job. This field also indicates if a reboot is required to activate the firmware.
  - **Cancel** – An indication of a canceled job.
- Completed Jobs list:
  - **User** – The user who updated the firmware.
  - **Host** – The IP address of the host to which the adapter or port belongs.
  - **Adapter** – The model of the adapter.
  - **Port** – The port WWN for an FC adapter. This field is blank for a CNA.
  - **Start Time** – The time the job is submitted.
  - **End Time** – The time the job is completed.
  - **Status Message** – The status of the job.

To cancel the jobs with the status "Waiting in queue to start", click **Cancel Updates**.

Chapter 8: DCB Configuration

This chapter describes how to configure DCB parameters on LPe16202/OCe15100 adapters in NIC+FCoE mode.

To view the DCB parameters for a NIC+FCoE port, perform these steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Emulex OneCommand** tab.
2. In the Emulex Device Management tree-view, select the port for which you want to view DCB parameters.
3. Select the **DCB** tab (Figure 57).

Figure 57 DCB Tab (NIC+FCoE)

The following **DCB** tab fields are displayed:

- **DCBX State** – The current DCBX state (enabled or disabled).
- **LLDP Transmit State** – DCBX uses LLDP to exchange parameters between two link peers. For the DCBX protocol to operate correctly, both LLDP receive and transmit states must be enabled. If either the receive or transmit state is disabled, DCBX is disabled.
- **LLDP Receive State** – DCBX uses LLDP to exchange parameters between two link peers. For the DCBX protocol to operate correctly, both LLDP receive and transmit states must be enabled. If either the receive or transmit states is disabled, DCBX is disabled.
- PFC Properties area:
  - **State** – Enabled indicates that there is flow control in both directions (receive and transmit).
  - **Sync'd** – If Yes, the PFC priorities have been set by the peer. This parameter cannot be set.
  - **Error** – Indicates whether an error occurred during the configuration exchange with the peer. Error is also set to Yes when the Compatible method fails.
  - **Active Priorities** – Lists the number of priorities with PFC set to enabled.
- Protocol Properties area:
  - **State** – Enabled indicates that there is flow control in both directions (receive and transmit).
  - **Sync'd** – If Yes, the PFC priorities have been set by the peer. This parameter cannot be set.

- **Error** – Indicates whether an error occurred during the configuration exchange with the peer. Error is also set to Yes when the Compatible method fails.
- **FCoE Priority** – The configured FCoE priority.

- ETS Priority Group Properties area:
  - **State** – The protocol state. It can be enabled or disabled.
  - **Sync'd** – If Yes, the protocol priority has been set by the peer. This parameter cannot be set.
  - **Error** – Indicates whether an error occurred during the configuration exchange with the peer. Error is also set to Yes when the Compatible method fails.
  - Active Groups table:
    - **Group ID** – The Priority Group number. It can be 0 to 7.
    - **Priority Membership** – The priorities that are assigned to each Priority Group. It is represented in comma separated format.
    - **Bandwidth%** – The percentage of available link bandwidth allocated to a particular Priority Group.

The **DCB Configuration** dialog allows you to configure physical port DCB parameters. Although there are eight ETS Priority Groups in which the priorities can be assigned, only two of those groups can be assigned priorities and bandwidths. The FCoE priority is assigned to one Priority Group and the other seven (NIC) priorities must be assigned to one and only one other Priority Group. If you are configuring Priority Groups on a DCBX enabled switch, in order to work properly with the adapter port, you must only configure two Priority Groups in the same way.

To configure DCB for NIC+FCoE adapter ports, perform these steps:

1. From the **DCB** tab, click **Configure DCB**. The **DCB Configuration** dialog appears (Figure 58).

Figure 58 DCB Configuration Dialog (NIC+FCoE)

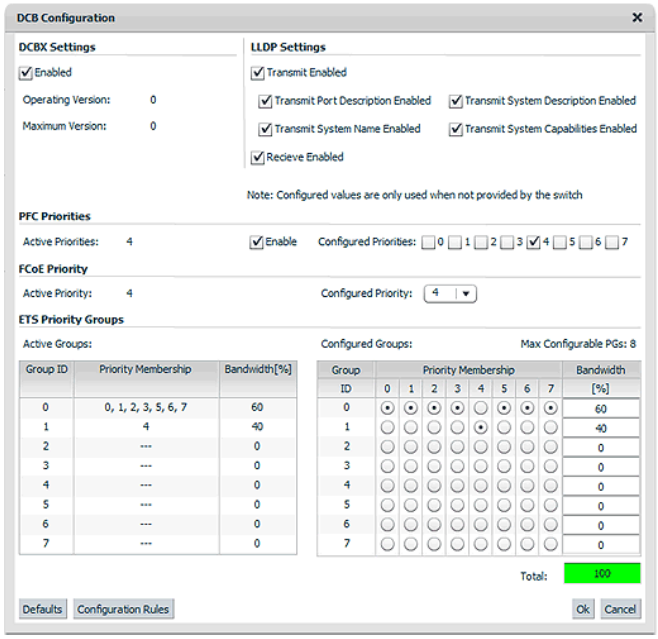

2. Configure the settings. The following **Configure DCB** dialog fields are displayed:

- DCBX Settings area:
  - **Enabled** – DCBX can be enabled or disabled. With DCBX enabled, the configured values are used only if the switch does not provide them. With DCBX disabled, the configured values are used.
  - **Operating Version** – The operating version of the DCBX protocol. The system adjusts as needed to operate at the highest version supported by both link partners. This setting cannot be changed.

- **Maximum Version** – The highest DCBX protocol version supported by the system. Version numbers start at 0. The DCBX protocol must be backward-compatible with all previous versions. This setting cannot be changed.

- LLDP Settings area:
  - **Transmit Enabled** – LLDP Transmit can be enabled or disabled.
  - **Transmit Port Description Enabled** – Provides a description of the port in an alpha-numeric format. If the LAN device supports RFC 2863, the value equals the ifDescr object.
  - **Transmit System Description Enabled** – Provides a description of the network entity in an alpha-numeric format. This description includes the system's name and versions of hardware, operating system, and networking software supported by the device. If the LAN device supports RFC 3418, the value equals the sysDescr object.
  - **Transmit System Name Enabled** – Provides the system's assigned name in an alpha-numeric format. If the LAN device supports RFC 3418, the value equals the sysName object.
  - **Transmit System Capabilities Enabled** – Indicates the primary functions of the device and whether these functions are enabled on the device. The capabilities are indicated by two octets. Bits 0 through 7 indicate Other, Repeater, Bridge, WLAN AP, Router, Telephone, DOCSIS cable device, and Station, respectively. Bits 8 through 15 are reserved.
  - **Receive Enabled** – LLDP Receive can be enabled or disabled.
- PFC Priorities area:
  - **Active Priorities** – The priorities that are marked active for PFC.
  - **Enable** – Enables PFC priorities.
    - To configure a NIC port, clear this check box.
    - To configure an FCoE port, select this check box.
  - **Configured Priorities** – If the Enable check box is selected, these are the priorities that are configured but that might not yet be active. You can select two PFC priority check boxes, one of which must match the FCoE priority. The other PFC priority is for Ethernet traffic. The PFC priority for Ethernet traffic must be assigned to a Priority Group that has no other priorities.
  - FCoE Priority area:
    - **Active Priority** – The active FCoE priority.
    - **Configured Priority** – The configured FCoE priority.
- ETS Priority Groups area:
  - Active Groups – This area displays the settings that are currently active:
    - **Max Configurable PGs** – The maximum number of Priority Groups that can be configured.
    - **Group ID** – The active Priority Group ID.
    - **Priority Membership** – The priorities that are currently assigned to the various Priority Groups.
    - **Bandwidth** – The bandwidths that are currently assigned to various Priority Groups.
  - Configured Groups – This area allows you to create a new configuration:
    - **Group ID** – The configured Priority Group ID.
    - **Priority Membership** – The configured priority membership grouping.
    - **Bandwidth%** – The configured bandwidth value for the various Priority Groups.

      NOTE An error message is displayed if you try to configure more priority groups than the adapter supports. The **Max Configurable PGs** field in the ETS Priority Groups area shows the number of priority groups supported by the adapter.

3. Click **OK** to apply and save your changes.

To see the configuration rules that apply to configuring priorities, click **Configuration Rules**. You must observe these rules when configuring Priority Groups for adapter ports. Configuration rules are displayed based on the type of adapter ports you are configuring.

- FCoE configuration rules:
  - One and only one priority is configured for the priority.
  - You can select a maximum of two PFC priorities, and one of them must match the priority.
  - The PFC Priority Group to which the priority is assigned must contain no other priorities.
  - The additional PFC priority must be assigned to a Priority Group that has no other priority.
  - Bandwidths of all the Priority Groups must total percent.
- NIC configuration rules:
  - One and only one priority is configured for the priority.
  - The PFC priority must be assigned to a Priority Group that has no other priorities.
  - Bandwidths of all the Priority Groups must total 100 percent.

To reset the DCB configuration to its default values, click **Defaults**.

Chapter 9: Exporting SAN Information in Cluster View

Exporting SAN information (creating a SAN report) processes in the background until all the jobs are completed. During this period you can still browse through the other tabs.

NOTE Creating a SAN report can take several minutes for a large SAN.

To export SAN information, perform these steps:

1. Select a cluster in the vSphere console tree-view, and if applicable, select the **Emulex OneCommand** tab.
2. From the Emulex Device Management tree-view, select **Export San Info**. Export SAN Information is displayed (Figure 59).

Figure 59 Export SAN Info Option in the Emulex Device Managment Tree-View

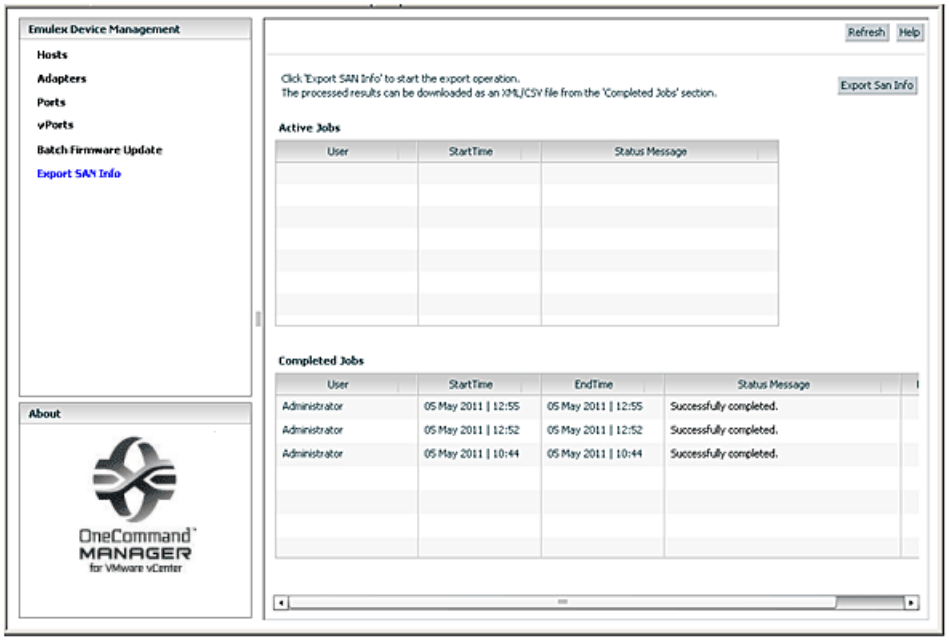

The following **Export SAN Info** window fields are displayed:

— Active Jobs column:
- **User** – The user who updated the firmware.
- **Start Time** – The time the job is submitted.
- **Status Message** – The status of the job.

— Completed Jobs column:
- **User** – The user who updated the firmware.
- **Start Time** – The time the job is submitted.
- **End Time** – The time the job is completed.
- **Status Message** – The status of the job.
- **Download** – Click **XML** or **CSV** to download the file with SAN information in the specified format.

3. Click **Export San Info**. The **Export SAN Info** dialog is displayed (Figure 60).

Figure 60 Export SAN Info Dialog

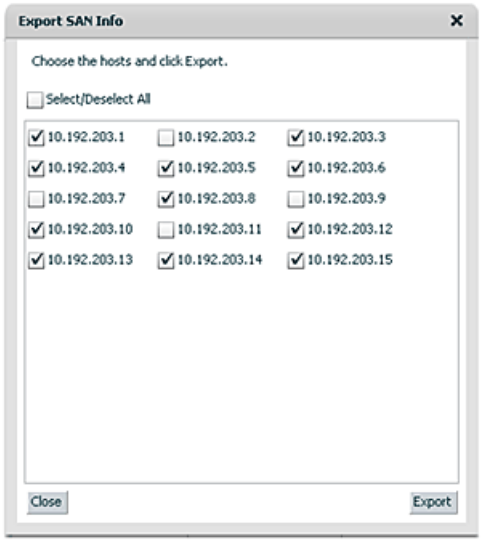

4. Select host or hosts to export. Optionally, select the **Select/Deselect All** check box to select all hosts. Clear the **Select/Deselect All** check box to clear all host check boxes.
5. Click **Export**.

   The **Export SAN Info Jobs** window is displayed. The status is displayed in the Active Jobs list (Figure 61).

Figure 61 Export SAN Info Jobs Window, Active Jobs List

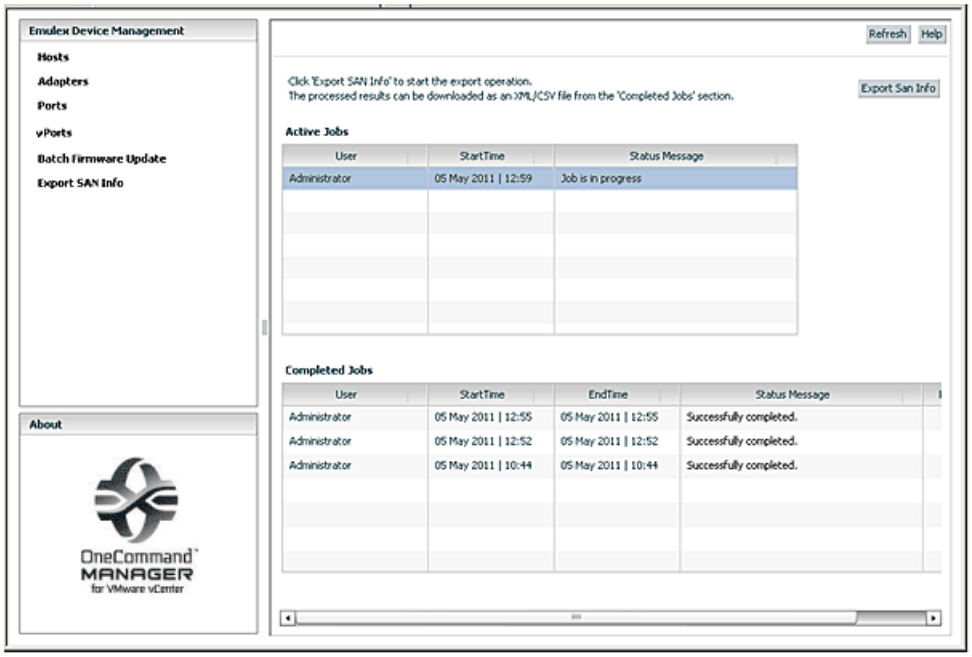

6. When the export job is completed, the entry is displayed in the Completed Jobs list with a **Successfully completed** status.

9.1 Capturing SAN Information in XML or CSV Format

1. Click the row that represents the information you want to capture.
2. Scroll to the right until the Download column appears.
3. In the Download column, click either **XML** or **CSV** to capture the information.

9.2 Considerations When Exporting SAN Information in a Cluster View

- At any time, only 10 completed jobs are available to be exported. If more than 10 jobs are completed, the first completed jobs are not available for download.
- If you click **Export SAN Info** and log out from the vSphere console before the export is completed, all pending jobs fail.

Chapter 10: Diagnostics

This section describes diagnostics for Emulex LightPulse adapters.

10.1 Emulex LightPulse Diagnostics

When running port diagnostic tests using OneCommand Manager for VMware vCenter, do not perform any active management operations on the ESXi host.

CAUTION Running a PCI Loopback, Internal Loopback, External Loopback, or POST test on an LPe12000-series adapter that is being used to boot from SAN is not advisable. After the tests have completed, the system performs an adapter reset, which may cause a loss of connectivity to the SAN and possible loss of data. To perform these tests, you must make sure that the adapter is not currently being used to boot from SAN.

Do one of the following:

- Move the target adapter to a non-boot from SAN host, and perform the tests from that location.
- If the host with the target adapter is also hosting other boot from SAN adapters, perform a boot from SAN using one of the other boot from SAN adapters. The target adapter can now be tested, because it is no longer being used for boot from SAN.

10.1.1 Running Loopback Tests

Yuo can run three loopback tests for the FC adapter port:

- PCI Loopback – A firmware-controlled diagnostic test in which a random data pattern is routed through the PCI bus without being sent to an adapter link port. The returned data is subsequently validated for integrity.
- Internal Loopback – A diagnostic test in which a random data pattern is sent down to an adapter link port, and then is immediately returned without actually going out on the port. The returned data is subsequently validated for integrity.
- External Loopback – A diagnostic test in which a random data pattern is sent down to an adapter link port. The data goes out of the port and immediately returns using a loopback connector. The returned data is subsequently validated for integrity.

To run loopback tests, perform these steps:

1. From the Emulex Device Management tree-view, select the FC adapter physical port on which you want to run the loopback test.
2. Select the **Diagnostics** tab (Figure 62). In the Loopback Test area of the dialog, choose the type of loopback test you want to run, and define the loopback test parameters.

Figure 62 Diagnostics Tab

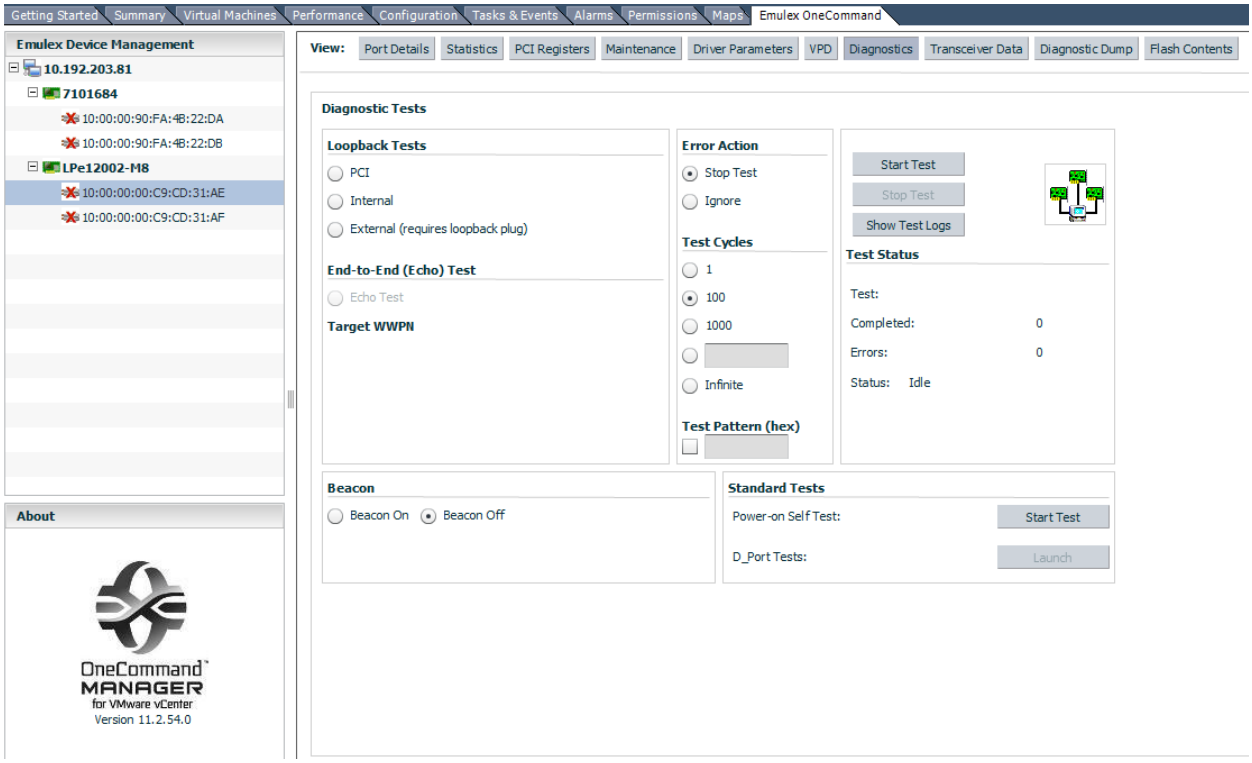

3. Click **Start Test**. The following warning is displayed.

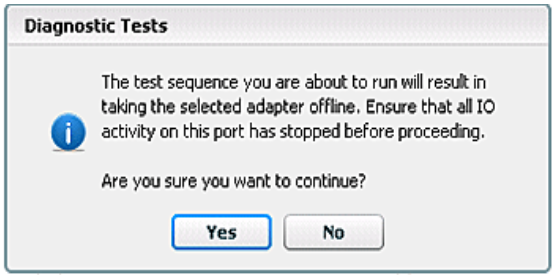

4. Click **Yes**. A progress bar shows that the test is running.
5. Periodic test feedback, consisting of the current loopback test/cycle plus the completion status of each type of test, is displayed in the Test Status section of the dialog. Click **Show Test Logs** to view and save the log file.

10.1.2 Running End to End (ECHO) Tests

Run echo tests using the End-to-End (ECHO) Test section of the **Diagnostics** tab. The end-to-end test enables you send an ECHO command/response sequence between an adapter port and a target port.

To run end-to-end echo tests, perform these steps:

1. From the Emulex Device Management tree-view, select the physical port on which you want to initiate the End-to-End (ECHO) test.
2. Select the **Diagnostics** tab (Figure 62). In the End-to-End (Echo) Test area, select **Echo Test**.

3. Enter the WWPN for the target. The following warning appears:

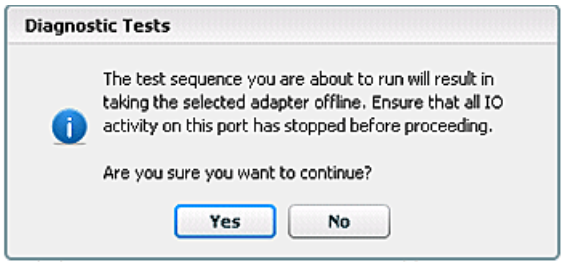

4. Click **Yes**. A result screen appears and the test results appear in the Test Log.
5. Either click **Clear** to erase the contents of the log display, or click **Save to File** to save the log file.

NOTE The **ECHO Test** button is enabled only if its port has targets connected.

10.1.3 Running D_Port Tests

D_Port is a diagnostic mode supported by Brocade switches for LPe16000-series, LPe31000-series, and LPe32000-series adapters with D_Port support.

NOTE D_Port is also referred to as ClearLink.

D_Port tests detect physical cabling issues that can result in increased error rates and intermittent behavior. When activated, D_Port tests include:

- Local electrical loopback
- Loopback to the remote optics
- Loopback from the remote port to the local optics
- A full device loopback test with data integrity checks
- An estimate of cable length (to validate that a proper buffering scheme is in place)

These tests allow a level of fault isolation to distinguish faults due to marginal cable, optics modules, and connector or optics seating.

To run D_Port tests, perform these steps:

1. From the Emulex Device Management tree-view, select the FC adapter physical port for which you want to run the D_Port tests.
2. Select the **Diagnostics** tab, and in the Standard Tests area, click **Launch (**Figure 63).

Figure 63 FC Adapter Diagnostics Tab with D_Port Tests Option

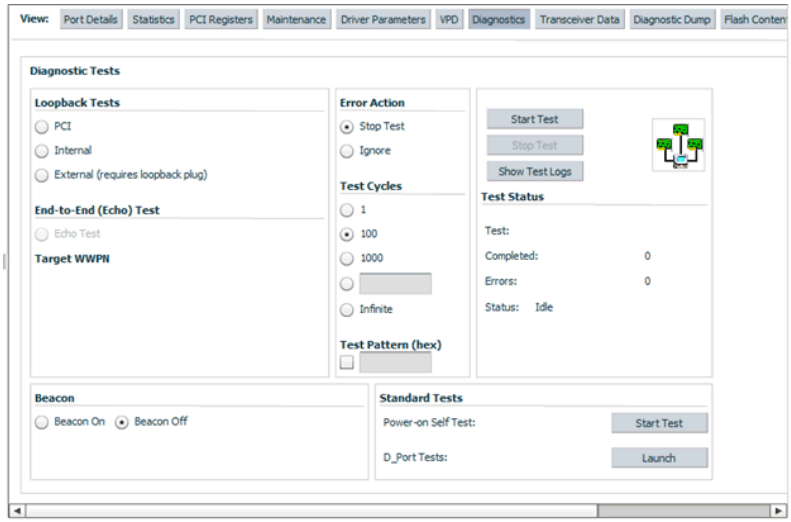

3. The D_Port Tests dialog is displayed (Figure 64). Click **Start Test**.

Figure 64 D_Port Tests Dialog

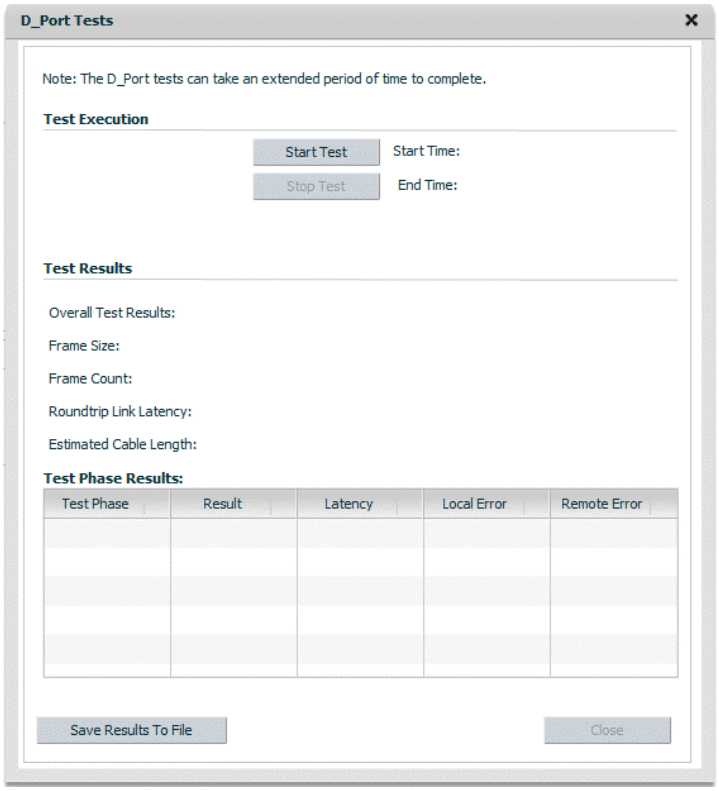

The D_Port tests are launched. If all tests pass, a dialog similar to Figure 65 is displayed. If all tests do not pass, the Failed result in shown in the Test Phase Results (Figure 66).

Click **Save Results to File** to save the test results to a text file. You can view this text file in any text editor.

To stop tests, click **Stop Test**. If a test phase fails, the D_Port diagnostics are automatically stopped. In this case, some of the phases may not be reported in the results. However, the failed phase will be reported.

If the Overall Test Result is FAILED, you must either rerun the tests successfully, or reset the HBA port to bring the link back up.

NOTES

- The D_Port tests can take an extended period of time to complete.
- If an older SFP version is detected by the OneCommand Manager for VMware vCenter, a message is displayed under the results box indicating that the SFP version does not fully support all D_Port tests.

Figure 65 D_Port Tests Dialog – Passed Result

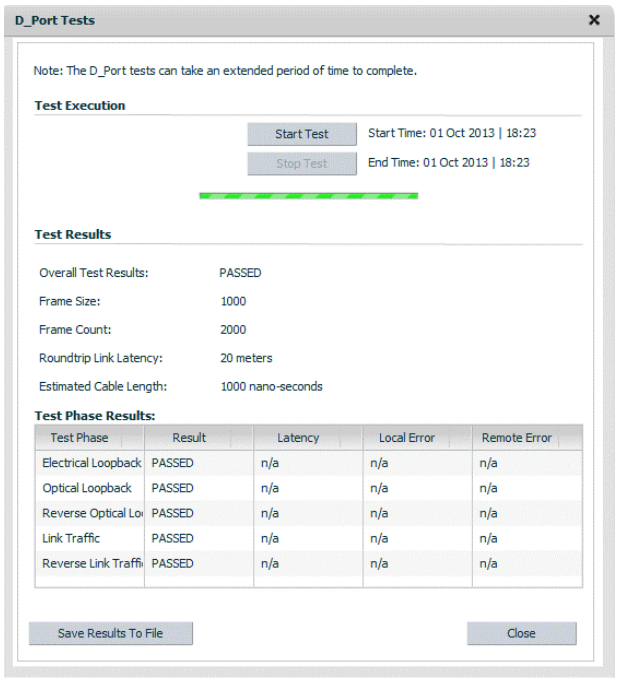

Figure 66 D_Port Tests Dialog – Failed Results

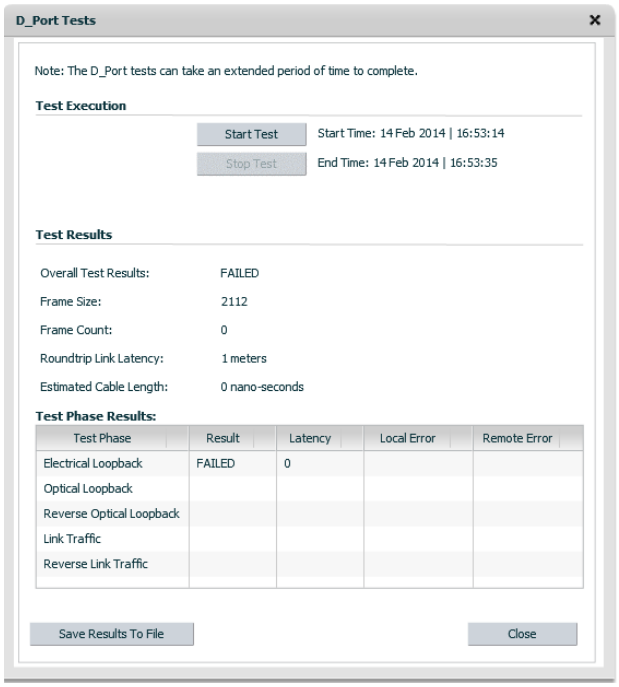

10.1.4 Running a POST

The power-on self-test (POST) is a firmware test normally performed on an adapter after a reset or restart. The POST does not require any configuration to run.

NOTE The POST test is available only for LPe12000-series adapters.

To run the POST, perform these steps:

1. From the Emulex Device Management tree-view, select the FC adapter physical port on which you want to run the POST.
2. Select the **Diagnostics** tab (Figure 62) and, in the Standard Tests area, click **Start Test**. A progress window appears, showing the progress of the POST test.
3. After the test is completed, the **Test Completion Status** window appears. Click **OK**. A POST window is displayed with the POST information.

10.1.5 Using Beaconing

Beaconing enables you to force a specific adapter's LEDs to blink in a particular sequence. The blinking pattern acts as a beacon, making it easier to locate a specific adapter among racks of other adapters.

NOTE Beaconing is disabled if the selected adapter does not support beaconing.

To enable or disable beaconing, perform these steps:

1. From the Emulex Device Management tree-view, select the adapter port whose LEDs you want to set.
2. Select the **Diagnostics** tab (Figure 62) and click either **Beacon On** or **Beacon Off**.

10.1.6 Setting Up Diagnostic Test Options

Setting up test options include error actions, test cycle counts, and test patterns.

10.1.6.1 Setting Up a Test Failure Error Action

Two error action options are available in the event of a test failure:

- Stop Test – Does not log the error and aborts the test. No further tests are run.
- Ignore – Logs the error and proceeds with the next test cycle.

10.1.6.2 Setting Up Test Cycles

Specify one of the following test cycles:

- Select an established cycle count by clicking on the corresponding radio button.
- Enter a custom cycle count in the blank field in the Test Cycles area.
- Select **Infinite** to set the test to run until you manually click **Stop Test**.

10.1.6.3 Setting Up a Test Pattern

Enter a custom test pattern to be used in tests that transfer data. The test pattern can be up to 8 hexadecimal bytes.

10.1.6.4 Test Status

The Test Status area shows how many completed cycles of each test ran, as well as the number of errors.

10.1.7 Saving the Log File

You can save the test log to a log file for later viewing or printing. When data is written to a saved file, the data is appended at the end of the file. Each entry has a two-line header with the adapter identifier and the date and time of the test. The data accumulates to form a chronological history of the diagnostics performed on the adapter.

- The default location is the OneCommand Manager for VMware vCenter install directory on your local drive.
- In the VMware Server, there is no default directory for ESXi.

After writing an entry into the log, you are prompted to clear the display. Figure 67 shows the Diagnostic Test log entries that will be saved to the log file.

Figure 67 Diagnostic Test Log Entries

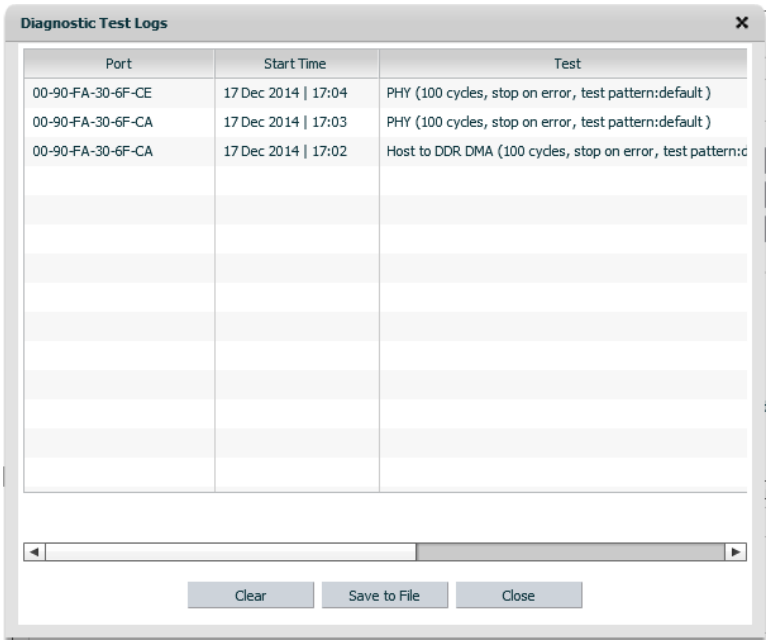

To save the log file, perform these steps:

1. After running a test from the **Diagnostic** tab, click **Show Test Logs**. The **Diagnostic Test Logs** dialog appears. The default name of a saved file is `DiagTest.log`.
2. Browse to the desired directory, change the log file name if you want, and click **Save to File**.

10.1.8 Creating Port Diagnostic Dumps

Diagnostic dump enables you to create and manage a diagnostic dump for a selected port. Dump files contain information, such as, firmware version and driver version, that is particularly useful when troubleshooting an adapter.

To start a diagnostic dump, perform these steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Emulex OneCommand** tab.
2. In the Emulex Device Management tree-view, select the port.
3. Select the **Diagnostic Dump** tab (Figure 68). Diagnostic dump information is displayed.

Figure 68 Port Diagnostic Dump Tab

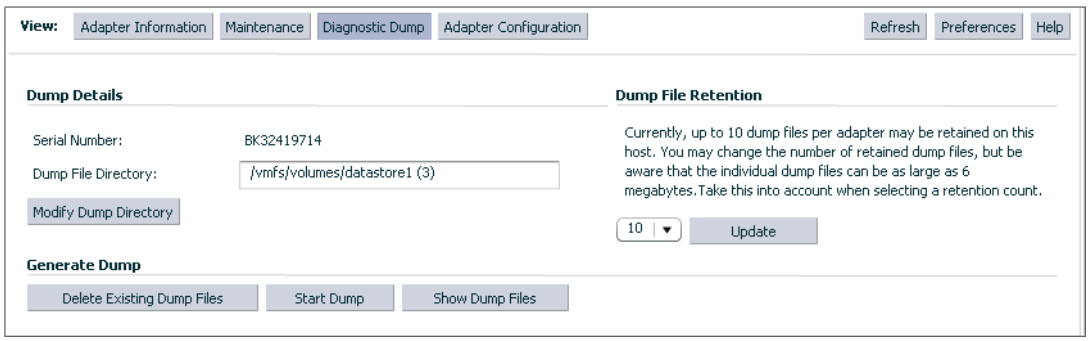

4. Enter a name in the **Dump File Directory** field in the Dump Details area to set the dump file directory. A prefix of `/vmfs/volumes` is added to the location, if it is not specified.
5. To specify up to 20 files to retain using the **Dump Files Retention** counter, enter the number of files and click **Update**.
6. Click **Start Dump** to initiate a diagnostic dump on the selected port. Click **Delete Existing Dump Files** to remove existing dump files for the selected port.

CAUTION Disruption of service can occur if a diagnostic dump is run during I/O activity.

10.1.8.1 Diagnostic Dump Tab Field Definitions

- **Port WWN** – The port WWN.
- **Dump File Directory** – The location where the dump files are created.
- **Dump File Retention** – The number of dump files per port to retain on the host.

10.1.9 Viewing Diagnostic Dump Files

You can view diagnostic dump file names using the OneCommand Manager for VMware vCenter. The dump files are stored on the host's data store, and the vSphere client can be used to download dump files by browsing the host data store.

To view the FC port diagnostic dump, perform these steps:

1. On the **Diagnostic Dump** tab, click **Show Dump Files**. The **Diagnostic Dump Files** window opens showing the diagnostic dump files currently on your system (Figure 69). These files are available in the dump directory configured from the **Diagnostic Dumps** tab. You can extract these files using the vSphere client.

Figure 69 Diagnostic Dump Files Window

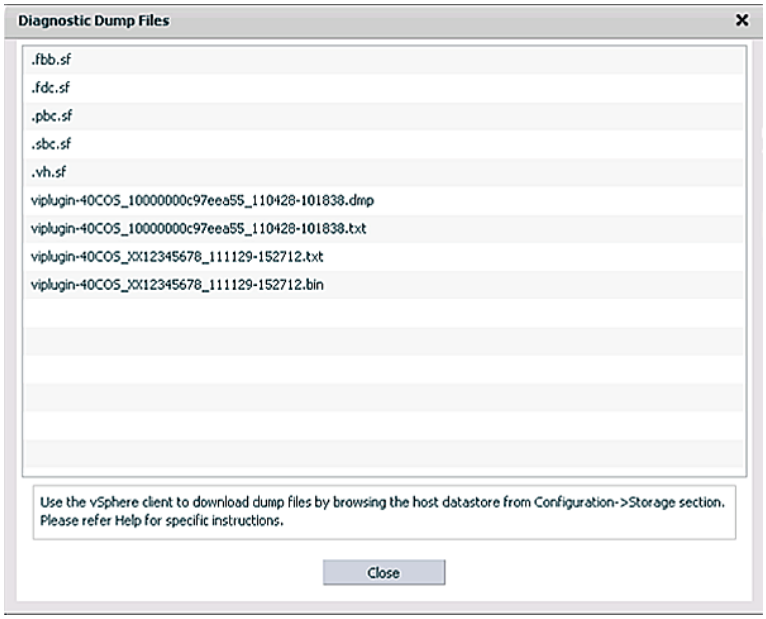

2. Extract the dump files by using the vSphere client to download the dump files by browsing the data store.
   a. Click the **Configuration** tab in the vSphere client. The Datastores view is displayed (Figure 70).

Figure 70 Datastores View of vSphere Client

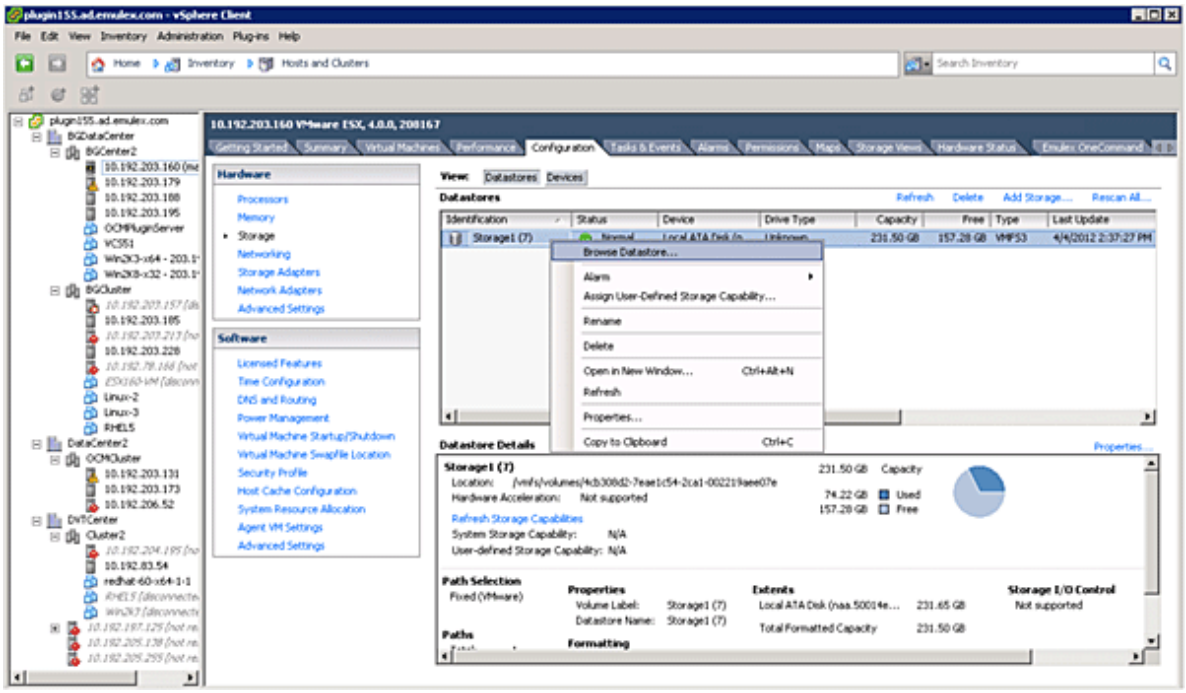

b. Right-click the datastore where the dump file is located, and select **Browse Datastore** from the context menu. A Datastore Browser window opens (Figure 71).

Figure 71 vSphere Datastore Browser

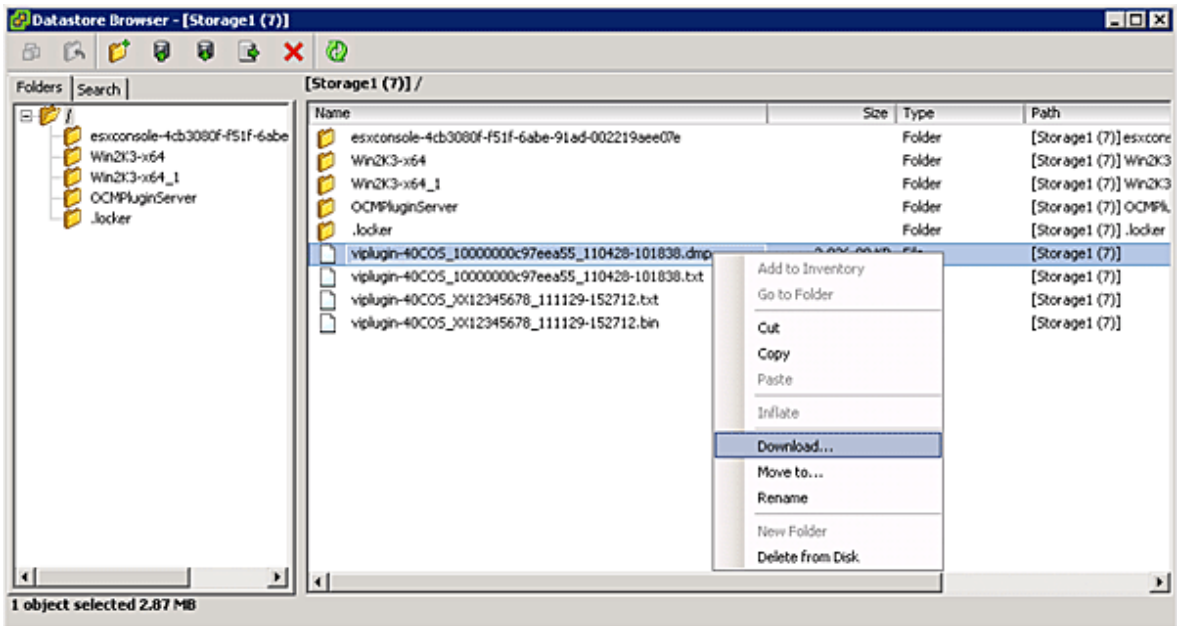

c. In the Datastore Browser, right-click the dump file that you want to review and select **Download...** from the context menu. A **Browse for Folder** window opens.

d. Select the desired location for the dump files in the **Browse for Folder** window and press **OK**. The file is downloaded to the location you select.

You can view the dump file in any text editor.

3. Click **Yes**. A progress bar shows that the test is running.

   Periodic test feedback, consisting of the current loopback test/cycle plus the completion status of each type of test, is displayed in the Test Status section of the dialog. Click **Show Test Logs** to view and save the log file.

Chapter 11: Generating and Installing Secured Certificates

OCMNG is a web application, based on a client-server model, running on the Apache Tomcat Web Server. Data is exchanged between the client (browser) and the server (on a remote machine), which requires a secure user logon to manage Emulex adapters on different and multiple hosts.

11.1 SSL Certificate

A Secure Sockets Layer (SSL) certificate establishes an encrypted connection between the web server and the web browser on a remote machine. This connection allows private information to be transmitted without eavesdropping, data tampering, or message forgery.

An SSL certificate provides security through encryption and authentication. Encryption is ensured by accessing the remote server using the HTTPS protocol and an SSL certificate.

NOTE If OneCommand Manager for VMware vCenter is running, the server must be configured to support HTTPS protocol access and provide a self-signed certificate.

The OneCommand Manager for VMware vCenter server is authenticated to the browser by a Public Key in the self-signed certificate.

11.1.1 Generating an SSL Certificate

To allow secured communication between the client and server, perform these steps:

1. Generate a self-signed certificate with a keystore file for each server providing the server's domain name and company details. See Section 11.1.2, Generating a Self-Signed Certificate, for instructions. For more information, refer to the X.509 attributes list on the International Telecommunications Union website.
2. Use this certificate to create a request to the customer's trusted certificate authority (CA). The request certificate is referred as a Certificate Signing Request (CSR). The CA issues a new SSL certificate. See Section 11.2.1, Generating a CSR for a Server Using the Java Tool, for instructions.
3. Import the new SSL certificate to the application server, and install the SSL certificate on the client's browser. See Section 11.2.4.1, Installing the Certificates to the Keystore of OneCommand Manager for VMware vCenter, for instructions.
4. Configure the server to use the keystore file. See Section 11.2.4.2, Configuring a Web Server, for instructions.
5. Access the server's content through the browser using HTTPS protocol.

   The browsers understand the certificate, and the browsers allow access to and from the remote server.

11.1.2 Generating a Self-Signed Certificate

A self-signed certificate is a certificate that is signed by itself (the server hosting OneCommand Manager for VMware vCenter) rather than a trusted CA. This self-signed certificate includes a public or private key that is distributed by the SSL to verify the identity of the server.

A self-signed certificate can also be used as an alternative to SSL certificates if the server is not running in a public domain.

If a self-signed certificate is used in place of an SSL, a warning is shown in the browser before accessing the server content.

For Java-based applications, a self-signed certificate can be generated using the tools provided by Java. This creates a keystore file that must be installed on the web server. This keystore includes a private key specific to the server used for generating a CSR and authenticating the server.

As the OneCommand Manager for VMware vCenter server is developed using Java, it leverages the keystore tool provided by Java to generate the self-signed certificates at no cost.

NOTE The self-signed certificate for the OneCommand Manager for VMware vCenter server is generated and installed on its server as part of the OneCommand Manager for VMware server installation on a Windows machine. This self-signed certificate is generated with Broadcom® organization details using RSA algorithm and private key of size 2048 bits.

To generate a self-signed certificate, perform these steps:

1. In the OCM for VMware installation directory, go to `ApacheTomcat\conf`.

   ```
   >>cd /d  "C:\Program Files\Emulex\OCM for VMware\ApacheTomcat\conf"
   ```

2. Run the following command:

   ```
   >> ..\..\JRE\bin\keytool.exe -genkey -alias <new-alias> -keyalg RSA
   -keystore emulex.vcplugin.jks  -keysize 2048
   ```

   NOTE You can change alias, keysize, and keystore name.

Example

```
Enter keystore password: (Enter "emulex" if using the same keystore name)
Re-enter new password:
What is your first and last name?
  [Unknown]:  pluginserver.ad.emulex.com  (Give the complete domain name of the
server [FQDN])
What is the name of your organizational unit?
  [Unknown]:  ocm
What is the name of your organization?
  [Unknown]:  elx
What is the name of your City or Locality?
  [Unknown]:  bg
What is the name of your State or Province?
  [Unknown]:  ka
What is the two-letter country code for this unit?
  [Unknown]:  in
Is CN=pluginserver.ad.emulex.com, OU=ocm, O=elx, L=bg, ST=ka, C=in correct?
  [no]:  yes

Enter key password for <elxocm>:
        (RETURN if same as keystore password)
```

11.2 Generating a CSR

A Certificate Signing Request (CSR) is a block of encrypted text that is generated on the server on which the certificate is used. A CSR contains information to be included in the SSL certificate, such as the organization name, common name (domain name), locality, country, and other X.509 attributes. It also contains the public key that is included in the certificate. The CA uses the CSR to create a new SSL certificate.

11.2.1 Generating a CSR for a Server Using the Java Tool

To generate a CSR for a server, use the Java tool available in the `jre/bin` folder. The syntax using the Java tool follows:

```
keytool -certreq -keyalg <algorithm>  -alias <alias-name>  -file <csr-name>
-keystore <keystore-name>
```

Example

```
keytool -certreq -keyalg RSA -alias selfsigned -file elxocmreq.csr -keystore
emulex.vcplugin.jks
```

11.2.2 Generating and Validating a CSR

To generate a CSR, perform these steps:

1. Generate a self-signed certificate (see Section 11.1.2, Generating a Self-Signed Certificate, for instructions).
2. Generate a CSR using the following syntax:

```
>>..\..\JRE\bin\keytool -certreq -v -alias <new-alias> -file elxocmreq.csr
-keypass elxocm -keystore emulex.vcplugin.jks
Enter the keystore password: (Enter "emulex" if using the default keystore
name)
Certification request stored in file <elxocmreq.csr>
```

To validate a CSR for its completeness, perform these steps:

You can validate the generated CSR for its completeness before submitting (with the help of the CA). Copy the CSR content from the following link for validation.

```
http://www.sslshopper.com/csr-decoder.html
```

NOTE The CSR must begin and end with the following tags:

```
-----BEGIN NEW CERTIFICATE REQUEST-----
-----END NEW CERTIFICATE REQUEST-----
```

Example

```
-----BEGIN NEW CERTIFICATE REQUEST-----
MIIC2TCCAcECAQAwZDELMAkGA1UEBhMCaW4xCzAJBgNVBAgTAmthMQswCQYDVQQHEwJiZzEMMAoG
A1UEChMDZWx4MQwwCgYDVQQLEwNvY20xHzAdBgNVBAMTFmJnc3N5ZWQxLmFkLmVtdWxleC5jb20w
ggEiMA0GCSqGSIb3DQEBAQUAA4IBDwAwggEKAoIBAQCDhovljXhfjNPM/5eBsX4280AIl3YARn0p
R6Z7eOqs1r5Qh07kT58M6T8fER+NpIN7WhOOF/TbsFsS0gmfYwJQqttvtq1dtxUGpvFe9lywbP+l
kY+w6GOyPTG2qnXgILtX5ArZbC2UBbz+J8WJ3SjPHXiSY35EZWnyZZmIN8v1vOe9e21f8vwRkn/4
fdfFrpoQa3H+GcAJMRSBRTd5H6mXQv6HaA5Z0BbsABisFx4scqSuM/HJKLP6GcSHR61bzHfiO/NH
4qU/s6I2LC5DvGs1hIW3PPbmb1rxBiEFpjPtWhfzPxPMKSU8uey+lE0UIPMS0FMTxo63oYnMeiSX
X5mxAgMBAAGgMDAuBgkqhkiG9w0BCQ4xITAfMB0GA1UdDgQWBBSvpKLBF3lY03Jin9kI4ym94bJi
zjANBgkqhkiG9w0BAQsFAAOCAQEAfs94wzEUlDAMq0jITi6fiD7YxK2KFWJgMBfjxZIGex2zxlHL
mOS14BGSWk5dvSMwqDBC14l4C79rUOlTUwwWs9zFqMHynndQ2Ze2vuJNTWUlnFyFb37/rEvbFufB
QVvFXgycaKRgUpWo2x5sekRJRAPxXI/vLWOFRLLrzcVykgZ/sg3QrO4ezlKFc49put0vKpvI1dY9
l9BN2REuWrlmq5y3L8nx8mKX9dRmP6CKzHBaVrvY+nVju+Vf/ikfTtQIDEXAIW2Q7AObpcOaudnf
Nsaey+u27vGy77gAv7092xBHsDyOrD7COy/83b194igmVBVY4dt0496oXkOHCA0txA==
-----END NEW CERTIFICATE REQUEST----
```

11.2.3 Getting an SSL Certificate

The CSR can be submitted to the trusted CA (as chosen by you). The CA validates the CSR and issues a new SSL certificate.

11.2.4 Installing the SSL into the Web Server

When you receive the SSL certificate from the CA, you must install the SSL certificate on the server to accept the secure connections.

NOTE The CSR must be generated on the same machine that the server is running on. The SSL certificate must also be installed on this same server.

11.2.4.1 Installing the Certificates to the Keystore of OneCommand Manager for VMware vCenter

The Root Certificate file, the Intermediate Certificate file, and the Primary Certificate file must all be installed in the keystore.

To install the certificates to the keystore of OneCommand Manager for VMware vCenter, perform these steps:

1. Download the SSL certificate file from the CA. Save the SSL certificate file to the same directory as the keystore (self-signed certificate) that was created for the CSR.

   NOTE The certificate only works with the same keystore that was initially created for the CSR. The certificates must be installed to your keystore in the correct order.

2. Install the Root Certificate file.

   Every time you install a certificate to the keystore, you must enter the keystore password that you chose when you generated it. Enter the following command to install the Root Certificate file:

   ```
   keytool -import -trustcacerts -alias root -file RootCertFileName.crt
   -keystore keystore.key
   ```

   If the following message is displayed, select `Yes`:

   ```
   Certificate already exists in system-wide CA keystore under alias <...> Do
   you still want to add it to your own keystore?
   ```

   If successful, the following message is displayed:

   ```
   Certificate was added to keystore.
   ```

3. Install the Intermediate Certificate file.

   If the CA provided an Intermediate Certificate file, you must install it here using the following command:

   ```
   keytool -import -trustcacerts -alias intermediate -file
   IntermediateCertFileName.crt -keystore keystore.key
   ```

   If successful, the following message is displayed:

   ```
   Certificate was added to keystore.
   ```

4. Install the Primary Certificate file.

   Use following command to install the Primary Certificate file (for your domain name):

   ```
   keytool -import -trustcacerts -alias tomcat -file PrimaryCertFileName.crt
   -keystore keystore.key
   ```

   If successful, the following message is displayed:

   ```
   Certificate reply was installed in keystore.
   ```

All the certificates are now installed to the keystore file. You must configure your server to use the keystore file.

11.2.4.2 Configuring a Web Server

NOTE These configuration changes are not required if the default keystore name and password are used. If they are different, you must change the configuration as needed.

1. Copy the keystore file or SSC to a directory (preferably, the `conf` folder) of the web server.
2. Open the file `${CATALINA_HOME}/conf/server.xml` in a text editor.
3. Uncomment the SSL Connector Configuration.
4. Make sure the keystorePass matches the password for the keystore, and the keystoreFile contains the path and file name of the keystore.

   Your connector should be displayed similar to the following:

   ```
   <Connector className="org.apache.catalina.connector.http.HttpConnector"
   port="8443" minProcessors="5" maxProcessors="75" enableLookups="true"
   acceptCount="10" debug="0" scheme="https" secure="true">

   <Factory className="org.apache.catalina.net.SSLServerSocketFactory"
   clientAuth="false" protocol="TLS" keystoreFile="./conf/emulex.vcplugin.jks"
   keystorePass="emulex"/>
   ```
5. Save the changes to `server.xml`.
6. Restart the web server.

   If you launch the OneCommand Manager for VMware vCenter URL in the browser, the application should be launched without any security warnings.

NOTE Use the host name with the domain name that you used to generate the CSR.

Chapter 12: Troubleshooting

This section includes information about certificate or insecure-content warnings that may be displayed on the console. This section also describes unexpected circumstances and some proposed solutions.

12.1 Security

The vSphere Web Client server and the OneCommand Manager for VMware vCenter can be installed on different machines. As a result, certificate or insecure-content warnings could occur. The two ways to remedy the issue are:

- Accept the blocked content – temporary solution
- Install a security certificate – permanent solution

12.1.1 Accepting the Blocked Content

The procedure for accepting blocked content depends on the type of browser you are using. This solution is not permanent, and you must repeat this procedure every time you load the web client. If you want a permanent solution, you must install the correct security certificate. See Section 12.1.2, Installing a Security Certificate.

12.1.1.1 Internet Explorer

Accept the blocked content (Figure 72).

Figure 72 Blocked Content in Internet Explorer

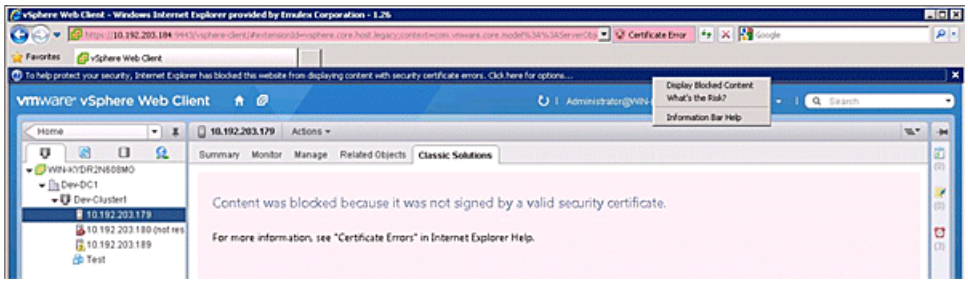

12.1.1.2 Chrome and Firefox

1. Load the plug-in URL in a separate tab or window.

   The plug-in URL format is:

   `https://<plugin-server>:<https-port>/elxvcplugin`

   For example:

   `https://pluginserverhostFQDN:443/elxvcplugin`

   NOTE You can extract the plug-in server, IP address, host name, and port number from the browser warning message.

2. Confirm or accept the certificate warning (Figure 73).
3. Refresh the vSphere web client tab or window.

Figure 73 Blocked Content in the Firefox or Chrome Browser

12.1.2 Installing a Security Certificate

A permanent solution to the security warnings is to install the correct security certificate.

To install a security certificate, perform these steps:

1. Open Internet Explorer in Administrative mode.
2. Load the plug-in URL, and accept the certificate warning.
   The plug-in URL must have the following format:
   `https://<plugin-server>:<https-port>/elxvcplugin`
   The page loads with a certificate error.
3. Click the **Certificate Error**. A list appears showing the untrusted certificate (Figure 74).

Figure 74 Untrusted Certificate

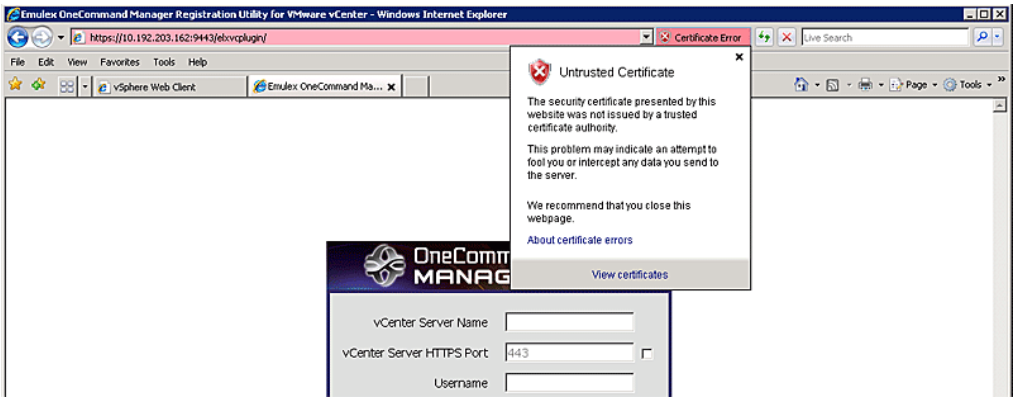

4. Click **View certificates** in the Certificate Error list. The **Certificate** dialog is displayed (Figure 75).

Figure 75 Certificate Dialog

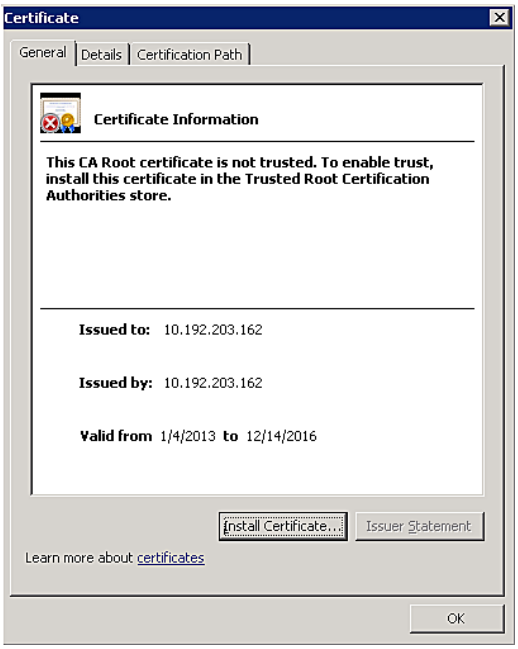

5. Click **Install Certificate...**. The **Certificate Import Wizard** is displayed (Figure 76).

Figure 76 Certificate Import Wizard

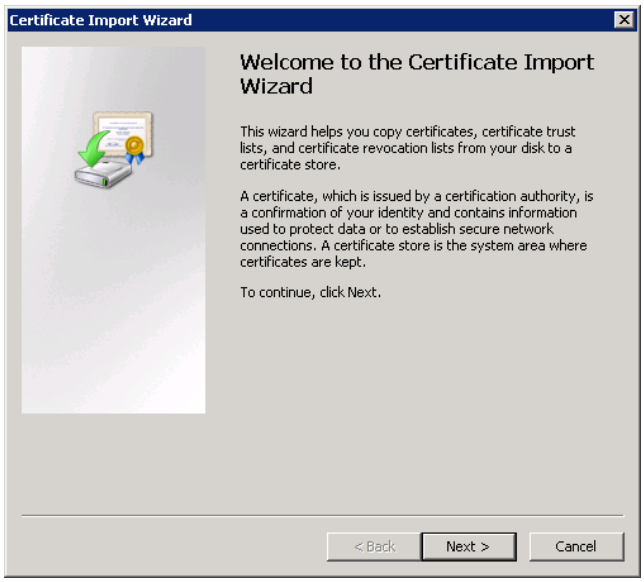

6. Follow the wizard instructions and install the certificate to the Trusted Root Certification Authorities location.

12.1.3 Issues and Resolutions

Your system may operate in an unexpected manner in several circumstances.. Table 2 explains some of these circumstances and offers one or more solutions for each issue.

Table 2 Troubleshooting Issues and Resolutions

Issue	Resolution
The **Emulex OneCommand** tab is not visible in the vSphere console.	In the vSphere console, select the **Plug-in** menu and choose **Manage Plug-ins**. In the **Plug-in Manager** window, check the status of the Emulex OneCommand Manager for VMware vCenter (Emulex OneCommand). The status must be Enabled. If it is not, enable it. On the machine where OneCommand Manager for VMware vCenter is installed, make sure the port numbers configured during the installation are open and dedicated to the plug-in server only. No other service should be listening on this port.
When you select the **Emulex OneCommand** tab in the vSphere console, a message appears indicating that the Adobe Flash player is required.	Ensure the version of Adobe Flash player installed is 11.2 or later. If you have not installed the Adobe Flash player, you can download it from the Adobe website.
When you select the **Emulex OneCommand** tab in the vSphere console, the Emulex Device Management tree-view does not display any elements.	Ensure that you have the required privileges to view information in the vSphere console.
There is slow response from OneCommand Manager for VMware vCenter.	Ensure that the following are on the same network: ■ ESXi servers managed by the OneCommand Manager VMware for vCenter server ■ Systems hosting the OneCommand Manager for VMware vCenter server ■ OneCommand Manager for VMware vCenter
Firmware update fails.	On the ESXi host, check the firewall settings and ensure that the HTTP/HTTPS ports are open. Use the following command to disable the firewall: `esxcli network firewall unload`
Firmware update fails with the error message `Error reading resource`.	Check the following: **NOTE** Make sure that you can run the ping command the host name, on which the OneCommand Manager for VMware vCenter is installed, from the ESXi host. If you cannot run the ping command on the host name, either reinstall providing the reachable IP or host name (with domain) or add the host name to the DNS. ■ Check the memory space in the ESXi host and clean up the old logs.
When you make any changes to the ESXi host, such as plugging cables, unplugging cables, or storage references, OneCommand Manager for VMware vCenter does not reflect the change immediately.	Click **Refresh** in the GUI. If the change is not reflected, restart `sfcb` on the ESXi host using the command: `/etc/init.d/sfcbd-watchdog restart` Click **Refresh** again in the GUI.

Table 2 Troubleshooting Issues and Resolutions (Continued)

Issue	Resolution
On a Windows 7 x64 operating system, executing the CLI commands using the executables results in unnecessary error traces.	The C disk is highly protected; even the administrator account has limited privileges. For example, the contents in the directory `C:\Program Files\` have no write and full control privileges. To remedy the problem: ■ Assign your account write and full control privileges to `C:\Program Files \`. *or* ■ Install OneCommand Manager for VMware vCenter on another disk, for example, D.
When OneCommand Manager for VMware vCenter loads within the vSphere console, it shows a Security warning.	See Section 12.1, Security.
When the VMware Web Client is launched, it requires OneCommand Manager for VMware vCenter to load in another browser tab and accept the Broadcom certificate.	See Section 12.1, Security.

Chapter 13: Using OneCommand Manager for VMware vCenter Command Line Interface

The CLI client component of OneCommand Manager for VMware vCenter is installed as part of OneCommand Manager for VMware vCenter installation.

13.1 elxvcpcmd Syntax Usage

- OneCommand Manager for VMware vCenter CLI runs only in TCP/IP mode.
- OneCommand Manager for VMware vCenter CLI can manage Emulex adapters in systems with VMware ESXi 5.5, 6.0, and 6.5 environments.
- CLI client commands are supported for Windows operating systems only.
- All commands must start with `elxvcpcmd.exe`. The `elxvcpcmd.exe` command is available in the OneCommand Manager for VMware vCenter installation directory (which is by default `C:\Program Files\Emulex`). This component is intended for use in scripted operations within batch files. Each time you run this script from the command line, a single operation is performed.
- Most operations retrieve information about an entity on the SAN and show that information on the console.
- Most of the CLI client commands require one or more additional parameters that specify the nature of the command.
- The requested operation must contain at least three characters, or as many as needed to distinguish it from any other operation.
- The parameters must be specified in the order indicated in the syntax.
- Parameters that are not required and can be omitted are in square brackets [ ].
- To run a command at the cluster level, use:

  ```
  elxvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd>
  c=<clustername> <ocm_cmd> [<ocm_cmd_arg>...]
  ```

- To run a command at the host level, use:

  ```
  elxvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd> h=<esx_host>
  <ocm_cmd> [<ocm_cmd_arg>...]
  ```

- For FC and FCoE ports, the WWPN of the adapter must be specified.

  NOTE When a WWPN is specified, individual fields are separated by colons (:).

For example, run the following command to display the port attributes for the adapter with the specified WWPN:

```
elxvcpcmd.exe  v=10.120.121.122 u=Administrator p=password h=10.120.121.123
portattributes 10:00:00:00:C9:39:6C:9D
```

- For NIC ports, the MAC address must be specified.

  NOTE When a MAC address is specified, the fields are separated by a dash (–).

For example, run the following command to show the port attributes for the CNA port with the specified MAC address:

```
elxvcpcmd.exe  v=10.120.121.122 u=Administrator p=password h=10.120.121.123
portattributes 00-11-22-33-44-55
```

- For help purposes only, commands are grouped together.

13.2 Help Commands

Help commands include help for a single command or a group of commands.

13.2.1 help (Single Command)

Syntax

`elxvcpcmd.exe [CmdName] help`

Description

This command shows the help for a specific command.

Parameters

CmdName Any CLI command.

13.2.2 help (Group)

Syntax

`elxvcpcmd.exe help`

Description

This command shows the help for a group. These help groups are categorized based on the functionality of the commands. Table 4 shows all supported groups, definitions, and supported commands. You can specify the group name in the `help` command to find the commands supported for a group.

Example Command

`elxvcpcmd.exe attributes help`

Parameters

None.

13.3 CLI Command Reference Tables

Table 3 lists CLI commands in alphabetical order, with the corresponding section number for details.

Table 4 lists the CLI commands within a help group. These help groups are categorized based on the functionality of the commands.

Table 3 CLI Commands

Command	Section
changewwn	Section 13.4.14.1
D_Port (Clearlink)	Section 13.4.6.1
deletedumpfiles	Section 13.4.8.1
driverconfig	Section 13.4.7.1
dump	Section 13.4.8.2
echotest	Section 13.4.6.2

Table 3 CLI Commands (Continued)

Command	Section
enablebootcode	Section 13.4.3.1
exportsaninfo	Section 13.4.4.1
firmwareupdate	Section 13.4.4.2
getbeacon	Section 13.4.6.3
getdcbparams	Section 13.4.5.1
getdriverparams	Section 13.4.7.2
getdriverparamsglobal	Section 13.4.7.3
getdumpdirectory	Section 13.4.8.3
getdumpfilenames	Section 13.4.8.4
getexpresslanelunlist	Section 13.4.12.1
getfcfinfo	Section 13.4.9.1
getfipparams	Section 13.4.9.2
getlunlist	Section 13.4.9.3
getpcidata	Section 13.4.2.1
getpginfo	Section 13.4.5.2
getportstatistics	Section 13.4.2.2
getretentioncount	Section 13.4.8.5
getvpd	Section 13.4.2.3
getwwncap	Section 13.4.2.4
getxcvrdata	Section 13.4.6.4
hbaattributes	Section 13.4.2.5
listhbas	Section 13.4.2.6
listvms	Section 13.4.13.1
listvports	Section 13.4.13.2
loadlist	Section 13.4.6.5
loopbacktest	Section 13.4.6.6
pcidata	Section 13.4.2.7
portattributes	Section 13.4.2.8
posttest	Section 13.4.6.7
readwwn	Section 13.4.14.3
reset	Section 13.4.1.2
restorewwn	Section 13.4.14.4
saveconfig	Section 13.4.7.4
serverattributes	Section 13.4.2.9
setbeacon	Section 13.4.6.8
setcnapgbw	Section 13.4.5.3
setdcbpriority	Section 13.4.5.5
setdriverparam	Section 13.4.7.5
setdriverparamdefaults	Section 13.4.7.6
setdumpdirectory	Section 13.4.8.6

Table 3 CLI Commands (Continued)

Command	Section
setexpresslanelunstate	Section 13.4.12.3
setfipparams	Section 13.4.9.4
setportenabled	Section 13.4.2.10
setretentioncount	Section 13.4.8.7
setvcred	Section 13.4.1.1
sriovenable	Section 13.4.11.2
targetmapping	Section 13.4.12.6
version	Section 13.4.1.3

Table 4 Help Groups and Supported CLI Commands

Help Group	Supported CLI Commands	Command Descriptions
General group – General commands that can be run on the OneCommand manager application for VMware servers.	■ setvccred ■ reset ■ version	setvcred – An optional command that saves vCenter credentials and executes subsequent commands without specifying vCenter credentials. reset – Resets the adapter. version – Shows the version of the installed CLI.
Attributes group – Commands to read and manage information about an adapter.	■ getpcidata ■ getportstatistics ■ getvpd ■ getwwncap ■ hbaattributes ■ listhbas ■ pcidata ■ portattributes ■ serverattributes ■ setportenabled	getpcidata – Lists PCI register details for FC, FCoE, and NIC ports. getportstatistics – Lists statistics for an FC or FCoE port. If the optional parameter, clear, is set, this command clears the 10GBASE-T counters. getvpd – Shows the VPD details for the specified port on the adapter. getwwncap – Shows the WWN capabilities of the specified port. hbaattributes – At the host level, shows adapter attributes for a port on the adapter. listhbas – For a cluster or a host, shows a list of manageable Emulex adapters. pcidata – Lists PCI attributes for a port on the adapter. portattributes – Shows a list of all port attributes for the port on the adapter. serverattributes – Lists basic information about the host. setportenabled – Enables or disables the port status for an FC, FCoE, and NIC port on a host.
Boot group – Commands to enable an ESXi host to manage the boot environment.	■ enablebootcode	enablebootcode - Enables or disables the bootBIOS state on a given port.
Cluster group – Commands that can be run on a cluster.	■ exportsaninfo ■ firmwareupdate ■ listhbas	exportsaninfo – Exports SAN information related to Emulex adapters in all the hosts in a cluster. firmwareupdate – Updates the firmware on the Emulex adapters found in a VMware cluster or the ESXi host. listhbas – Lists manageable Emulex adapters.

Table 4 Help Groups and Supported CLI Commands (Continued)

Help Group	Supported CLI Commands	Command Descriptions
DCB group – Commands to enable an ESXi host to read and manage data center bridging parameters for an LPe16202/OCe15100 adapter in NIC+FCoE mode.	■ getdcbparams ■ getpginfo ■ setcnapgbw ■ setdcbparam ■ setdcbpriority	`getdcbparams` – Extracts the DCB parameters associated with a given FC or NIC port. `getpginfo` – Lists the priorities and bandwidth percentages for all the priority groups for the port. `setcnapgbw` – Sets the bandwidth percentage of a priority group on a port. `setdcbparam` – Configures DCB and LLDP settings. `setdcbpriority` – Sets the priorities for a priority group.
Diagnostic group – Commands to run diagnostic tests for an Emulex adapter or port.	■ D_Port ■ echotest ■ getbeacon ■ loadlist ■ loopbacktest ■ posttest ■ setbeacon	`D_Port` – Runs D_Port diagnostics. Also known as ClearLink. `echotest` – Runs an echo test on a port. `getbeacon` – Shows the current beacon state for a port on an adapter. `loadlist` – Lists the flash memory load list data for an FC port. `loopbacktest` – Runs a loopback test on a port. `posttest` – Runs a POST on a specified FC port. `setbeacon` – Turns the beacon state on or off for a port on an adapter.
Driver Parameters group – Commands to enable an ESXi host to read and manage driver parameters for the host and an Emulex adapter.	■ driverconfig ■ getdumpdirectory ■ getdriverparamsglobal ■ saveconfig ■ setdriverparam ■ setdriverparamdefaults	`driverconfig` – Sets all driver parameters to the values in the `.dpv` file on an ESXi host. `getdumpdirectory` – Lists driver parameters for a port. `getdriverparamsglobal` – Lists global driver parameters for a port. `saveconfig` – Saves an FC or FCoE adapter's driver parameters to a file. `setdriverparam` – Sets a driver parameter for a port and designates the scope of the change. `setdriverparamdefaults` – Restores all driver parameters to the default value, at the port or global level (temporarily or permanently).
Dump group – Commands to manage dump files for a selected adapter. Dump files are useful when troubleshooting.	■ deletedumpfiles ■ dump ■ getdumpdirectory ■ getdumpfilenames ■ getretentioncount ■ setdumpdirectory ■ setretentioncount	`deletedumpfiles` – Deletes all diagnostic dump files for a port. `dump` - Performs a core dump of a local port. `getdumpdirectory` – Shows a dump file directory for a port in the host. `getdumpfilenames` – Lists all dump file names for a port. `getretentioncount` – Shows the maximum number of diagnostic dump files to keep for a port. `setdumpdirectory` – Sets the dump directory for all adapters in the server. `setretentioncount` – Specifies the maximum number of diagnostic dump files for the adapter.

Table 4 Help Groups and Supported CLI Commands (Continued)

Help Group	Supported CLI Commands	Command Descriptions
FCoE group – Commands to enable an ESXi host to read and manage FCoE functions for an Emulex adapter.	■ getfcfinfo ■ getfipparams ■ getlunlist ■ setfipparams ■ targetmapping	`getfcfinfo` – Lists Fibre Channel Forwarding information for an adapter in FCoE mode. `getfipparams` – Lists FIP parameters for an LPe16202/OCe15100 adapter in NIC+FCoE mode. `getlunlist` – Lists the LUNs attached to the FC and FCoE target, of the specified port. `setfipparams` – Sets FIP parameters for an LPe16202/OCe15100 adapter in NIC+FCoE mode. `targetmapping` – Lists the targets attached to the FC and FCoE of the specified port.
Firmware group – A command for an ESXi host to update the firmware for an Emulex adapter or port.	■ firmwareupdate	`firmwareupdate` – Updates the firmware on Emulex adapters in a VMware cluster or ESXi host.
SR-IOV group – Commands that allow multiple VFs to be created on a NIC PF for an adapter port.	■ listvfuncs ■ sriovenable	`listvfuncs` – Lists all of the virtual functions discovered. `sriovenable` – Enables or disables SR-IOV support for an adapter's physical port.
Target And Luns group – Commands to enable an ESXi host to read targets and LUNs attached to the port on an Emulex adapter, enable ExpressLane LUNs, and set frame priority levels for ExpressLane LUNs.	■ getexpresslanelunlist ■ getlunlist ■ setexpresslanelunstate ■ GetLunXLaneConfig ■ SetLunXLaneConfig ■ targetmapping	`getexpresslanelunlist` – On LPe16000-series, LPe31000-series, and LPe32000-series adapters, lists the LUNs that support ExpressLane and their respective ExpressLane status. `getlunlist` – Lists the LUNs attached to the FC and FCoE target of the specified port. `setexpresslanelunstate` – Enables or disables ExpressLane on a LUN. `GetLunXLaneConfig` – Displays the frame priority value for ExpressLane LUNs on the specified target. `SetLunXLaneConfig` – Configures the ExpressLane state and the frame priority levels, or values, for ExpressLane LUNs. `targetmapping` – Lists the targets attached to the FC and FCoE of the specified port.
Virtual Machines group – Commands to enable an ESXi host to find Virtual Machines and their attached ports.	■ listvms ■ listvports	`listvms` – Lists all virtual machines and their information for all manageable ports. `listvports` – Lists all virtual ports on the specified physical port.
WWWN Management group – Commands to enable an ESXi host to read and manage the WWN details for the port on an Emulex adapter.	■ changewwn ■ getwwncap ■ readwwn ■ restorewwn	`changewwn` – Changes the WWN of the port. `getwwncap` – Lists the WWN capabilities of the port. `readwwn` – Lists the WWN details of the port and category. `restorewwn` – Restores the WWN value of the port.

13.4 Group Commands and CLI Command Descriptions

This section provides syntax and descriptions for group and CLI commands.

13.4.1 General Group Commands

The General group commands save vCenter credentials, reset the adapter, and show the version of the installed CLI.

13.4.1.1 setvccred

13.4.1.1.1 Syntax

```
elxvcpcmd.exe setvccred v=<vCenter IP/Name> u=<username> p=<password>
```

13.4.1.1.2 Description

While executing a set of CLI commands, you must enter vCenter credentials repeatedly. By executing the `setvccred` command first, you can save these credentials including vCenter server name/IP, user name and password to a file in an encrypted format and execute subsequent commands without use of vCenter credentials.

Using this command is optional. You can continue to execute the commands providing all credentials.

NOTE This command does not apply to ports or adapters.

13.4.1.1.3 Parameters

`v` The vCenter server IP address.

`u` The user name for the vCenter server.

`p` The user password for the vCenter server.

13.4.1.1.4 Examples

Execute `setvccred` first:

```
elxvcpcmd.exe setvccred v=12.345.678.901 u=username p=password
```

Subsequent commands can be:

New format:

```
elxvcpcmd.exe h=12.345.678.123 listhbas
```

or

Old format:

```
elxvcpcmd.exe v=12.345.678.901 u=username p=password h=12.345.678.123 listhbas
```

13.4.1.2 reset

This command resets the adapter. An adapter reset can require several seconds to complete, especially for remote devices. When the reset is completed, the system command prompt is displayed.

NOTES

- Supported only for FC and FCoE ports. Not supported for NIC ports.
- For FCoE ports, this command only resets the driver to update changed driver parameters that require a driver reset. It does not cause a hardware reset of the adapter port.

13.4.1.2.1 Syntax

```
elxvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
reset <WWPN>
```

13.4.1.2.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FC or FCoE port.

13.4.1.3 version

This command shows the version of the CLI installed.

13.4.1.3.1 Syntax

```
elxvcpcmd.exe version
```

13.4.1.3.2 Parameters

None.

13.4.2 Attribute Commands

The Attribute commands show port information, adapter attributes, pci data, and server attributes. These commands also enable a port on a host and set physical port speed.

13.4.2.1 getpcidata

This command shows PCI register details for FC, FCoE, and NIC ports.

13.4.2.1.1 Syntax

```
elxvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password> h=<esx_host>
getpcidata <WWPN | MAC>
```

13.4.2.1.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of an FC or FCoE port.
MAC	The MAC address of the NIC port.

13.4.2.2 getportstatistics

This command extracts the statistics for a designated FC or FCoE port.

13.4.2.2.1 Syntax

```
elxvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password> h=<esx_host>
getportstatistics <WWPN> [clear]
```

13.4.2.2.1.1 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of an FC or FCoE port.

13.4.2.3 getvpd

This command shows the VPD details for the specified port on the adapter.

13.4.2.3.1 Syntax

```
elxvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd> h=<ESXhostIP> getvpd
<WWPN | MAC>
```

13.4.2.3.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host
WWPN	The WWPN of an FC or FCoE port.
MAC	The MAC address of the NIC port.

13.4.2.4 getwwncap

This command shows the WWN capabilities of the specified port.

13.4.2.4.1 Syntax

```
elxvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd> h=<ESXhostIP>
getwwncap <WWPN>
```

13.4.2.4.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of an FC or FCoE port.

13.4.2.5 hbaattributes

This command shows a list of all adapter attributes for the specified port on the adapter. This command is supported only at the host level.

13.4.2.5.1 Syntax

```
elxvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd> h=<ESXhostIP>
hbaattributes  <WWPN | MAC>
```

13.4.2.5.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of an FC or FCoE port.
MAC	The MAC address of the NIC port.

13.4.2.6 listhbas

This command shows a list of the manageable Emulex adapters found by remote discovery. For NIC ports, instead of the port WWN, the MAC address is displayed. The node WWN and fabric WWN are not displayed.

13.4.2.6.1 Syntax

For a cluster:

```
elxvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd> c=<clustername>
listhbas
```

For a host:

```
elxvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd> h=<ESXhostIP>
listhbas
```

13.4.2.6.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
c	The cluster name in the vSphere console.
h	The IP address of the ESXi host.

13.4.2.7 pcidata

This command shows the PCI attributes for the port specified on the adapter.

13.4.2.7.1 Syntax

```
elxvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd> h=<ESXhostIP>
pcidata <WWPN | MAC>
```

13.4.2.7.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of an FC or FCoE port.
MAC	The MAC address of the NIC port.

13.4.2.8 portattributes

This command shows a list of all port attributes for the port on the adapter.

13.4.2.8.1 Syntax

```
elxvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd> h=<ESXhostIP>
PortAttributes <WWPN | MAC>
```

13.4.2.8.2 Parameters

`v`	The vCenter server IP address.
`u`	The user name for the vCenter server.
`p`	The user password for the vCenter server.
`h`	The IP address of the ESXi host.
WWPN	The WWPN of an FC or FCoE port.
MAC	The MAC address of the NIC port.

13.4.2.9 serverattributes

This command retrieves basic information about the host such as the operating system version and CIM Provider version.

13.4.2.9.1 Syntax

```
elxvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
serverattributes
```

13.4.2.9.2 Parameters

`v`	The vCenter server IP address.
`u`	The user name for the vCenter server.
`p`	The user password for the vCenter server.
`h`	The IP address of the ESXi host.

13.4.2.10 setportenabled

This command enables or disables an FC, FCoE, or NIC port on a host.

NOTES

- Ensure that all I/O traffic on the port is stopped before disabling the port.
- When the setportenabled command disables an FC port, the adapter must be reset to activate the new setting. Only CNAs do not require a reset when the adapter port is enabled or disabled.

13.4.2.10.1 Syntax

```
elxvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
setportenabled <WWPN | MAC> <Flag>
```

13.4.2.10.2 Parameters

`v`	The vCenter server IP address.
`u`	The user name for the vCenter server.
`p`	The user password for the vCenter server.
`h`	The IP address of the ESXi host.
`WWPN`	The WWPN of an FC or FCoE port.
`MAC`	The MAC address of the NIC port.
`Flag`	`0` = Disabled `1` = Enabled

13.4.3 Boot Command

The `enablebootcode` command enables or disables the bootBIOS state on a port.

CAUTION Using the `enablebootcode` command on an LPe12000-series adapter that is being used to boot from SAN is not advisable. After the command has completed, the system performs an adapter reset, which may cause a loss of connectivity to the SAN and possible loss of data. To perform this command, you must make sure that the adapter is not currently being used to boot from SAN.

Do one of the following:

- Move the target adapter to a non-boot from SAN host.
- If the host with the target adapter is also hosting other boot from SAN adapters, perform a boot from SAN using one of the other boot from SAN adapters. The target adapter can now be used.

13.4.3.1 enablebootcode

This command enables or disables the bootBIOS state on a given port by enabling or disabling the boot code on the FC or FCoE adapter, or the PXE BIOS.

13.4.3.1.1 Syntax

```
elxvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
enablebootcode <WWPN | MAC> [Flag][NetworkBootMethod]
```

13.4.3.1.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of an FC or FCoE port.
MAC	The MAC address of the CNA port.
Flag	0 = disable BootBIOS state. 1 = enable BootBIOS state.
NetworkBootMethod	Network boot method to be used by the NIC.

13.4.4 Cluster Commands

The Cluster commands export SAN information, update firmware, and list adapters.

13.4.4.1 exportsaninfo

This command exports all SAN information related to Emulex adapters in all the hosts in a cluster. This command is supported only at the cluster level.

NOTE Due to the amount of information that must be obtained and reported, this command can take a long time to run on large SAN configurations. You can redirect this output to a file with a proper extension, .xml for XML-formatted files and .csv for CSV-formatted files.

13.4.4.1.1 Syntax

```
elxvcpcmd.exe v=<vcenter server> u=<vc_username> p=<vc_pwd> c=<clustername>
exportsaninfo [format]
```

NOTE [format] is optional. If the format parameter is specified as csv, the adapter information is shown in CSV format. If the format parameter is specified as xml, the adapter information is shown in XML format. Leaving the format parameter blank shows the data in XML format.

13.4.4.1.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
c	The cluster name.
Format	csv – The output information in CSV format. xml – The output information in XML format (default).

13.4.4.2 firmwareupdate

This command updates the firmware on the Emulex adapters found in the specified VMware cluster or the ESXi host.

13.4.4.2.1 Procedure

1. Run the `firmwareupdate` command. A list of adapter serial numbers (for LPe16000-series, LPe31000-series, and LPe32000-series adapters) or port WWNs (for LPe12000-series adapters) are displayed, applicable to the firmware file specified.
2. Select the adapter or port option to use for the update. `List`, `Range`, `All`, or `Choice` options are displayed.
3. The firmware update process begins and returns the result for each adapter or port.
4. If you press **Ctrl + C** and the firmware update process has started on any adapter or port, the update continues. But if the firmware update process is queued, the update is canceled.

NOTE The optional argument [`all`|`WWNs`|`MACs`] updates the firmware without any user prompt. Either `all` or a combination of `WWNs` and `MACs` can be given as an option.

You can view the status of submitted firmware jobs on the OneCommand Manager VMware for vCenter, **Maintenance** tab of the cluster or host.

13.4.4.2.2 Syntax

For a cluster:

```
elxvcpcmd.exe v=<vcenter server> u=<vc_username> p=<vc_pwd> c=<clustername>
firmwareupdate [all|WWNs|MACs] <filelocation>
```

For a host:

```
elxvcpcmd.exe v=<vcenter server> u=<vc_username> p=<vc_pwd> h=<ESXhostIP>
firmwareupdate [all|WWNs|MACs] <file location>
```

13.4.4.2.3 Parameters

`v`	The vCenter server IP address.
`u`	The user name for the vCenter server.
`p`	The user password for the vCenter server.
`c`	The cluster name in the vSphere console.
`h`	The IP address of the ESXi host.
`all|WWNs|MACs` (optional)	The all optional argument updates all compatible adapters or ports without any user prompt. The WWNs or MACs optional argument updates the port WWNs or MACs belonging to a specified cluster or host without any user prompt.
`file_location`	The firmware file path on the local disk.

13.4.4.3 listhbas

See Section 13.4.2.6, listhbas.

13.4.5 DCB Commands

The DCB commands show DCB parameter information, show priority and bandwidth percentages for all priority groups, configure the DCB and LLDP settings on an adapter, and set bandwidth percentage of a priority group on a port.

NOTE DCB commands apply to LPe16202/OCe15100 adapters in NIC+FCoE mode only.

13.4.5.1 getdcbparams

This command extracts the DCB parameters associated with a given FCoE or NIC port.

13.4.5.1.1 Syntax

```
elxvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
getdcbparams <WWPN | MAC>
```

13.4.5.1.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC port.

13.4.5.2 getpginfo

This command shows the priorities and bandwidth percentages for all the priority groups for the port. Additionally, this command shows the number of priority groups supported by an adapter.

13.4.5.2.1 Syntax

```
elxvcpcmd.exe  v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
getpginfo <WWPN>
```

13.4.5.2.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of an FCoE port.

13.4.5.3 setcnapgbw

This command sets the bandwidth percentage of a priority group on a given FC or NIC port according to the following rules:

- BW0 to BW7 must add up to 100 percent.
- Bandwidth can be assigned to a priority group that has priorities.

13.4.5.3.1 Syntax

```
elxvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
setcnapgbw <WWPN | MAC> <BW0>...<BW7>
```

13.4.5.3.2 Parameters

`v`	The vCenter server IP address.
`u`	The user name for the vCenter server.
`p`	The user password for the vCenter server.
`h`	The IP address of the ESXi host.
`WWPN`	The WWPN of an FCoE port.
`MAC`	The MAC address of the NIC port.
`BW0` `to` *`BW7`*	The bandwidth for priority group 0 to 7.

13.4.5.3.3 Example

This command sets the bandwidth of PGID0 to 50, PGID1 to 50 ,and the rest to 0 percent.

```
elxvcpcmd.exe setcnapgbw 10:00:00:00:c9:3c:f7:88 50 50 0 0 0 0 0 0
```

13.4.5.4 setdcbparam

This command configures the DCB and LLDP settings.

13.4.5.4.1 Syntax

```
elxvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
setdcbparam <WWPN | MAC> <Param> <Value>
```

13.4.5.4.2 Parameters

`v`	The vCenter server IP address.
`u`	The user name for the vCenter server.
`p`	The user password for the vCenter server.
`h`	The IP address of the ESXi host.
`WWPN`	The WWPN of an FCoE port.
`MAC`	The MAC address of the NIC port.

Param and Value Parameters

DCB Parameters and Values (`1` = Enabled and `0`= Disabled)

`DCBXState`	The current DCBX (Data Center discovery and Capability exchange protocol) state.
`PFCEnable`	If set to `1`, there is flow control in both directions (Tx and Rx).
`FCoEPriority`	You must specify a single priority (0 to 7). For FC ports only.
`PFCPriorities`	Specify `PFCPriorities` as a single priority or as a list of comma separated values (up to 7 values ranging from 0 to 7).
`Default`	Setting the SetDCBParam arguments to `default` sets all CNA DCB parameters (including priority groups) to their default values.

Link Layer Discovery Protocol (LLDP) Parameters and Values

`TxState`	Transmit State – LLDP to exchange parameters between two link peers. For the DCBX protocol to operate correctly, both LLDP Rx and Tx must be enabled. If either Rx or Tx is disabled, DCBX is disabled.
`RxState`	Receive State – DCBX uses LLDP to exchange parameters between two link peers. For the DCBX protocol to operate correctly, both LLDP Rx and Tx must be enabled. If either Rx or Tx is disabled, DCBX is disabled.
`TxPortDesc`	Transmit Port Description – Provides a description of the port in an alpha-numeric format. If the LAN device supports RFC 2863, the value equals the ifDescr object.
`TxSysDesc`	Transmit System Description – Provides a description of the network entity in an alpha-numeric format. This includes the system's name and versions of hardware, operating system, and networking software supported by the device. If the LAN device supports RFC 3418, the value equals the sysDescr object.
`TxSysName`	Provides the system's assigned name in an alpha-numeric format. If the LAN device supports RFC 3418, the value equals the sysName object.
`TxSysCap`	Transmit System Capabilities – Indicates the primary functions of the device and whether these functions are enabled on the device. The capabilities are indicated by two octets. Bits 0 through 7 indicate Other, Repeater, Bridge, WLAN AP, Router, Telephone, DOCSIS cable device, and Station respectively. Bits 8 through 15 are reserved. Where multiple values are possible, specify them with a comma-separated list.

13.4.5.4.3 Example Command

```
elxvcpcmd v=12.345.678.901 u=username p=password h=12.345.678.123 setdcbparam
00-00-c9-3c-f7-88 fcoepriority 3
```

13.4.5.5 setdcbpriority

This command sets the priorities for a priority group.

13.4.5.5.1 General Rules

- The priorities range from 0 to 7.
- The PGIDs range from 0 to 7.
- A priority can exist in only one priority group.
- All priorities must appear once in any of the eight (PG0 – PG7) priority groups.
- Each set of priorities for a group must be separated by a space.
- Specify multiple priorities for the same group by a comma-separated list.

- To specify none, use – for the argument.
- The same priority values cannot be specified to different groups.
- All priorities (0 to 7) must be assigned to a PGID.
- Not all PGIDs must be assigned a priority.
- Not all adapters support two PFC priorities and eight priority groups. For adapters, if you exceed the PFC priorities or priority groups, an error message appears.

13.4.5.5.2 FCoE Port-Specific Rules

- A minimum of one and a maximum of two PFC priorities can be configured.
- One of the PFC priorities must match FCoE priority.
- The additional PFC priority must be assigned to a priority group that has no other priorities.
- The FCoE priority must be assigned to a priority group that has no other priorities.

13.4.5.5.3 NIC Port-Specific Rules

- Only one PFC priority can be configured.
- The PFC priority must be assigned to a priority group that has no other priorities.

13.4.5.5.4 Syntax

```
elxvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
setdcbpriority <WWN | MAC> <PFC Priorities> <Priorities of PGID0> < Priorities of
PGID1>...<Priorities of PGID7>
```

13.4.5.5.5 Parameters

`v`	The vCenter server IP address.
`u`	The user name for the vCenter server.
`p`	The user password for the vCenter server.
`h`	The IP address of the ESXi host.
`WWPN`	The WWPN of an FCoE port.
`MAC`	The MAC address of the NIC port.
`PFCPriorities`	PFC priority that is a comma separated list of up to 7 values ranging from 0–7.
`Priorities of PGID`	Priority group membership that is a comma separated list of priorities ranging in value for 0–7.

13.4.5.5.6 Example Command

```
elxvcpcmd v=12.345.678.901 u=username p=password h=12.345.678.123 setdcbpriority
10:00:00:00:c9:3c:f7:88 3 0,1,2,4,5,6,7 3 0 0 0 0 0 0
```

13.4.6 Diagnostic Commands

The Diagnostic commands run diagnostics, including POST and loopback. Diagnostic commands also show and set beacon, and run the D_Port diagnostic (LPe16000-series, LPe31000-series, and LPe32000-series adapters).

CAUTION Using the loopback or POST test commands on an LPe12000-series adapter that is being used to boot from SAN is not advisable. After the command has completed, the system performs an adapter reset, which may cause a loss of connectivity to the SAN and possible loss of data. To perform these commands, you must make sure that the adapter is not currently being used to boot from SAN.

Do one of the following:

- Move the target adapter to a non-boot from SAN host.
- If the host with the target adapter is also hosting other boot from SAN adapters, perform a boot from SAN using one of the other boot from SAN adapters. The target adapter can now be used.

13.4.6.1 D_Port

The D_Port diagnostic is also known as ClearLink. The D_Port diagnostic tests are run from the OneCommand Manager for VMware vCenter CLI by specifying the D_Port command. D_Port is a diagnostic mode supported by Brocade switches for LPe16000-series, LPe31000-series, and LPe32000-series adapters with D_Port support.

D_Port tests detect physical cabling issues that can result in increased error rates and intermittent behavior. When activated, D_Port tests include:

- Local electrical loopback
- Loopback to the remote optics
- Loopback from the remote port to the local optics
- A full device loopback test with data integrity checks
- An estimate of cable length (to validate that a proper buffering scheme is in place)

These tests allow a level of fault isolation to distinguish faults due to marginal cable, optics modules, and connector or optics seating.

Considerations for running D_Port

- It is not possible to detect if the switch can run D_Port tests before running the tests. Therefore, a test failure occurs if the D_Port command is run with a switch that does not support D_Port.
- To terminate tests while they are running, type <CTL> + <C>. In this case, no results are given.
- If the overall test result is FAILED, you must re-run the tests successfully or reset the HBA port to bring the link back up. A message is displayed instructing you to perform one of these actions if the overall test result is FAILED.
- If a test phase fails, the D_Port tests are automatically stopped. In this case, some of the phases may not be reported in the results. However, the failed phase is reported.
- More than one error can be reported. In this case, multiple lines are displayed for the test phase showing each error.

13.4.6.1.1 Example

```
elxvcpcmd.exe v=10.192.000.000 u=root p=password d_Port WWPN
```

13.4.6.1.2 Parameters

`v`	The vCenter server IP address.
`u`	The user name for the vCenter server.
`p`	The user password for the vCenter server.
`WWPN`	The WWPN of the FC or FCoE port on which to run tests.

13.4.6.2 echotest

This command runs a loopback test on a given port.

13.4.6.2.1 Syntax

```
elxvcpcmd.exe  v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
echo <WWN | MAC> <Destination WWPN> <Count> <StopOnError> [Pattern]
```

13.4.6.2.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of an FC or FCoE port.
MAC	The MAC address of the NIC port.
Destination WWPN	The WWPN of the destination (echoing) adapter.
Count	The number of times to run the test (0 = run test infinitely, Range = 1 to 99,999).
StopOnError	Checks if the test must be halted on error. 0 = No halt 1 = Halt
Pattern (optional)	1 to 8 hexadecimal bytes to use for loopback data (for example: 1a2b3c4d).

13.4.6.2.3 Example

```
elxvcpcmd.exe  v=10.20.30.40 u=user p=password  h=1.2.3.4 echo
10:00:00:c9:12:34:56 10:00:00:c9:ab:cd:ee 100 1 1a2b3c4d5e
```

13.4.6.3 getbeacon

This command shows the current beacon state, on or off.

13.4.6.3.1 Syntax

```
elxvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
getbeacon <WWPN | MAC>
```

13.4.6.3.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FC or FCoE port.
MAC	The MAC address of the NIC port.

13.4.6.4 getxcvrdata

This command shows transceiver data, such as vendor name and serial number.

13.4.6.4.1 Syntax

```
elxvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
getxcvrdata <WWN | MAC>
```

13.4.6.4.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FC or FCoE port.
MAC	The MAC address of the NIC port.

13.4.6.5 loadlist

This command shows the flash parameters for a given FC port.

13.4.6.5.1 Syntax

```
elxvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
loadlist <WWPN>
```

13.4.6.5.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FC port.

13.4.6.6 loopbacktest

This command runs a loopback test on a given FC, FCoE, or NIC port.

13.4.6.6.1 Syntax

```
elxvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
loopbacktest <WWPN> <Type> <Count> <StopOnError> [Pattern]
```

13.4.6.6.2 Parameters

`v`	The vCenter server IP address.
`u`	The user name for the vCenter server.
`p`	The user password for the vCenter server.
`h`	The IP address of the ESXi host.
`WWPN`	The WWPN of an FC or FCoE port.
`Type`	The type of loopback test to run: `0` = PCI Loopback Test `1` = Internal Loopback Test `2` = External Loopback Test (requires loopback plug)
`Count`	The number of times to run the test (0 = run test infinitely, Range = 1 to 99,999)
`StopOnError`	Checks if the test must be halted on error. `0` = No halt `1` = Halt
`Pattern` `(optional)`	1 to 8 hexadecimal bytes to use for loopback data (for example: 1a2b3c4d)

13.4.6.6.3 Example Command

```
elxvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password> h=<esx_host>
loopbacktest 00-00-c9-93-2f-9f 4 120 0
```

13.4.6.7 posttest

This command runs the POST on a specified FC port.

NOTE The `posttest` command is available only for LPe12000-series adapters.

13.4.6.7.1 Syntax

```
elxvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
posttest <WWPN>
```

13.4.6.7.2 Parameters

`v`	The vCenter server IP address.
`u`	The user name for the vCenter server.
`p`	The user password for the vCenter server.
`h`	The IP address of the ESXi host.
`WWPN`	The WWPN of the FC port.

13.4.6.8 setbeacon

This command turns the beacon on or off.

13.4.6.8.1 Syntax

```
elxvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
setbeacon <WWN | MAC> <BeaconState>
```

13.4.6.8.2 Parameters

`v`	The vCenter server IP address.
`u`	The user name for the vCenter server.
`p`	The user password for the vCenter server.
`h`	The IP address of the ESXi host.
`WWPN`	The WWPN of an FC or FCoE port.
`MAC`	MAC address of the NIC port.
`BeaconState`	New state of the beacon: `0` = Off `1` = On

13.4.7 Driver Parameter Commands

The Driver Parameter commands show, set, and save driver parameter values. You can also change the parameters back to factory default values.

13.4.7.1 driverconfig

This command sets all driver parameters to the values in the `.dpv` file on a particular ESXi host. The `.dpv` file's driver type must match the driver type of the host platform FC or FCoE adapter.

13.4.7.1.1 Syntax

```
elxvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
driverconfig <WWPN> <FileName> <Flag>
```

13.4.7.1.2 Parameters

`v`	The vCenter server IP address.
`u`	The user name for the vCenter server.
`p`	The user password for the vCenter server.
`h`	The IP address of the ESXi host.
`WWPN`	The WWPN of the FC or FCoE port.
`FileName`	The name of the `.dpv` file (stored in the Emulex Repository directory)
`Flag`	`G` = Makes change global (all adapters on this host) N = Makes change non-global (adapter-specific)

13.4.7.2 getdriverparams

This command shows the driver parameters of the specified port.

13.4.7.2.1 Syntax

```
elxvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd> h=<ESXhostIP>
getdriverparams <WWPN | MAC>
```

13.4.7.2.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FC or FCoE port.
MAC	The MAC address of the NIC port.

13.4.7.3 getdriverparamsglobal

This command shows the global driver parameters of the specified port.

13.4.7.3.1 Syntax

```
elxvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd> h=<ESXhostIP>
getdriverparamsglobal <WWPN | MAC>
```

13.4.7.3.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FC or FCoE port.
MAC	The MAC address of the NIC port.

13.4.7.4 saveconfig

This command saves the specified FC or FCoE adapter's driver parameters to a file on an ESXi host. The resulting file contains a list of driver parameter definitions in ASCII file format with definitions delimited by a comma. Each definition is of the form:

<parameter-name>=<parameter-value>

The command saves either the values of the global set, or those specific to the adapter in the Emulex Repository directory.

NOTE Driver parameters that are set temporarily and globally (using the G and T flags) must be read using the `getdriverparamsel xvcpcmd` command to view the current value of the parameter. The `getdriverparamsglobal elxvcpcmd` command returns only permanently set driver parameter values. Additionally, if temporary, global values have been set for one or more driver parameters, the `saveconfig elxvcpcmd` command must be run with the N flag (using the N flag is analogous to the `elxvcpcmd` command `getdriverparams`) to force the driver parameter values for the specified adapter to be saved. Inaccurate values may be saved if the G flag is used for this command.

13.4.7.4.1 Syntax

```
elxvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
saveconfig <WWN> <FileName> <Flag>
```

13.4.7.4.2 Parameters

`v`	The vCenter server IP address.
`u`	The user name for the vCenter server.
`p`	The user password for the vCenter server.
`h`	The IP address of the ESXi host.
`WWPN`	The WWPN of the FC or FCoE port.
`FileName`	The name of the local `.dpv` file.
`Flag`	Valid types are: `G` = Make change global (all adapters on this host) `N` = Make change non-global (adapter-specific)

13.4.7.5 setdriverparam

This command sets the driver parameter at the port or global level, either permanently or temporarily, for the specified port.

13.4.7.5.1 Syntax

```
elxvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd> h=<ESXhostIP>
setdriverparam <WWPN | MAC> <Flag1> <Flag2> <Param> <Value>
```

13.4.7.5.2 Parameters

`v`	The vCenter server IP address.
`u`	The user name for the vCenter server.
`p`	The user password for the vCenter server.
`h`	The IP address of the ESXi host.
`WWPN`	The WWPN of an FC or FCoE port.
`MAC`	The MAC address of the NIC port.
`Flag1`	`L` = Local (all adapters on this host) `G` = Global (all adapters on this host)
`Flag2`	`P` = Permanent (persists across re-boot) `T` = Temporary.
`Param`	The name of the driver parameter.
`Value`	The value of the driver parameter.

13.4.7.6 setdriverparamdefaults

This command restores the driver parameter to the default value at the port or global level, either permanently or temporarily, for the specified port.

13.4.7.6.1 Syntax

```
elxvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd> h=<ESXhostIP>
setdriverparamdefaults <WWPN | MAC> <Flag1> <Flag2>
```

13.4.7.6.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of an FC or FCoE port.
MAC	The MAC address of the NIC port.
Flag1	L = Local (all adapters on this host) G = Global (all adapters on this host)
Flag2	P = Permanent (persists across re-boot) T = Temporary.

13.4.8 Dump Commands

The Dump commands initiate a core dump on a local port, show the dump file for the port on the host, show the diagnostic dump file retention count set on a port, and specify the maximum number of diagnostic dump files for the adapter.

13.4.8.1 deletedumpfiles

This command deletes all diagnostic dump files for a given port.

13.4.8.1.1 Syntax

```
elxvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
deletedumpfiles <WWPN | MAC>
```

13.4.8.1.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of an FC or FCoE port.
MAC	The MAC address of the NIC port.

13.4.8.2 dump

This command performs a core dump on a local port. The dump file is placed in the dump directory with the following file name format:

```
<Hostname>_<Adapter_serial_number>_<datetimestamp>.core
```

If the command is successful, the following message is displayed:

```
Dump Successful.
```

NOTE Because this command is dumping memory, it can take time while generating large files.

13.4.8.2.1 Syntax

```
elxvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
dump <WWPN | MAC> core [Options]
```

13.4.8.2.2 Parameters

`v`	The vCenter server IP address.
`u`	The user name for the vCenter server.
`p`	The user password for the vCenter server.
`h`	The IP address of the ESXi host.
`WWPN`	The WWPN of the FC or FCoE port.
`MAC`	The MAC address of the CNA port.
`Options`	Additional options are available under the direction of Broadcom Technical Support.

13.4.8.3 getdumpdirectory

This command shows the dump file directory.

NOTE The dump directory applies to all adapters in the server. There is not a separate dump directory for each adapter.

13.4.8.3.1 Syntax

```
elxvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
getdumpdirectory <WWPN | MAC>
```

13.4.8.3.2 Parameters

`v`	The vCenter server IP address.
`u`	The user name for the vCenter server.
`p`	The user password for the vCenter server.
`h`	The IP address of the ESXi host.
`WWPN`	The WWPN of an FC or FCoE port.
`MAC`	The MAC address of the NIC port.

13.4.8.4 getdumpfilenames

This command shows a list of all the dump file names for a given port.

13.4.8.4.1 Syntax

```
elxvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
getdumpfilenames <WWPN | MAC>
```

13.4.8.4.2 Parameters

`v`	The vCenter server IP address.
`u`	The user name for the vCenter server.
`p`	The user password for the vCenter server.
`h`	The IP address of the ESXi host.
`WWPN`	The WWPN of an FC or FCoE port.
`MAC`	The MAC address of the NIC port.

13.4.8.5 getretentioncount

This command shows the diagnostic dump file retention count set on a port.

NOTE The retention count applies to all adapters in the server.

13.4.8.5.1 Syntax

```
elxvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
getretentioncount <WWPN | MAC>
```

13.4.8.5.2 Parameters

`v`	The vCenter server IP address.
`u`	The user name for the vCenter server.
`p`	The user password for the vCenter server.
`h`	The IP address of the ESXi host.
`WWPN`	The WWPN of an FC or FCoE port.
`MAC`	The MAC address of the NIC port.

13.4.8.6 setdumpdirectory

Use the `setdumpdirectory` command to set the dump directory for a given port. To use the `setdumpdirectory` command, you must have a directory mapped under `/vmfs/volumes/` where the files will be placed.

NOTE The dump directory applies to all adapters in the server. There is no separate dump directory for each adapter.

13.4.8.6.1 Syntax

```
elxvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
setdumpdirectory <WWPN | MAC> <DumpDirectory>
```

13.4.8.6.2 Parameters

`v`	The vCenter server IP address.
`u`	The user name for the vCenter server.
`p`	The user password for the vCenter server.
`h`	The IP address of the ESXi host.
`WWPN`	The WWPN of an FC or FCoE port.
`MAC`	The MAC address of the NIC port.
`DumpDirectory`	The directory under `/vmfs/volumes` that you created to store the dump files.

13.4.8.6.3 Example Command

```
elxvcpcmd v=12.345.678.901 u=username p=password h=12.345.678.123
setdumpdirectory 10:00:00:00:c9:61:f2:64 vcenter-datastore
```

This example shows the dump directory set to `/vmfs/volumes/vcenter-datastore`.

13.4.8.7 setretentioncount

This command specifies the maximum number of diagnostic dump files for the adapter. When the count reaches the limit, the next dump operation deletes the oldest file.

The retention count applies to all adapters in the server.

13.4.8.7.1 Syntax

```
elxvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
setretentioncount <WWPN | MAC> <RetentionCount>
```

13.4.8.7.2 Parameters

`v`	The vCenter server IP address.
`u`	The user name for the vCenter server.
`p`	The user password for the vCenter server.
`h`	The IP address of the ESXi host.
`WWPN`	The WWPN of an FC or FCoE port.
`MAC`	The MAC address of the NIC port.
`RetentionCount`	The number of dump files to retain.

13.4.9 FCoE Commands

The FCoE commands, which are only available for LPe16202/OCe15100 adapters in NIC+FCoE mode, show the FCF information of an adapter, the FIP parameters of the adapter, and the LUNs attached to the target of the specified port. The FCoE commands also set the FIP parameters of the adapter in FCoE mode and show the targets attached to the FCoE port.

13.4.9.1 getfcfinfo

This command shows the FCF information of an adapter in FCoE mode

13.4.9.1.1 Syntax

```
elxvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>
h=<esx_host> getfcfinfo <WWPN>
```

13.4.9.1.2 Parameters

`v`	The vCenter server IP address.
`u`	The user name for the vCenter server.
`p`	The user password for the vCenter server.
`h`	The IP address of the ESXi host.
`WWPN`	The WWPN of an FC or FCoE port.

13.4.9.2 getfipparams

This command shows the FIP parameters.

13.4.9.2.1 Syntax

```
elxvcpcmd.exe  v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
getfipparams <WWPN>
```

13.4.9.2.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of an FC or FCoE port.

13.4.9.3 getlunlist

This command shows the LUNs attached to an FC or FCoE target of the specified FC or FCoE port.

13.4.9.3.1 Syntax

```
elxvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd> h=<ESXhostIP>
getlunlist <WWPN> <TargetWWN>
```

13.4.9.3.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of an FC or FCoE port.
Target WWPN	The WWPN of the FC or FCoE target.

13.4.9.4 setfipparams

This command sets the FIP parameters.

13.4.9.4.1 Syntax

```
elxvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
setfipparams <WWPN> <Param> <Value>
```

13.4.9.4.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of an FCoE port.
Param	FIP parameter name (`pfabric`, `pswitch`, `vlanid`, `fcmap`, and `cinvlanid`).
Value	FIP parameter value and the valid range for the FIP parameter. `pfabric` – 8-byte fabric name (format XX:XX:XX:XX:XX:XX:XX:XX) `pswitch` – 8-byte switch name (format XX:XX:XX:XX:XX:XX:XX:XX) `vlanid` – 2-byte VLAN_ID [0-4095] OR 'any' for any valid VLAN_ID `fcmap` – 3-byte FC_map, 0x0EFCxx `cinvlanid` – 2-byte VLAN_ID [0–4095]

13.4.9.4.3 Example Command

```
elxvcpcmd  v=12.345.678.901 u=username p=password h=12.345.678.123 setfipparams
10:00:00:00:c9:5b:3a:6d fcmap 0x0efc99
```

13.4.9.5 targetmapping

This command shows the targets attached to the specified FC or FCoE port.

13.4.9.5.1 Syntax

```
elxvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd> h=<ESXhostIP>
targetmapping <WWPN>
```

13.4.9.5.2 Parameters

`v`	The vCenter server IP address.
`u`	The user name for the vCenter server.
`p`	The user password for the vCenter server.
`h`	The IP address of the ESXi host.
`WWPN`	The WWPN of an FC or FCoE port.

13.4.10 Firmware Commands

The `firmwareupdate` command updates firmware on the Emulex adapters found in the specified VMware cluster or the ESXi host.

13.4.10.1 firmwareupdate

See Section 13.4.4.2, firmwareupdate.

13.4.11 SR-IOV Commands

The SR-IOV commands show a list of all the discovered virtual functions and also enables or disables SR-IOV support for an adapter's physical port.

13.4.11.1 listvfuncs

This command shows a list of all the virtual functions discovered.

13.4.11.1.1 Syntax

```
elxvcpcmd.exe  v=<vcenter_server> u=<vc_username> p=<vc_password> h=<esx_host>
listvfuncs <MAC>
```

13.4.11.1.2 Parameters

`v`	The vCenter server IP address.
`u`	The user name for the vCenter server.
`p`	The user password for the vCenter server.
`h`	The IP address of the ESXi host with Emulex adapters.
`MAC`	MAC address of any NIC function on the physical adapter port

13.4.11.2 sriovenable

This command enables or disables SR-IOV support for an adapter's physical port.

13.4.11.2.1 Syntax

```
elxvcpcmd.exe  v=<vcenter_server> u=<vc_username> p=<vc_password> h=<esx_host>
sriovenable <MAC>  <0 | 1>
```

13.4.11.2.2 Parameters

`v`	The vCenter server IP address.
`u`	The user name for the vCenter server.
`p`	The user password for the vCenter server.
`h`	The IP address of the ESXi host with Emulex adapters.
`MAC`	The MAC Address of any NIC function on the physical adapter port.
`SR-IOV Enable state`	`0` = Disable SR-IOV `1` = Enable SR-IOV

NOTE SR-IOV is not supported with UMC channel management.

13.4.12 Target and LUN Commands

The Target and LUN commands show the LUNs that support ExpressLane and their status, and also show LUNs attached to the target of the port.

13.4.12.1 getexpresslanelunlist

This command shows the LUNs that support ExpressLane and their respective ExpressLane status, and is supported on LPe16000-series, LPe31000-series, and LPe32000-series adapters only.

NOTE Only LUNs that support ExpressLane are shown by this command. LUNS without ExpressLane support are not shown.

13.4.12.1.1 Syntax

```
elxvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_pwd> h=<ESXhostIP>
getexpresslanelunlist <WWPN> <Target_WWPN> <Option>
```

13.4.12.1.2 Parameters

`v`	The vCenter server IP address.
`u`	The user name for the vCenter server.
`p`	The user password for the vCenter server.
`h`	The IP address of the ESXi host.
`WWPN`	The WWPN of the FC or FCoE port connected to the target.
`Target WWPN`	The WWPN of the FC or FCoE target.
`Option`	Valid options: `0` = Get information from the driver `1` = Get information from configuration

13.4.12.1.3 Example

```
elxvcpcmd.exe v=10.192.000.000 u=user p=password h=10.192.87.198
getexpresslanelunlist 10:00:00:00:00:87:01:98 20:22:d4:ae:52:6e:6f:08 0
```

13.4.12.2 getlunlist

See Section 13.4.9.3, getlunlist.

13.4.12.3 setexpresslanelunstate

NOTE This command has been deprecated. Use the `SetLunXLaneConfig` command to configure ExpressLane. The setting has been saved to both the driver and configuration to match the behavior of the SetLunXLaneConfig command.

This command enables and disables ExpressLane on a particular LUN, and is supported on LPe16000-series, LPe31000-series, and LPe32000-series adapters only.

NOTE A masked LUN cannot be enabled for ExpressLane.

13.4.12.3.1 Syntax

```
setexpresslanelunstate <WWPN> <Target_WWPN> <Lun> <State> <Option>
```

13.4.12.3.2 Parameters

`WWPN`	The WWPN of the FC or FCoE port connected to the target.
`Target WWPN`	The WWPN of the target.
`Lun`	LUN number to enable/disable ExpressLane.
`State`	`0` = Disable ExpressLane `1` = Enable ExpressLane
`Option`	Valid ExpressLane options: `0` = Set ExpressLane LUN state in driver (active) `1` = Set ExpressLaneLUN state in configuration (persist across reboots) `2` = Set ExpressLane LUN state in both (driver and configuration)

13.4.12.3.3 Example Commands

The following command enables ExpressLane for LUN 2 on specified target, in the driver and configuration.

```
elxvcpcmd.exe v=10.192.000.000 u=user p=password h=10.192.87.198
setexpresslanelunstate 10:00:00:00:00:87:01:98 20:22:d4:ae:52:6e:6f:08 2 1 2
```

The following command enables ExpressLane for LUN 2 on a specified target, in the driver only. After a reboot, a LUN is not an ExpressLane LUN.

```
elxvcpcmd.exe v=10.192.000.000 u=user p=password h=10.192.87.198 setoaslunstate
10:00:00:00:00:87:01:98 20:22:d4:ae:52:6e:6f:08 2 1 0
```

The following command disables ExpressLane for LUN 2 on a specified target, in the driver and configuration.

```
elxvcpcmd.exe v=10.192.000.000 u=user p=password h=10.192.87.198 setoaslunstate
10:00:00:00:00:87:01:98 20:22:d4:ae:52:6e:6f:08 2 0 2
```

The following command disables ExpressLane for LUN 2 on a specified target, in the driver only. Assuming this is configured to be an ExpressLane LUN, after a reboot LUN 2 is an ExpressLane LUN again.

```
elxvcpcmd.exe v=10.192.000.000 u=user p=password h=10.192.87.198 setoaslunstate
10:00:00:00:00:87:01:98 20:22:d4:ae:52:6e:6f:08 2 0 0
```

13.4.12.4 GetLunXLaneConfig

This command displays the frame priority value for ExpressLane LUNs on the specified target.

Syntax

`GetLunXLaneConfig <WWPN> [vport=<vPort_WWPN>] <Target_WWPN> <Option>`

Parameters

`WWPN`	World-wide port name of any FC/FCoE function on the adapter.
`vPort WWPN`	The WWPN of an optional vPort.
`Target WWPN`	The WWPN of the target connected to the LUNs.
`Option`	`0` = Get the information from the driver. `1` = Get the information from the configuration.

13.4.12.4.1 Example Command

```
GetLunXLaneConfig 10:00:00:00:c9:55:55:56 50:06:01:60:10:20:5C:38 0

Supported Priority Levels: High, Medium, Low

                                Priority   Priority
    FCP_LUN           OS_LUN    Level      Value
------------------- -------- ---------- ----------
0000 0000 0000 0000    0        High       113
0001 0000 0000 0000    1        High       113
0002 0000 0000 0000    2        Medium     92
0009 0000 0000 0000    9        Low        53
```

13.4.12.5 SetLunXLaneConfig

If the adapter and switch support it, the `SetLunXLaneConfig` command enables you to configure the ExpressLane Optimized Access Storage (OAS) state and the frame priority levels, or values, for ExpressLane LUNs. Switches can provide up to three priority levels; Low, Medium, and High, but they might provide fewer options.

NOTE If the switch connected to the FC or FCoE initiator does not support LUN-specific frame priority levels using the Get Fabric Object (GFO), you must manually enter the frame priority values in the range of 0–127 for all ExpressLane-enabled LUNs.

You can also use the `SetLunXLaneConfig` command to disable ExpressLane on all LUNs in a single operation.

Use the `GetLunXLaneConfig` command to determine if frame priority levels are supported.

The following rules apply when using the `SetLunXLaneConfig` command:

- The ExpressLane (OAS) state and priority levels, or values, will be automatically saved to both the driver and configuration settings. You cannot specify where to save the configuration.
- The priority parameter is only required if the ExpressLane state parameter is set to `1` (enable).
  - — You cannot disable the ExpressLane OAS state with the priority level, or value, parameter set.
  - — You cannot disable all LUNs with the priority level, or value, parameter set.
- The `EnableXLane` driver parameter must be enabled.

Syntax

`SetLunXLaneConfig <WWPN> [vport=<vPort_WWPN>] <Target_WWPN> <LUN> <State> <PriorityLevel|PriorityValue>`

Parameters

WWPN	The World-wide port name of any FC or FCoE function on the adapter.
vPort WWPN	The WWPN of an optional vPort.
Target WWPN	The WWPN of the target connected to the LUNs.
LUN	The LUN number. (The OS_LUN from the GetExpressLaneLunList command to set.) Use ALL to disable ExpressLane for all LUNs.
State	0=disable ExpressLane, 1=enable ExpressLane
PriorityLevel	The levels are high, medium, or low. (Use the GetSmartSanPriorities command to get an accurate list of the supported priority levels). This value must be omitted if the *State* = 0.
PriorityValue	A value within the range of 0–127. This value is only allowed if priority levels are not supported by the switch. This value must be omitted if the *State* =0.

13.4.12.5.1 Example Command

```
SetLunXLaneConfig 10:00:00:00:c9:55:55:56 50:06:01:60:10:20:5C:38 0 1 low

ExpressLane configuration successfully changed on the specified LUN(s)
```

13.4.12.6 targetmapping

See Section 13.4.9.5, targetmapping.

13.4.13 Virtual Machines Commands

13.4.13.1 listvms

The Virtual Machines commands list all virtual machines and their information for all manageable ports.

If you specify the host with the h=<*esx_host*> option or just give the physical WWPN, only the virtual machines for that host are shown. If you specify the physical port and the virtual port, only the virtual machine for the specified virtual port returns.

The virtual machine name is only displayed if the virtual port is associated with a virtual machine on VMware ESXi and 5.1. If you are running this command on any other server that has virtual ports, you will not see the virtual machine name.

13.4.13.1.1 Syntax

```
elxvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
listvms <WWPN>
```

13.4.13.1.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of an FC or FCoE port.

13.4.13.2 listvports

This command lists virtual ports on the specified physical port. Leaving the physical WWPN parameter blank lists all virtual ports on all manageable hosts that support virtual ports.

13.4.13.2.1 Syntax

```
elxvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
listvports <WWPN>
```

13.4.13.2.2 Parameters

`v`	The vCenter server IP address.
`u`	The user name for the vCenter server.
`p`	The user password for the vCenter server.
`h`	The IP address of the ESXi host.
`WWPN`	The WWPN of an FC or FCoE port.

13.4.14 WWWN Management Commands

CAUTION Using the `changewwn` or `restorewwn` commands on an LPe12000-series adapter) that is being used to boot from SAN is not advisable. After the command has completed, the system performs an adapter reset, which may cause a loss of connectivity to the SAN and possible loss of data. To perform these commands, you must make sure that the adapter is not currently being used to boot from SAN.

Do one of the following:

- Move the target adapter to a non-boot from SAN host.
- If the host with the target adapter is also hosting other boot from SAN adapters, perform a boot from SAN using one of the other boot from SAN adapters. The target adapter can now be used.

13.4.14.1 changewwn

This command changes the WWN of the specified port.

13.4.14.1.1 Syntax

```
elxvcpcmd.exe  v=<vcenter_server> u=<vc_username> p=<vc_pwd> h=<ESXhostIP>
changewwn <WWPN> <New_WWPN> <New_WWNN> <ReadType>
```

13.4.14.1.2 Parameters

`h`	The IP address of the ESXi host.
`WWPN`	The WWPN of the FC or FCoE port.
`New WWPN`	The new WWPN of the FC or FCoE port.
`New WWNN`	The new WWNN of the FC or FCoE port.
`ReadType`	`0` = Volatile, `1` = Non-volatile.
`v`	The vCenter server IP address.
`u`	The user name for the vCenter server.
`p`	The user password for the vCenter server.

13.4.14.2 getwwncap

See Section 13.4.2.4, getwwncap.

13.4.14.3 readwwn

This command shows the WWN details of the specified port and category.

13.4.14.3.1 Syntax

```
elxvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_pwd> h=<ESXhostIP>
showpersonalities <WWPN | MAC> [ReadType]
```

13.4.14.3.2 Parameters

`v`	The vCenter server IP address.
`u`	The user name for the vCenter server.
`p`	The user password for the vCenter server.
`h`	The IP address of the ESXi host.
`WWPN`	The WWPN of an FC or FCoE port.
`MAC`	The MAC address of the NIC port.
`ReadType`	Valid types are: `0` = Volatile `1` = Non-volatile `2` = Factory default `3` = Current `4`= Configured

13.4.14.4 restorewwn

This command restores the WWN value of the specified port.

13.4.14.4.1 Syntax

```
elxvcpcmd.exe  v=<vcenter_server> u=<vc_username> p=<vc_pwd> h=<ESXhostIP>
restorewwn <WWPN> <RestoreType>
```

13.4.14.4.2 Parameters

`v`	The vCenter server IP address.
`u`	The user name for the vCenter server.
`p`	The user password for the vCenter server.
`h`	The IP address of the ESXi host.
`WWPN`	The WWPN of an FC or FCoE port.
`RestoreType`	Valid types are: `0` = Restore default WWNs `1` = Restore NVRAM WWNs

13.5 Viewing Audit Logs Using the CLI Command

You can use the `elxvcpaudit.exe` script to log all the historical active management performed through the vSphere console on the Emulex adapters. To see the usage information, run the script with no parameters specified. All supported events are displayed.

NOTE All active management actions performed are saved to a log file specific to the action. The maximum size of a log file is 2 MB. If the size of the log file exceeds this limit, old log entries are deleted for the particular event.

13.5.0.0.1 Syntax

```
elxvcpaudit.exe <event_name>
```

NOTE If an event name is not specified, all events are displayed.

13.5.0.0.2 Parameter

`event name`	The command name describing the appropriate active management action performed.

13.5.0.0.3 List of Supported Event Names

`loopbacktest`	Diagnostic tests.
`setdriverparam`	Driver parameters changes, at both port and global levels.
`setdcbparam`	DCB parameter changes (LPe16202/OCe15100 adapters in NIC+FCoE mode).
`setbeacon`	Beacon setting changes.
`setfipparam`	FIP parameter changes (LPe16202/OCe15100 adapters in NIC+FCoE mode).
`download`	Firmware download activities.
`reset`	Port reset activities.
`changewwn`	WWN change activities.

13.5.0.0.4 Example

```
Audit log  for        : changewwn
--------------------

User Name             : Administrator
Date and Time         : 2011-06-16T19:25:12Z
Operation             : ChangeWWNJobInfo
Host Name             : 10.192.203.179
Adapter Id            : BT11161224
Port Id               : 10:00:00:00:C9:BB:1E:77
Message               : Successfully changed WWN

New State
                      New WWPN = 10:00:00:00:C9:BB:1E:78
                      New WWNN = 20:00:00:00:C9:BB:1E:78
                      Volatile = false

Old State
                      Old WWPN = 10:00:00:00:C9:BB:1E:77
                      Old WWNN = 20:00:00:00:C9:BB:1E:77
```

Appendix A: License Notices

A.1 VI Java SDK


BROADCOM®

Emulex® Boot for the Fibre Channel Protocol

User Guide

Version 11.4
September 6, 2017

BT-FC-LPE-UG114-100

Corporate Headquarters
San Jose, CA

Website
www.broadcom.com

Chapter 1: Introduction

This guide describes installing, updating, enabling, and configuring Emulex® FC boot code. This guide also describes the Emulex boot from SAN implementation and its operation with distinct hardware and operating system requirements. Boot from SAN is booting servers directly from disk operating system images located on a SAN by way of Emulex adapters using Emulex FC boot code.

When booting from SAN, the storage device is typically identified by its WWPN and a LUN. By extending the server system boot BIOS, boot from SAN functionality is provided by the bootBIOS contained on an Emulex adapter in the server. When properly configured, the adapter then permanently directs the server to boot from a logical unit (disk) on the SAN as if it were a local disk.

NOTE Emulex drivers support multipath boot configurations. Refer to your storage vendor's documentation for details on how to configure multipath booting.

NOTE Screenshots in this guide are for illustrative purposes only. Your system information can vary.

1.1 Emulex Boot Code Files

Emulex is unique in providing Universal Boot and Pair Boot images of boot code. Universal Boot images contain x86 BootBIOS, OpenBoot, and EFIBoot boot code. Pair Boot boot code images contain x86 BootBIOS and EFIBoot boot code. These images provide multi-platform support for boot from SAN.

Universal Boot and Pair Boot transparently determine your system platform type and automatically execute the proper boot code image in the adapter. These boot code images reside in the adapter flash memory, which allows for easier adapter portability and configuration between servers.

Emulex provides the following types of boot code:

- x86 BootBIOS works with the existing system BIOS on x64 and x86 systems.
- OpenBoot works with the existing system firmware on SPARC systems. OpenBoot is also called FCode.
- UEFIBoot provides system boot capability through the use of the Unified Extensible Firmware Interface (UEFI) shell. It also functions on UEFI 2.*x*-based x64 platforms through the Human Interface Infrastructure (HII) interface.

1.2 Abbreviations

AL_PA	Arbitrated Loop Physical Address
BBS	BIOS Boot Specification
BIOS	basic input/output system
CLI	command line interface
DHCP	Dynamic Host Control Protocol
DID	device ID
DNS	Domain Name System
EDD	Enhanced Disk Device
FA-PWWN	Fabric Assigned Port Word Wide Name

FC	Fibre Channel
FC-AL	Fibre Channel Arbitrated Loop
FL_Port	fabric loop port
Gb/s	gigabits per second
GPT	GUID partition table
GUI	Graphical User Interface
GUID	Globally Unique Identifier
HBA	host bus adapter
HII	Human Interface Infrastructure
HTTP	Hypertext Transfer Protocol
I/O	input/output
JBOD	just a bunch of disks
LDAP	Lightweight Directory Access Protocol
LUN	logical unit number
MBR	master boot record
NIS/NIS+	Network Information Service/Network Information Service Plus
NVRAM	non-volatile random-access memory
OBP	OpenBoot PROM
OCSD	Option Card Sensor Data
OEM	original equipment manufacturer
PCI	Peripheral Component Interconnect
PCIe	Peripheral Component Interconnect Express
PLOGI	port login
POST	power-on self-test
RAID	redundant array of independent disks
SAN	storage area network
SCSI	Small Computer System Interface
TB	terabyte
UEFI	Unified Extensible Firmware Interface
USB	Universal Serial Bus
WWN	World Wide Name
WWPN	World Wide Port Name
ZB	zettabyte

Chapter 2: Booting from SAN

This section provides instructions for installing and using all of the types of boot code provided by Emulex.

When booting from SAN, the storage device is typically identified by its WWPN and a LUN. By extending the server system BIOS, boot from SAN capability is provided by the boot BIOS contained on an Emulex adapter in the server. If properly configured, the adapter then permanently directs the server to boot from a logical unit (disk) on the SAN as if it were a local disk.

NOTE Not all procedures are required. Emulex HBAs usually ship from the factory with the latest version of boot code installed and enabled, so you do not need to install or enable boot code in those cases. However, if boot code is not installed, you must install it, and if it is not enabled, you must enable it. Check http://www.broadcom.com for the latest version of boot code.

This section describes setting up a system to boot from SAN. The specific procedure to follow is determined by the system architecture and the operating system.

Table 1 Boot from SAN Procedures

Architecture	Operating System	Procedure
x86 and x64	Windows Server	Section 2.2, Linux, Citrix, and VMware.
	Linux	Section 2.2.1, Configuring Boot from SAN on Linux or VMware (x86 and x64/UEFI).
	VMware	
	Solaris (SFS driver)	Section 2.3.1, Configuring Boot from SAN on Solaris (x86 and x64).
SPARC and PowerPC	Linux	Section 2.2.1, Configuring Boot from SAN on Linux or VMware (x86 and x64/UEFI).
	Solaris (SFS driver)	Section 2.3.2, Configuring Boot from SAN on Solaris (SPARC).

2.1 Windows Server

This section describes configuring boot from SAN for Windows systems and installing Windows Server on an FC boot disk.

2.1.1 Configuring Boot from SAN on Windows Server (x64)

To configure boot from SAN, perform these steps:

1. If necessary, update the boot code on the adapter (see Chapter 3, Updating and Enabling Boot Code).
2. If necessary, enable the boot code on the adapter (see Chapter 3, Updating and Enabling Boot Code).
3. Enable the adapter to boot from SAN (see Section 4.2.1, Enabling an Adapter to Boot from SAN).
4. If you want to use a topology other than the default, you must change the topology setting before configuring boot devices (see Section 4.4.3, Changing the Topology).
5. Configure boot devices (see Section 4.3, Configuring Boot Devices).
6. If desired, configure the boot options on the adapter (see Section 4.4, Configuring Adapter Parameters).
7. Install the operating system on an FC boot disk (see Section 2.2, Linux, Citrix, and VMware).

2.1.1.1 Installing Windows Server on an FC Boot Disk

This installation procedure installs the Windows Server onto an unformatted FC disk drive and configures the system to boot from the SAN disk drive. The LUNs must be created in the SAN storage device and zoned appropriately to the host adapter WWN before starting the installation.

NOTE The computer system BIOS might require that another controller take precedence over the Emulex adapter during boot. If this occurs, you must disconnect or disable the other adapter. This allows you to configure and build the operating system on the drive connected to the Emulex adapter.

NOTE Restrict the number of paths (from the HBA to the boot disk) to one during installation. Additional paths can be added after the operating system is installed and the multipath/MPIO software is configured.

1. From the server system UEFI setup, ensure that the CD/DVD is the first device in the boot order list.
2. Enable the adapter BIOS setting to allow SAN boot in the Emulex x86 and x64/UEFI configuration utility.
3. Configure the boot target and LUN in the Emulex x86 and x64/UEFI configuration utility to point to the desired target (for configuring x86 and x64, see Section 4.3, Configuring Boot Devices, and for configuring UEFI, see Section 6.10, Configuring Boot Parameters).
4. Boot the host server with the Windows Server operating system DVD inserted. Follow the on-screen prompts to install the appropriate version of Windows Server.
5. The Windows installation exposes all available and visible LUNs as disks and partitions numbered 0 to N, where N is the highest number available. These numbers typically are the LUN numbers assigned by the array.
6. Select the disk on which you want to install the operating system.

   NOTE If the FC disk is not visible at this point, download the driver kit from the Broadcom website to your system. Extract/copy the driver files to a formatted disk or USB drive.

   Click **Load Driver**. Browse to the disk or USB device specified in step 1 where the driver is located to load the Storport Miniport driver for the appropriate operating system. When selected, the correct driver location and driver are displayed in the **Select driver to be installed** window. After the driver is loaded, the FC disks should be visible.
7. Select the disk and follow the system prompts in the Windows installation.
8. After the installation is complete, a Boot Option variable called Windows Boot Manager is populated with a Media Device path pointing to the Windows boot loader utility.

   The Windows Boot Manager option is inserted as the first boot option in the boot order list of the Host Server UEFI. The CD/DVD boot is the second device in the boot order list.

   Upon reboot, the system boots from the LUN set up on the SAN.
9. If multipath to the HBA is required, install or enable multipathing software (or Windows native MPIO Feature) and reboot the server.
10. After rebooting, verify the multipath configuration.

2.1.1.2 Directing a UEFI-based Server to a Windows Server Operating System Image (Installed as UEFI-Aware) Already Installed on the SAN

This installation procedure assumes a LUN exists in the SAN storage device, is zoned appropriately to the host adapter WWN, and a UEFI-aware operating system resides on the target LUN.

To direct a UEFI-based server to a Windows Server operating system image, perform these steps:

1. Enable boot from SAN in the Emulex UEFI configuration utility.

2. Configure the boot target and LUN in the Emulex UEFI configuration utility to point to the desired target (see Section 6.10, Configuring Boot Parameters).
3. Select **Boot Manager** from the System UEFI configuration manager.
4. Select **Add Boot Option**.
5. Identify the desired target in the list, and continue down the explorer path until you locate the `bootmgfw.efi` file. This file is the boot loader utility for your Windows Server UEFI-aware operating system installation.
6. Input a boot device description (such as Win2K8_UEFI_SAN) and optional data (if desired) for this device and select **Commit Changes**.
7. From the Boot Manager, select **Change Boot Order**.
8. Move the previous input description name (Win2K8_UEFI_SAN) to the desired position in the boot order.
9. Select **Commit Changes**.

   The Start Options list now reflects the boot order changes. Upon reboot, the server is able to boot from this target LUN on the SAN.

2.1.1.3 GUID Partition Table

The GPT provides a more flexible mechanism for partitioning disks than the older MBR partitioning scheme that has been common to PCs. MBR supports four primary partitions per hard drive and a maximum partition size of 2 TB. If the disk is larger than 2 TB (the maximum partition size in a legacy MBR), the size of this partition is marked as 2 TB and the rest of the disk is ignored.

The GPT disk itself can support a volume up to 2^{64} blocks in length (for 512-byte blocks, this is 9.44 ZB). The GPT disk can also theoretically support unlimited partitions.

NOTE By default, Microsoft Windows Server operating systems install with a GPT-formatted disk on a UEFI-aware server.

For more information on the GUID partition table, refer to the Microsoft website and search for the terms *Windows* and *GPT FAQ*.

2.2 Linux, Citrix, and VMware

This section describes configuring boot from SAN on Linux, Citrix, and VMware operating systems.

2.2.1 Configuring Boot from SAN on Linux or VMware (x86 and x64/UEFI)

To configure boot from SAN on Linux or VMware, perform these steps:

1. If necessary, install or update the boot code on the adapter. For x86 and x64, see Chapter 3, Updating and Enabling Boot Code, and for UEFI, see Chapter 6 Configuring Boot using UEFI HII in a UEFI 2.1 System.
2. If necessary, enable the boot code on the adapter. For x86 and x64, see Chapter 3, Updating and Enabling Boot Code, and for UEFI, see Chapter 6 Configuring Boot using UEFI HII in a UEFI 2.1 System.
3. Enable the adapter to boot from SAN. For x86 and x64, see Section 4.2.1, Enabling an Adapter to Boot from SAN, and for UEFI, see Section 6.4, Setting Boot from SAN.
4. If you want to use a topology other than the default, you must change the topology setting before configuring boot devices. For x86 and x64, see Section 4.4.3, Changing the Topology, and for UEFI, see Section 6.9.1, Changing the Topology.
5. Configure boot devices. For x86 and x64, see Section 4.3, Configuring Boot Devices, and for UEFI, see Section 6.10, Configuring Boot Parameters.

6. If desired, configure the boot options on the adapter. For x86 and x64, see Section 4.4, Configuring Adapter Parameters, and for UEFI, see Section 6.9, Configuring Adapter Parameters.
7. Use the driver on the operating system distribution disk to boot the system. If necessary, you can then update the driver to the desired version.

2.2.1.1 Installing Linux or VMware on an FC Boot Disk

The LUNs should be created in the SAN storage device and zoned appropriately to the host adapter's WWN before starting the installation.

NOTE The computer system BIOS might require that another controller take precedence over the Emulex adapter during boot. If this occurs, you must disconnect or disable the other adapter. This allows you to configure and build the operating system on the drive connected to the Emulex adapter.

Ensure that all paths (from the HBA to the boot disk) are enabled during installation.

1. From the server system UEFI setup, ensure that CD/DVD is the first device in the boot order list.
2. Enable the adapter BIOS setting to allow SAN boot in the Emulex x86 and x64/UEFI configuration utility.
3. Configure the boot target and LUN in the Emulex x86 and x64/UEFI configuration utility to point to the desired target. For configuring x86 and x64, see Section 4.3, Configuring Boot Devices, and for configuring UEFI, see Section 6.10, Configuring Boot Parameters.
4. Boot the host server with the operating system DVD inserted. Follow the on-screen prompts to install the OS.
5. The installation exposes all available and visible LUNs as disks and partitions.
6. Select the disk on which you want to install the OS.

   NOTE On Linux distributions, if the FC disk is not visible at this point, contact Broadcom Technical Support for the DUD (driver update diskette). The DUD iso image contains the drivers.

   Restart the installation process. On the installation splash screen, press **Tab** and in the boot command line, add `linux dd` option to boot the command line. This will prompt you for driver disks, at which point you must provide the DUD iso image.
7. Select the disk and follow the system prompts in the Linux or VMware installation.
8. After the installation is complete, a Boot Option variable specific to the OS installed is populated with a Media Device path pointing to the boot loader utility.

   Upon reboot, the system boots from the LUN set up on the SAN.
9. After rebooting, on Linux or VMware systems verify the multipath configuration.

2.3 Solaris

This section explains how to configure boot from SAN on the Solaris operating system.

2.3.1 Configuring Boot from SAN on Solaris (x86 and x64)

To configure boot from SAN, perform these steps:

1. If necessary, update the boot code on the adapter (see Chapter 3, Updating and Enabling Boot Code).
2. If necessary, enable the boot code on the adapter (see Chapter 3, Updating and Enabling Boot Code).
3. Enable the adapter to boot from SAN (see Section 4.2.1, Enabling an Adapter to Boot from SAN).

4. If you want to use a topology other than the default, you must change the topology setting before configuring boot devices (see Section 4.4.3, Changing the Topology).
5. Configure boot devices (see Section 4.3, Configuring Boot Devices).
6. If desired, configure the boot options on the adapter (see Section 4.4, Configuring Adapter Parameters).
7. Boot the Solaris installation CD and follow the prompts.

NOTE If you need help determining the LUNs to select for boot from SAN, see Section 2.3.1.1, Determining LUNs to Select for Boot from SAN.

2.3.1.1 Determining LUNs to Select for Boot from SAN

To determine which LUNs to select, perform these steps:

1. Open a terminal window and leave it open.
2. In the terminal window, select the LUN you are going to use as the SAN boot disk (not the local drive) using the `luxadm` probe command. This command shows all the available LUNs. Record this LUN information, which is used throughout this procedure. LUN 0 is used in the example:

```
luxadm probe

Found Fibre Channel device(s):
Node WWN:50060e8003823800 Device Type:Disk device
Logical Path:/dev/rdsk/c5t226000C0FF9833AFd6s2
Node WWN:50060e8003823800 Device Type:Disk device
Logical Path:/dev/rdsk/c5t226000C0FF9833AFd6s2
Node WWN:50060e8003823800 Device Type:Disk device
```

3. Copy the /dev/rdsk/nnn part of the path statement for a drive.
4. In the terminal window, use the luxadm display command to show the WWPN or the LUN for which you selected the path in the prior step:

```
luxadm display </dev/rdsk/nnn>
```

5. Record this LUN or WWPN information for use in the procedure.

2.3.2 Configuring Boot from SAN on Solaris (SPARC)

To configure boot from SAN, perform these steps:

1. If necessary, update the boot code on the adapter (see Chapter 3, Updating and Enabling Boot Code).
2. If necessary, enable the boot code on the adapter (see Chapter 3, Updating and Enabling Boot Code).
3. Type the following at the OBP prompt:

```
show-devs
```

The ID information for each found adapter is displayed, such as:

```
/pci@5d,700000/lpfc@1
Enable boot from SAN on each Emulex adapter in the system by typing the
following set of commands, replacing adapter_id with the ID information (as
shown above), for each Emulex adapter. There is a space between the first
quotation mark and the first character of the adapter ID.
" adapter_id" select-dev [for example, "/pci@5d,700000/lpfc@1" select-dev]
set-sfs-boot
unselect-dev
```

4. After all Emulex adapters have been enabled to boot from SAN, reboot the system with the following command:

```
reset-all
```

5. After the system reboots, boot the Solaris installation CD and follow the prompts.
6. After the installation completes successfully, you will be prompted to reboot or exit the system. Press **!** and then press **Enter** to go to the UNIX prompt.
7. After the UNIX prompt appears, append the following line to the system file at `/a/etc/system`:

   ```
   set pcie:pcie_max_mps=0
   ```
8. Save and reboot server.

2.3.3 Installing Solaris from a Network Image

The system must have a DVD drive and must be part of the site's network and naming service. If you use a naming service, the system must already be in a service, such as NIS, NIS+, DNS, or LDAP. If you do not use a naming service, you must distribute information about this system by following your site's policies.

NOTE This procedure assumes that the system is running the Volume Manager. If you are not using the Volume Manager to manage media, refer to the *Oracle Microsystems System Administration Guide: Devices and File Systems*.

To install from a network image, perform these steps:

1. Log on as a superuser or equivalent.
2. Insert the Solaris DVD in the system's drive.
3. Create a directory to contain the DVD image.

   ```
   # mkdir -p install_dir_path
   ```

   `Install_dir_path` specifies the directory where the DVD image is to be copied.
4. Change to the `Tools` directory on the mounted disk.

   ```
   # cd /cdrom/cdrom0/Solaris_10/Tools
   ```
5. For Solaris 10 only:
   a. Remove the `SUNWemlxu` and `SUNWemlxs` from the `/install_dir_path/Solaris/Tools/Boot` directory.
   b. Unzip the `lpfc` driver to a temporary directory:

      ```
      pkgadd -R/install_dir_path/Solaris/Tools/Boot -d/tmp
      ```
   c. Modify the `elxfc.conf` file to use persistent binding. For more information, refer to the *Emulex Drivers for Solaris for LightPulse Adapters User Guide*, which is available at http://www.broadcom.com.
6. Copy the DVD image in the drive to the install server's hard disk.

   ```
   # ./setup_install_server install_dir_path
   ```

   `install_dir_path` specifies the directory where the DVD image is to be copied.

   The `setup_install_server` command indicates whether you have enough disk space available for the Solaris software disk images. To determine available disk space, use the `df -kl` command.
7. Decide whether you need to make the install server available for mounting:
   — If the install server is on the same subnet as the system to be installed or you are using Dynamic Host Control Protocol (DHCP), you do not need to create a boot server. Proceed to step 8.

— If the install server is not on the same subnet as the system to be installed and you are not using DHCP, complete the following steps:

a. Verify that the path to the install server's image is shared appropriately.

```
# share | grep install_dir_path
```

install_dir_path specifies the path to the installation image where the DVD image was copied:

- If the path to the install server's directory is displayed and **anon=0** is displayed in the options, proceed to step 8.
- If the path to the install server's directory is not displayed or **anon=0** is not in the options, continue and make the install server available to the boot server. Using the `share` command, add this entry to the `/etc/dfs/dfstab` (all on one line).

```
share -F nfs -o ro,anon=0 -d "install server directory" install_dir_path
```

b. Verify that the `nfsd` daemon is running, or start the `nfsd` daemon.

- If the install server is running the current Solaris release, or a compatible version, type the following command.

```
# svcs -l svc:/network/nfs/server:default
```

If the nfsd daemon is online, continue to step c. If the nfsd daemon is not online, start it.

```
# svcadm enable svc:/network/nfs/server
```

c. Share the install server.

```
# shareall
```

8. Change directories to root (/).

```
# cd /
```

9. Eject the Solaris DVD.
10. (Optional) Patch the files that are located in the miniroot on the net install image that was created by `setup_install_server`. Patching a file might be necessary if a boot image has problems. For more information, refer to the Solaris operating system documentation.

2.3.3.1 Installing Solaris by Migrating an Image from a Local SCSI Disk

To install Solaris by migrating an image from a local SCSI disk, perform these steps:

1. Type the following at the OBP prompt:

```
show-devs
```

The ID information for each found adapter is displayed, such as:

```
/pci@5d,700000/lpfc@1 select-dev
```

2. Select the Emulex adapter on which you want to enable boot from SAN by entering the path to the adapter, for example:

```
" /pci@5d,700000/lpfc@1" select-dev
```

3. To view the current boot device ID, type:

```
show-devs
" /pci@5d,700000/lpfc@1" select-dev /* to select lpfc@1 (for example) */
.boot-id
```

Make a note of the WWPN, device ID (DID), or AL_PA returned from the probe and write down the corresponding boot entry.

4. To enable boot from the SAN, set the boot device ID to the SAN device from which you want to boot, for example:

```
" /pci@5d,700000/lpfc@1" select-dev
wwpn|did|alpa lun target_id set-boot-id
unselect-dev
```

where:

- `wwpn|did|alpa` is the device WWPN, DID, or AL_PA of the storage device.
- *`lun`* is the LUN number in hexadecimal. To enter it in decimal, enter `d# [`*`lun`*`]`.
- *`target_id`* is the target ID in hexadecimal. To enter it in decimal, enter `d# [`*`target_id`*`]`.

NOTE Use the WWPN in most cases. The DID and AL_PA might change between boots, causing the SAN boot to fail, unless the DID and AL_PA are specifically configured to not change between boots.

Example 1: alpa = e1, lun = 100 (decimal) and target id = 10 (decimal):

```
alpa e1 d# 100 d# 10 set-boot-id
```

Example 2: wwpn = 50000034987AFE, lun = af (hexadecimal) and target id = 10 (decimal):

```
wwpn 50000034987AFE af d# 10 set-boot-id
```

Example 3: did = 6312200, lun = 25 (hexadecimal) and target id = f (hexadecimal):

```
did 6312200 25 f set-boot-id
```

5. Boot to the original local disk to set up the newly defined FC disk. Type:

   ```
   boot local_disk
   ```

   where local_disk is the complete path or the alias of the original boot disk.
6. Run the format utility:

   ```
   format
   ```
7. Select the target disk to become the new boot disk (for example, `c1t1d0`).
8. Select the partition option and partition the disk as desired.
9. Select the label option and write a volume label to the target disk.

   For help with the format utility, refer to the man page, `man format`.
10. Install the boot on partition 0 of the target disk. (Type this command on one line.)

    ```
    installboot /usr/platform/ `uname -i`/lib/fs/ufs/bootblk /dev/rdsk/clt1d0s0
    ```
11. Create a file system for each partition that contains a mounted file system.

    ```
    newfs -v /dev/rdsk/clt1d0s0 (becomes root)
    newfs -v /dev/rdsk/clt1d0s6 (becomes usr)
    newfs -v /dev/rdsk/clt1d0s7 (becomes export/home)
    ```
12. Create temporary mount points for the new partitions.

    ```
    mkdir root2
    mkdir usr2
    mkdir export2
    ```
13. Mount, copy, then unmount the `usr2` file system.

    ```
    mount /dev/dsk/clt1d0s6 /usr2
    c0t0d0s6 ufsdump 0f - /dev/rdsk/c0t0d0s6 | (cd /usr2; ufsrestore rf -)
    umount /usr2
    ```
14. Copy the `export/home` file system.

    ```
    mount /dev/dsk/clt1d0s7 /export2
    ufsdump 0f - /dev/rdsk/c0t0d0s7 | (cd /export2; ufsrestore rf -)
    umount /export2
    ```
15. Perform copy.

    ```
    mount /dev/dsk/clt1d0s0 /root2
    ufsdump 0f - /dev/rdsk/c0t0d0s0 | (cd /root2; ufsrestore rf -)
    ```
16. Edit `/root2/etc/vfstab`, changing the controller number, target number, and LUN number to point to the new FC boot disk. For example, if the FC boot disk is c1t1d0, replace all local disk entries of c0t0d0 with c1t1d0.

Currently the file shows.

```
/dev/dsk/c0t0d0s1 (swap)

/dev/dsk/c0t0d0s0 and /dev/rdsk/c0t0d0s0 (root)
/dev/dsk/c0t0d0s6 and /dev/rdsk/c0t0d0s6 (usr)
/dev/dsk/c0t0d0s7 and /dev/rdsk/c0t0d0s7 (export)
```

Edit the file to show.

```
/dev/dsk/c1t1d1s1 (swap)

/dev/dsk/c1t1d0s0 and /dev/rdsk/c1t1d0s1 (root)
/dev/dsk/c1t1d0s6 and /dev/rdsk/c1t1d0s6 (usr)
/dev/dsk/c1t1d0s7 and /dev/rdsk/c1t1d0s7 (export)
```

17. Reboot the system.

```
sync
sync
halt
reset-all
```

18. Boot to disk.

```
boot disk
```

The system should boot to the FC disk.

19. View the current dump device setting.

```
dumpadm
```

20. Change the dump device to the swap area of the FC drive.

```
dumpadm -d /dev/dsk/c1t1d0s1
```

where `/dev/dsk/c1t1d0s1` is a sample path to the swap area of the FC drive.

Chapter 3: Updating and Enabling Boot Code

Use the Emulex utilities to update boot code to a newer version and enable boot code. The utility that you use depends on the operating system and, in some cases, the driver type or system architecture. Table 2 indicates the utilities you can use to install and update boot code.

Table 2 Utilities that Update and Enable Boot Code

Operating System	OneCommand® Manager GUI	OneCommand Manager CLI	UEFI Utility	CIM Provider[1]	FCA emlxadm	Offline Utilities
Windows	X	X	X			X
Linux	X	X	X			X
Solaris	X	X			X	
VMware	X	X	X	X		

1. VMware only; refer to the *Emulex CIM Provider Package Installation Guide*.

After you decide which utility to use, refer to the appropriate procedure:

- OneCommand Manager GUI: Refer to the *Emulex OneCommand Manager Application User Guide*.
- OneCommand Manager CLI: Refer to the *Emulex OneCommand Manager CLI User Guide*.
- UEFI utility: See Section 6, Configuring Boot using UEFI HII in a UEFI 2.1 System.
- FCA emlxadm utility: Refer to the *Emulex FCA Utilities for LightPulse Adapters User Guide*.
- Offline utility: Refer to the *Emulex Elxflash and LpCfg Management Utilities User Guide*.

Chapter 4: Emulex BIOS Utility

Before using the Emulex BIOS utility, ensure that the boot code is loaded and enabled on the adapter as described in Section 3, Updating and Enabling Boot Code.

NOTE This section reflects the most recent release of the BIOS utility. Some selections might not be available if you are using an older version of the utility.

NOTE Changes made to parameters common to UEFI and x86 drivers are changed in both driver configuration utilities.

NOTE After exiting the Emulex BIOS configuration utility, the system will reboot regardless of whether any changes were made.

4.1 Navigating the FC BIOS Utility

Use the following methods to navigate the BIOS utility:

- Press the up and down arrows on your keyboard to move through and select menu options or configuration fields. If multiple adapters are listed, use the up and down arrows to scroll to the additional adapters.
- Press **PageUp** to scroll to the previous page.
- Press **PageDn** to scroll to the next page.
- Press **Enter** to select a menu option, to select a changed value, to select a row in a screen, or to change a configuration default.
- Press **Esc** to go back to the previous menu.

4.2 Starting the Emulex BIOS Utility

To start the Emulex BIOS utility, perform these steps, perform these steps:

1. Turn on the computer and press and hold down **Alt** or **Ctrl** and press **E** immediately (within five seconds) when the Emulex bootup message to enter the BIOS utility is displayed.

   An adapter listing screen is displayed (Figure 1).

   NOTE Links must be connected and established before entering the BIOS utility; otherwise, you will receive an error message.

   NOTE If the bootup message does not appear, you must enable x86 BootBIOS. See Section 4.2.1, Enabling an Adapter to Boot from SAN for more information.

   NOTE After exiting the Emulex BIOS configuration utility, the system will reboot regardless of whether any changes were made.

Figure 1 Adapter Listing Screen

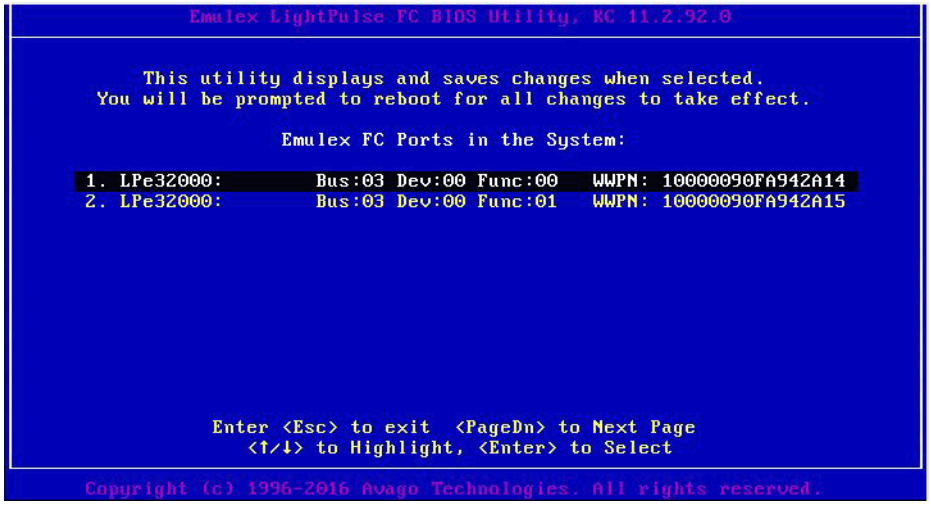

2. Select the adapter to configure and press **Enter**.
   The main configuration menu is displayed (Figure 2).

Figure 2 Main Configuration Menu

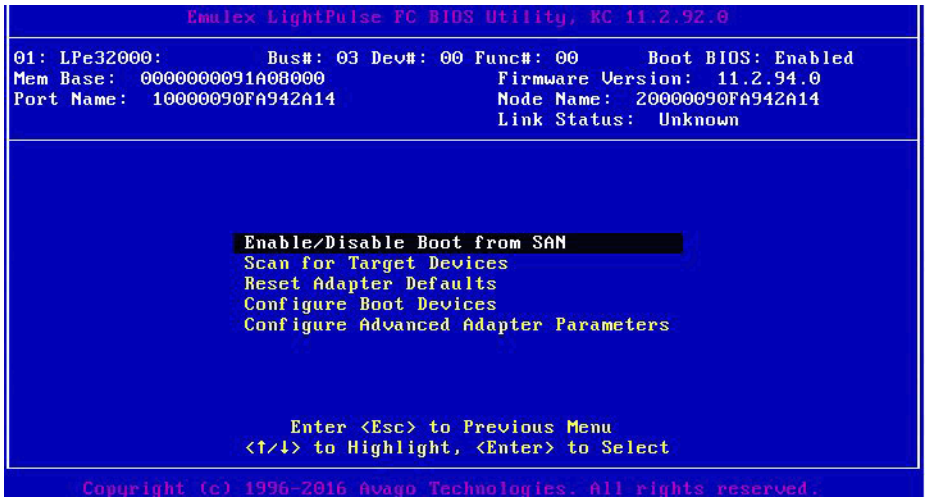

Under normal circumstances, you would first configure boot devices using the `BIOS` utility (see Section 4.3, Configuring Boot Devices). However, in the following two situations, you must perform the indicated procedure first:

- The adapter is not enabled to boot from SAN: You must enable the adapter's BIOS to boot from SAN (see Section 4.2.1, Enabling an Adapter to Boot from SAN).
- If you want to use a topology other than the default, you must change the topology setting before configuring boot devices (see Section 4.4.3, Changing the Topology).

4.2.1 Enabling an Adapter to Boot from SAN

To enable an adapter to boot from SAN, perform these steps:

1. From the main configuration menu (Figure 2), select **Enable/Disable Boot from SAN** and press **Enter**.

NOTE Adapters are disabled by default.

Boot from SAN must be enabled on one adapter to use remote boot functionality. After you enable an adapter, the status of the Boot BIOS changes as shown in Figure 3.

NOTE x86 BootBIOS supports a maximum of 16 or 32 ports depending on the adapter.

Figure 3 BIOS Status Screen

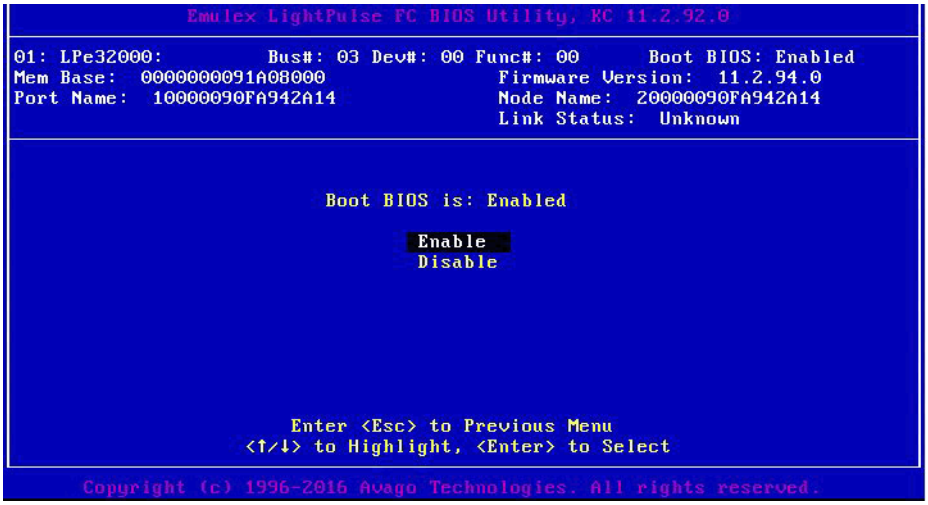

1. From the Adapter Listing screen (Figure 1), select the adapter that you want to scan for target devices and press **Enter**.
2. From the main configuration menu (Figure 2), select **Scan for Target Devices** and press **Enter**. A list of the discovered targets is displayed. This is only a list of discovered target devices to determine SAN connectivity.

4.3 Configuring Boot Devices

This option supports FC_AL: (public and private loop) and fabric point-to-point. When operating in loop (FC_AL) topology, the system automatically determines whether you are configured for a public or private loop. The BIOS looks for a fabric loop port (FL_Port) first. If a fabric loop is not detected, the BIOS looks for a private loop. For the **Configure Boot Devices** option, the eight boot entries are zero by default (D key).

NOTE If you want to use a topology other than the default, you must change the topology setting before configuring boot devices (see Section 4.4.3, Changing the Topology). For Fibre Channel Arbitrated Loop (FC-AL), each adapter has a default AL_PA of 01 (hexadecimal).

To configure boot devices, perform these steps:

1. On the main configuration menu (Figure 2), select **Configure Boot Devices**.

   A list of eight boot devices is shown (Figure 4). The primary boot device is the first entry shown, and it is the first bootable device.

   If the first boot entry fails due to a hardware error, the system can boot from the second bootable entry. If the second boot entry fails, the system boots from the third bootable entry and so on, if it is configured in the system BIOS on a BBS system.

   NOTE The Emulex `BIOS` utility only presents the boot devices to the system BIOS. The system BIOS must enumerate and attempt to boot from the drive as the primary hard drive from which to boot. See Section 4.4.13, Using Multipath Boot from SAN.

Figure 4 List of Saved Boot Devices Screen

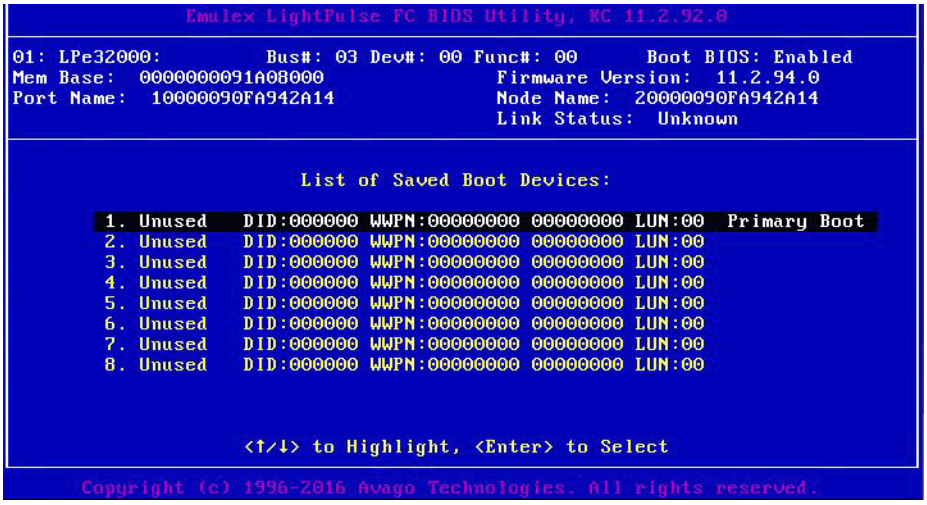

2. Select a boot entry and press **Enter**.

   A screen similar to Figure 5 is displayed.

Figure 5 Device Selection List Example (Array) Screen

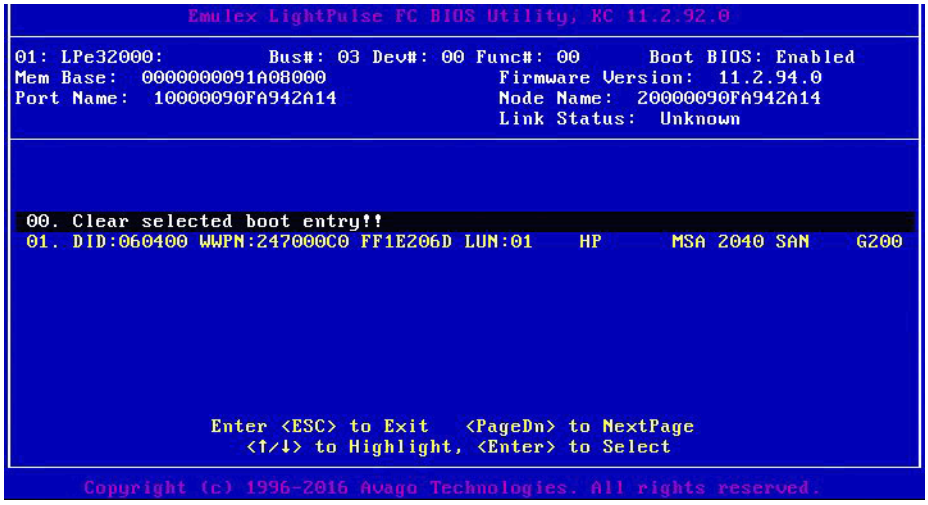

NOTE To minimize the amount of time needed to locate the boot device, select the drive with the lowest AL_PA as the boot device when connected to devices on a FC arbitrated loop.

3. Select **00** and press **Enter** to clear the selected boot entry, or select the WWPN or DID of a device to configure.

Figure 6 Starting LUN Offset Screen

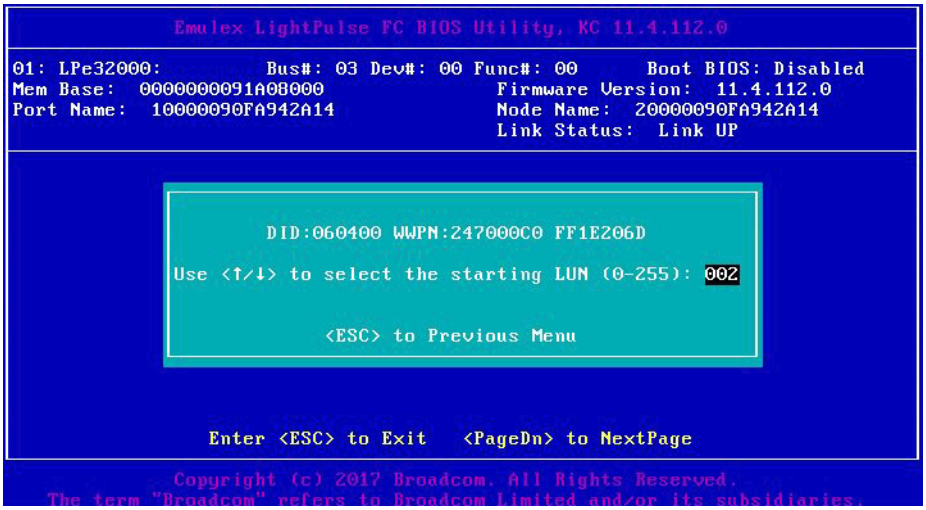

4. If you select a device, you are prompted for the starting LUN. Select the starting LUN, by using the up and down arrow keys (Figure 6). The starting LUN can be any number from 0 to 255.

NOTE The LUN value is displayed in HEX not decimal.

Figure 7 LUN Listing Screen

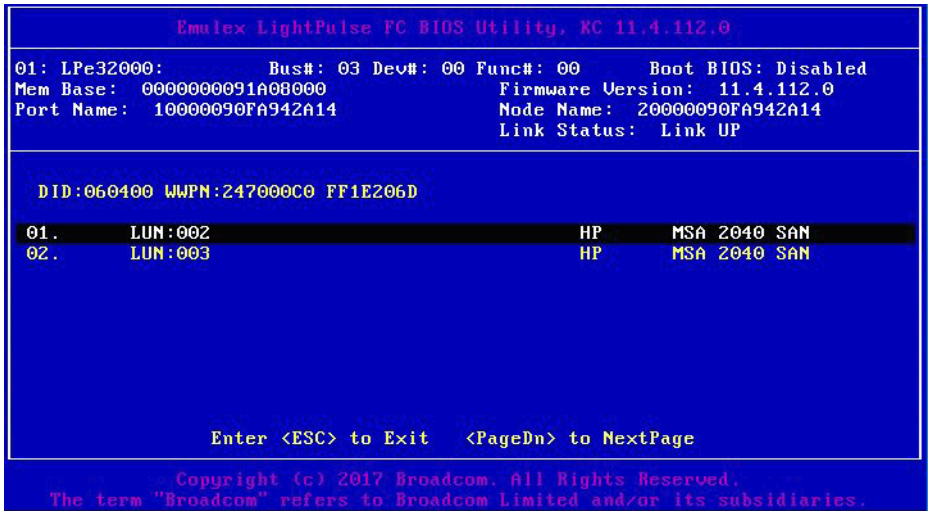

You can define 256 LUNs per adapter. The screen shows only 8 LUNs per page, from the starting LUN offset, up to 256 consecutive LUNs.

5. Use the arrow key up and down to select the LUN you want to setup as a boot device.

   The Boot Devices menu (Figure 8) is displayed.

Figure 8 Boot Devices Menu

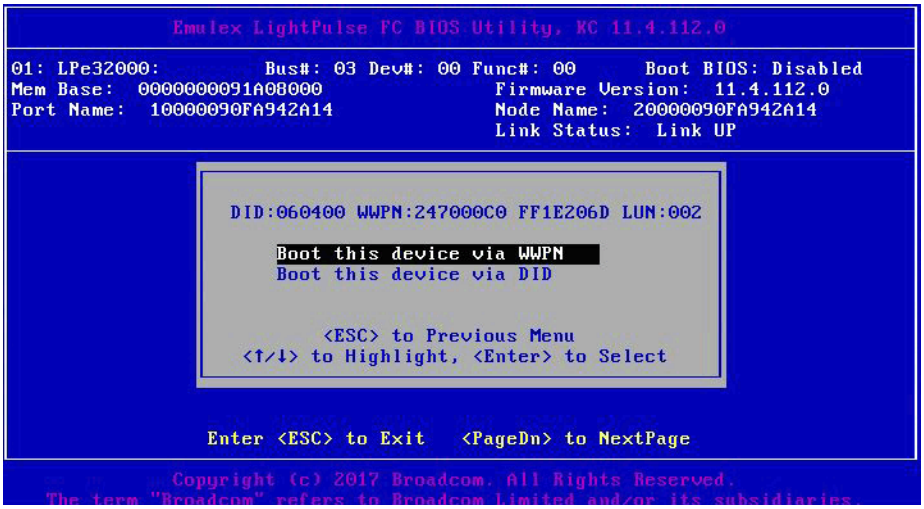

6. Use the up arrow and down arrow to select the boot method you want. If you select to boot the device by WWPN, the WWPN of the earlier selected entry is saved in the flash memory. However, during the initial BIOS scan, the utility issues a Name Server Inquiry GID_PN (Get Port Identifier), converting the WWPN to a DID. Then, based on this DID, it continues with the remainder of the scanning.

   If you select to boot this device by DID, the earlier selected entry is saved in the flash memory (Figure 9).

Figure 9 Primary Boot Device Set Up Screen

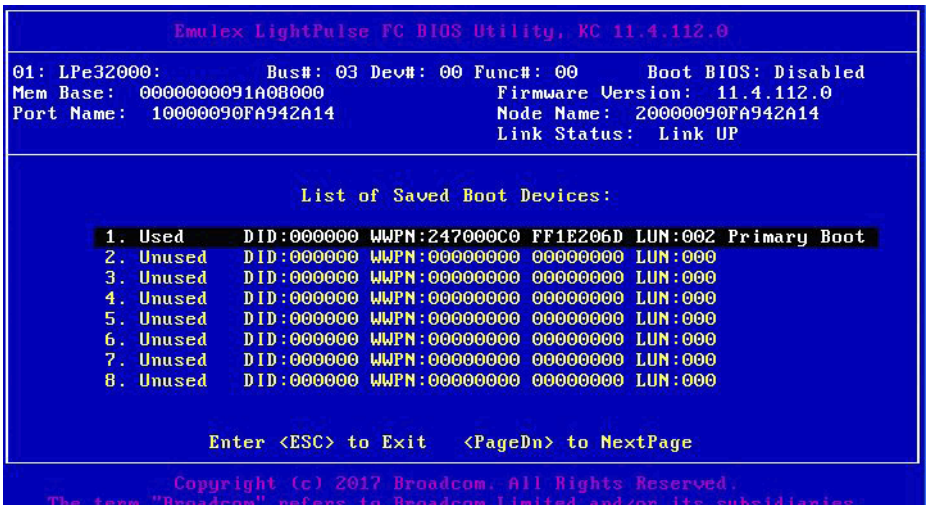

7. Press **Enter** to select the change.
8. Press **Esc** until you exit the `BIOS` utility.
9. Reboot the system for the new boot path to take effect.

4.4 Configuring Adapter Parameters

The BIOS utility has numerous options that can be modified to provide for different behavior. Use the BIOS utility to perform the following tasks:

- Change the default AL_PA of the adapter.
- Change the PLOGI retry timer.
- Select a topology.
- Enable or disable spinup delay.
- Set autoscan.
- Enable or disable EDD 3.0.
- Enable or disable the start unit command.
- Enable or disable the environment variable.
- Enable or disable the auto boot sector.
- Enable or disable Brocade FA-PWWN.
- Select a link speed.

NOTE If you want to use a topology other than the default, you must change the topology setting before configuring boot devices (see Section 4.4.3, Changing the Topology).

To access the adapter configuration menu, from the main configuration menu (Figure 2), select **Configure Advanced Adapter Parameters** and press **Enter**.

The adapter configuration menu is displayed (Figure 10).

Figure 10 Adapter Configuration Menu in LPe12000-Series, LPe16000-Series, LPe31000-Series and LPe32000-Series Adapters

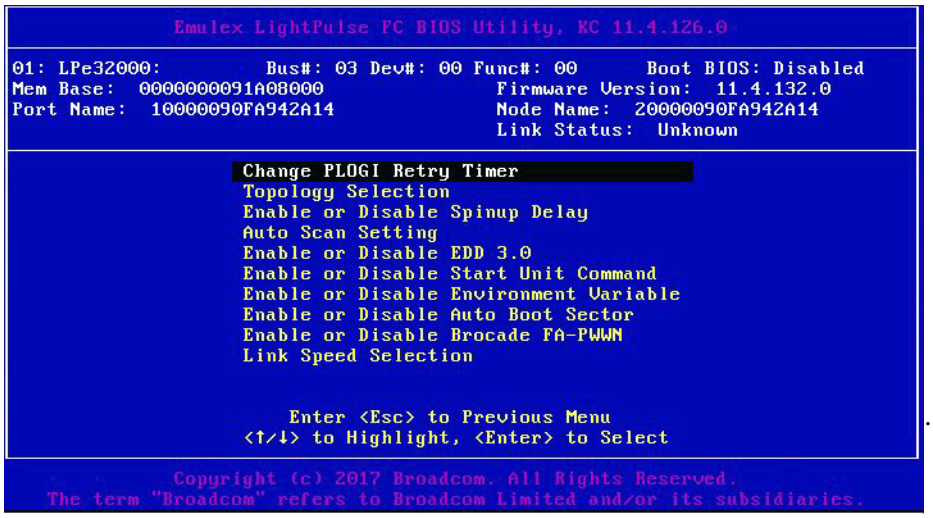

Default settings are acceptable for most installations.

To reset all values to their defaults, from the main configuration menu (Figure 2), select **Reset Adapter Defaults** and press **Enter**.

4.4.1 Changing the Default AL_PA

NOTE This option is available only on LPe12000-series and LPe16000-series adapters and is not available on LPe31000-series and LPe32000-series adapters.

The default value of the AL_PA for the adapter BIOS is 00 (hexadecimal). All adapters or boot drives can be configured to other AL_PAs rather than their default values.

NOTE This option applies only to arbitrated loop (FC-AL).

To change the default AL_PA, perform these steps:

1. From the main configuration menu (Figure 2), select **Configure Advanced Adapter Parameters**.
   The adapter configuration menu is displayed (Figure 10).
2. Select **Change Default ALPA of this Adapter** and press **Enter**.
   Information similar to Figure 11 is displayed.

Figure 11 Change Default ALPA Screen

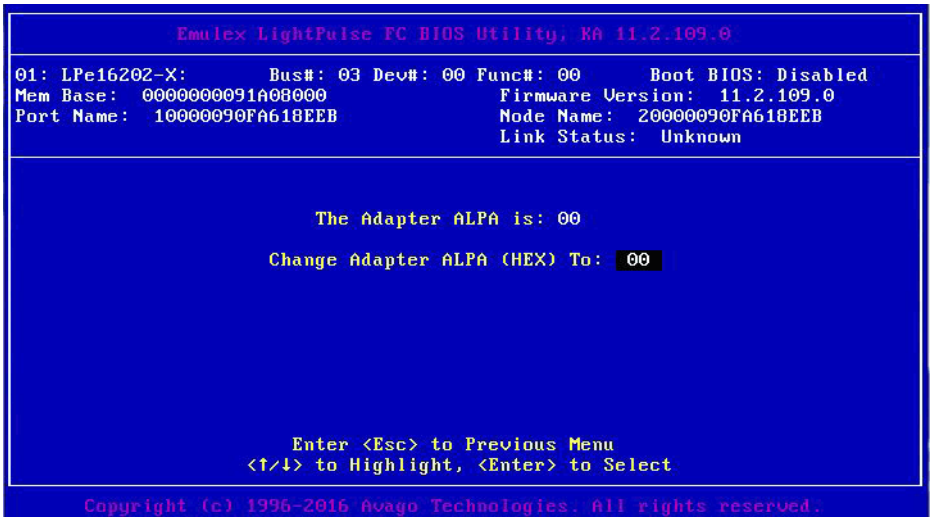

NOTE To change the default AL_PA, use the up arrow and down arrow keys to scroll through the valid ALPAs. Table 3 on page 22 lists the valid AL_PA values. If the adapter's AL_PA is changed, it does not appear on the NVRAM AL_PA until the system has been reset.

Table 3 Valid AL_PA Values

0x00	0x01	0x02	0x04	0x08	0x0F	0x10	0x17
0x18	0x1B	0x1D	0x1E	0x1F	0x23	0x25	0x26
0x27	0x29	0x2A	0x2B	0x2C	0x2D	0x2E	0x31
0x32	0x33	0x34	0x35	0x36	0x39	0x3A	0x3C
0x43	0x45	0x46	0x47	0x49	0x4A	0x4B	0x4C
0x4D	0x4E	0x51	0x52	0x53	0x54	0x55	0x56
0x59	0x5A	0x5C	0x63	0x65	0x66	0x67	0x69
0x6A	0x6B	0x6C	0x6D	0x6E	0x71	0x72	0x73
0x74	0x75	0x76	0x79	0x7A	0x7C	0x80	0x81

Table 3 Valid AL_PA Values (Continued)

0x82	0x84	0x88	0x8F	0x90	0x97	0x98	0x9B
0x9D	0x9E	0x9F	0xA3	0xA5	0xA6	0xA7	0xA9
0xAA	0xAB	0xAC	0xAD	0xAE	0xB1	0xB2	0xB3
0xB4	0xB5	0xB6	0xB9	0xBA	0xBC	0xC3	0xC5
0xC6	0xC7	0xC9	0xCA	0xCB	0xCC	0xCD	0xCE
0xD1	0xD2	0xD3	0xD4	0xD5	0xD6	0xD9	0xDA
0xDC	0xE0	0xE1	0xE2	0xE4	0xE8	0xEF	

3. Press **Enter** to accept the new value.
4. Press **Esc** until you exit the `BIOS` utility.
5. Reboot the system.

4.4.2 Changing the PLOGI Retry Timer

This option is useful for Tachyon-based RAID arrays. In rare situations, a Tachyon-based RAID array resets itself and the port goes offline temporarily. When the port returns to operation, the port login scans the loop to discover this device. The port login retry interval is the time it takes for one port login to scan the whole loop (if 126 AL_PAs are on the loop). You can choose:

- No PLOGI Retry: 0 ms – Default
- 50 ms takes 5 to 6 seconds per device
- 100 ms takes 12 seconds per device
- 200 ms takes 22 seconds per device

To set the interval for the port login retry timer, perform these steps:

1. On the main configuration menu (Figure 2), select **Configure Advanced Adapter Parameters** and press **Enter**. The adapter configuration menu is displayed (Figure 10).
2. Select **Change PLOGI Retry Timer** and press **Enter**. Information similar to Figure 12 is displayed.

Figure 12 Change the PLOGI Retry Timer Screen

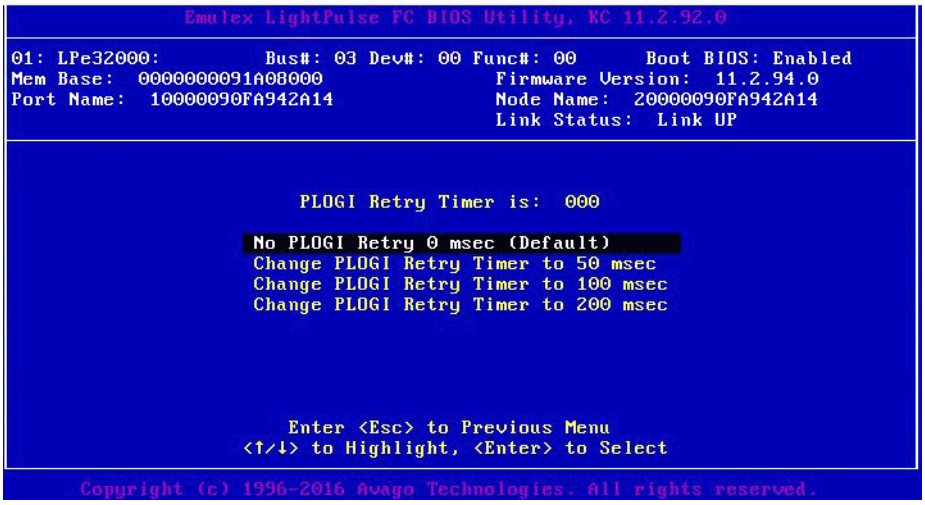

3. Select the retry timer interval.

4. Press **Enter** to accept the new interval.
5. Press **Esc** until you exit the BIOS utility.
6. Reboot the system.

4.4.3 Changing the Topology

NOTE Point to point is the only topology supported on LPe31000-series and LPe32000-series adapters.

By default, the LPe12000-series and LPe16000-series adapters use auto topology with loop first. If you want to use a topology other than the default, you must change the topology setting before configuring boot devices. For FC-AL, each adapter has a default AL_PA of 01 (hexadecimal).

The topology options available in LPe12000-series and LPe16000-series adapters are shown in Figure 13 .

You can configure these items:

- Auto topology with loop first – default
- Auto topology with point-to-point first
- FC-AL
- Fabric point-to-point

To select the adapter topology, perform these steps:

1. On the main configuration menu (Figure 2), select **Configure Advanced Adapter Parameters** and press **Enter**. The adapter configuration menu is displayed (Figure 10).
2. Select **Topology Selection** and press **Enter**. Information similar to Figure 13 is displayed.

Figure 13 Topology Menu in LPe12000-Series and LPe16000-Series Adapters

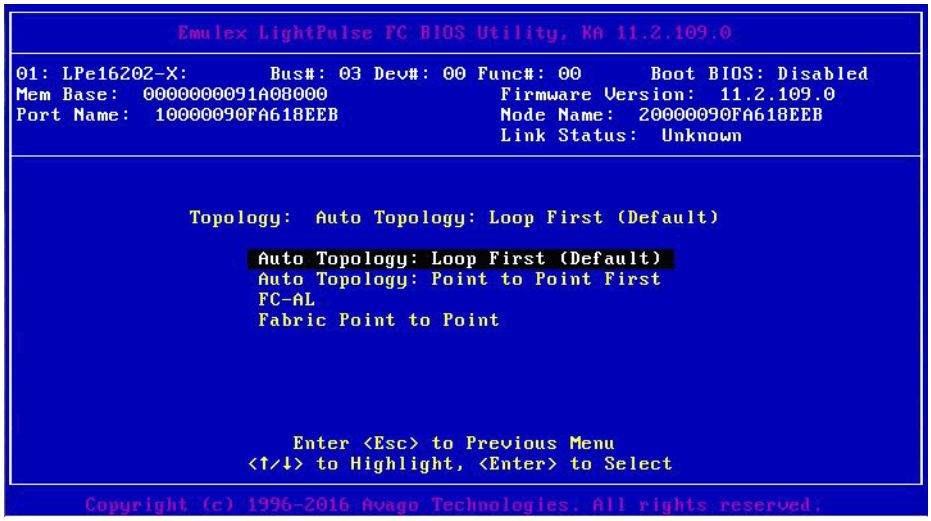

3. Select the topology for the adapter.
4. Press **Enter** to accept the new topology.
5. Press **Esc** until you exit the BIOS utility.
6. Reboot the system.

The LPe31000-series and LPe32000-series adapters use point-to-point topology and cannot be changed. The topology options available in LPe31000-series and LPe32000-series adapters is shown in Figure 14.

Figure 14 Topology Menu in LPe31000-Series and LPe32000-Series Adapters

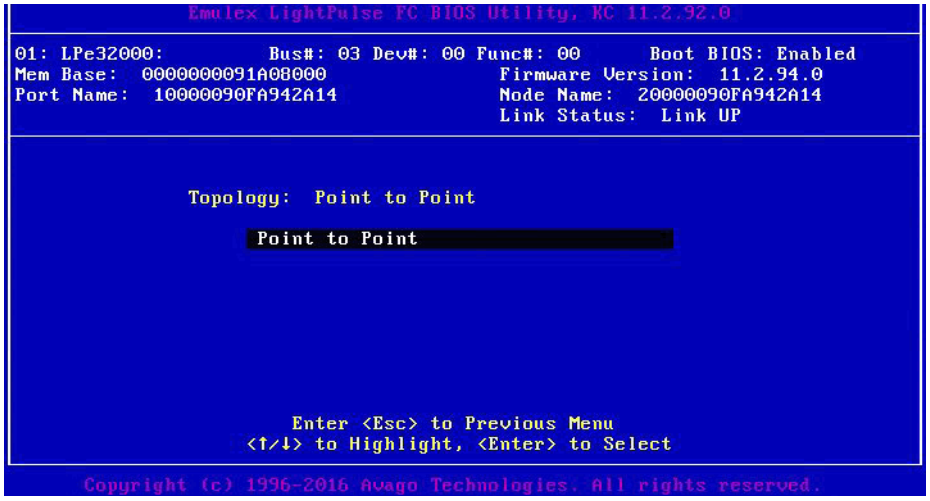

4.4.4 Enabling or Disabling the Spinup Delay

This option allows you to enable or disable the disk spinup delay. The factory default setting is disabled.

If at least one boot device has been defined, and the spinup delay is enabled, the BIOS searches for the first available boot device.

- If a boot device is present, the BIOS boots from it immediately.
- If a boot device is not ready, the BIOS waits for the spinup delay and, for up to three additional minutes, continues the boot scanning algorithm to find another multi-boot device.

If boot devices have not been defined and auto scan is enabled, the BIOS waits for five minutes before scanning for devices.

- In a private loop, the BIOS attempts to boot from the lowest target AL_PA it finds.
- In an attached fabric, the BIOS attempts to boot from the first target found in the NameServer data.

To enable or disable the spinup delay, perform these steps:

1. On the main configuration menu (Figure 2), select **Configure Advanced Adapter Parameters** and press **Enter**.
   The adapter configuration menu is displayed (Figure 10).
2. Select **Enable or Disable Spinup Delay** and press **Enter**.
   The Enable or Disable Spinup Delay screen (Figure 15) is displayed.

Figure 15 Enable or Disable Spinup Delay Screen

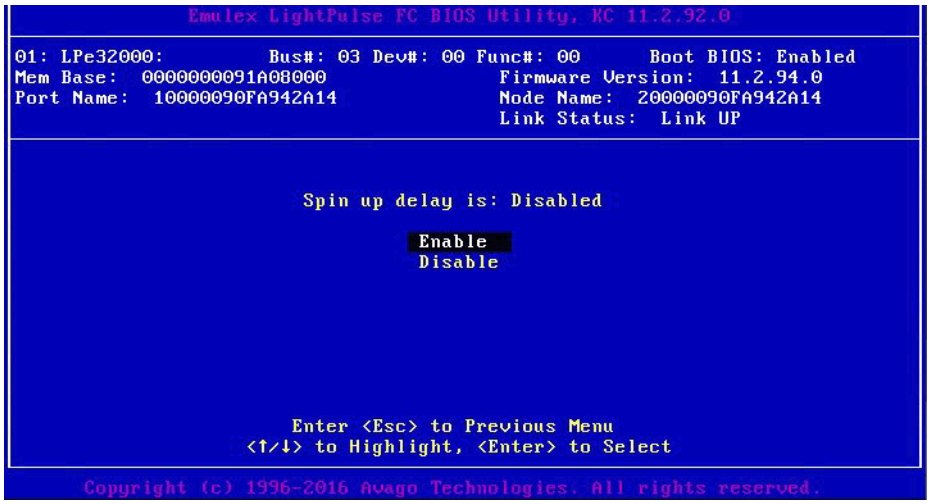

3. Select whether to enable or disable spinup delay.
4. Press **Enter** to accept the new value.
5. Press **Esc** until you exit the `BIOS` utility.
6. Reboot the system.

4.4.5 Setting Auto Scan

This option allows you to set auto scan and enable the first device in the boot entry list to issue a Name Server Inquiry. Auto scan is available only if none of the eight boot entries are configured to boot from DID or WWPN. The factory default is disabled. If there is more than one adapter with the same PCI Bus number exists in the system, and each has a boot drive attached, the first PCI-scanned adapter is the boot adapter.

Use the Boot Devices menu (Figure 8) to configure up to eight boot entries for fabric point-to-point, public loop, or private loop configurations. The first adapter is usually in the lowest PCI slot in the system. This device is the only boot device and it is the only device exported to the multi-boot menu.

The following auto scan options are available:

- Autoscan disabled – Default.
- Any first device – The first adapter issues a Name Server Inquiry and the first D_ID from the inquiry becomes the boot device. The adapter attempts to log in to a public loop first. If it fails, it logs in to a private loop. The first successfully scanned device becomes the boot device. Only this device is exported to the multi-boot menu.
- First LUN 0 device.
- First NOT LUN 0 device (a device other than LUN 0).

To set auto scan, perform these steps:

1. From the main configuration menu (Figure 2), select **Configure Advanced Adapter Parameters** and press **Enter**.
   The adapter configuration menu is displayed (Figure 10).
2. Select **Auto Scan Setting** and press **Enter**.
   Figure 16 is displayed.

Figure 16 Set Auto Scan Menu

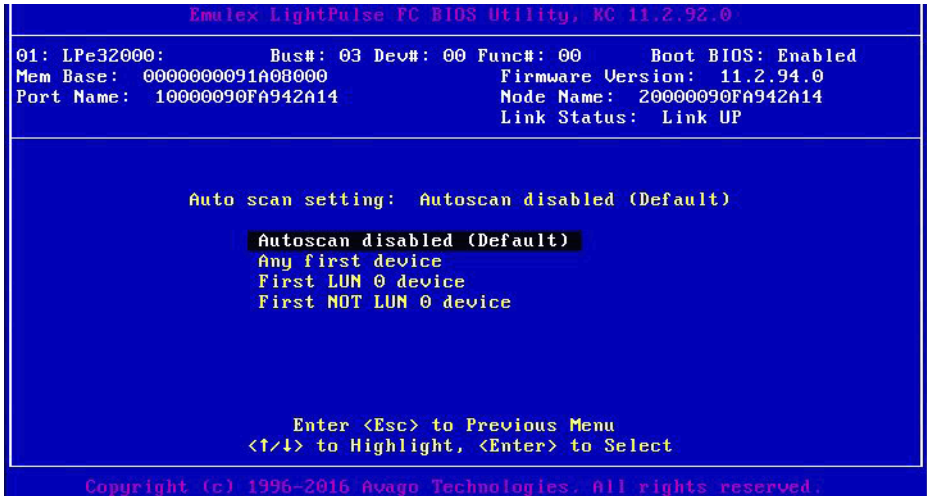

3. Select the desired auto scan option.
4. Press **Enter** to accept the new option.
5. Press **Esc** until you exit the `BIOS` utility.
6. Reboot the system.

4.4.6 Enabling or Disabling EDD 3.0

EDD 3.0 provides additional data to the operating system boot loader during the INT-13h function 48h (get device parameters) call. This information includes the path to the boot device and disk size. The default setting for EDD 3.0 is enabled.

NOTE If EDD 3.0 is disabled, then EDD 2.1 is enabled.

To enable or disable EDD 3.0, perform these steps:

1. From the main configuration menu (Figure 2), select **Configure Advanced Adapter Parameters** and press **Enter**.
   The adapter configuration menu is displayed (Figure 10).
2. Select **Enable or Disable EDD 3.0**.
   The Enable or Disable EDD 3.0 screen (Figure 17) is displayed.

Figure 17 Enable or Disable EDD 3.0 Screen

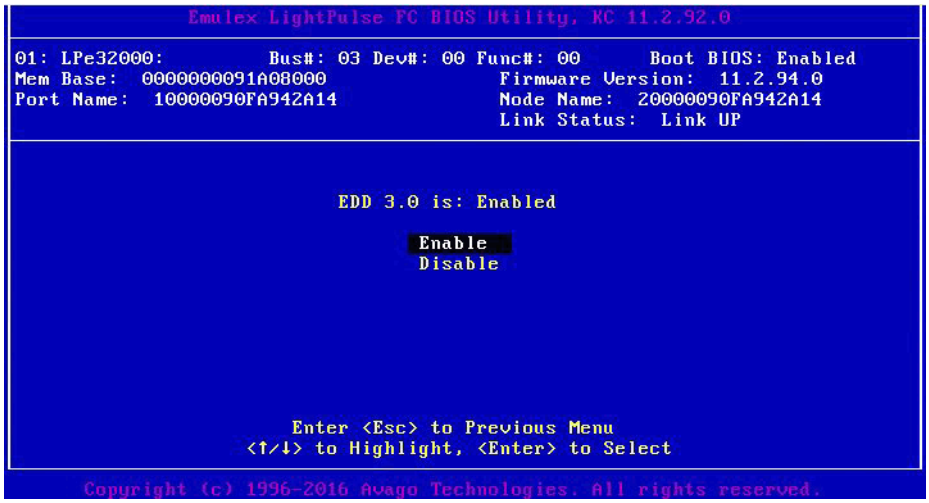

3. Select the desired EDD 3.0 setting.
4. Press **Enter** to accept the new setting.
5. Press **Esc** until you exit the `BIOS` utility.
6. Reboot the system.

4.4.7 Enabling or Disabling the Start Unit Command

This option allows a start unit command to be sent to a particular disk. This option might be needed if the disk you want to boot from is not yet spun up or started. You must know the specific LUN to issue the SCSI start unit command. The default setting is disabled.

To enable or disable the start unit command, perform these steps:

1. From the main configuration menu (Figure 2), select **Configure Advanced Adapter Parameters** and press **Enter**.
   The adapter configuration menu is displayed (Figure 10).
2. Select **Enable or Disable Start Unit Command** and press **Enter**.
   The Enable or Disable Start Unit Command screen (Figure 18) is displayed.

Figure 18 Enable or Disable Start Unit Command Screen

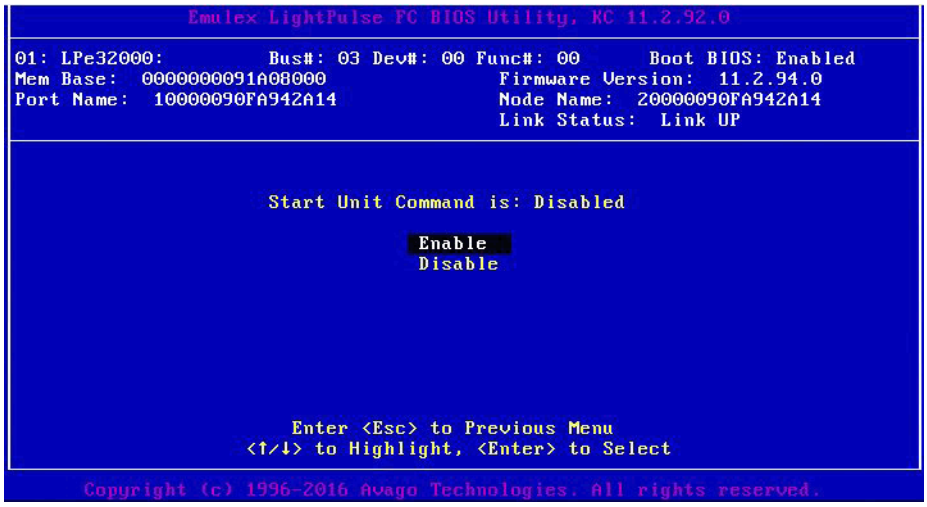

3. Select the desired Start Unit Command setting.
4. Press **Enter** to accept the new setting.
5. Press **Esc** until you exit the `BIOS` utility.
6. Reboot the system.

4.4.8 Enabling or Disabling the Environment Variable

This option sets the boot controller order if the system supports the environment variable. The default setting is disabled.

To enable or disable the environment variable, perform these steps:

1. From the main configuration menu (Figure 2), select **Configure Advanced Adapter Parameters** and press **Enter**.
   The adapter configuration menu is displayed (Figure 10).
2. Select **Enable or Disable** and press **Enter**.
   The Enable or Disable Environment Variable screen (Figure 19) is displayed.

Figure 19 Enable or Disable Environment Variable Screen

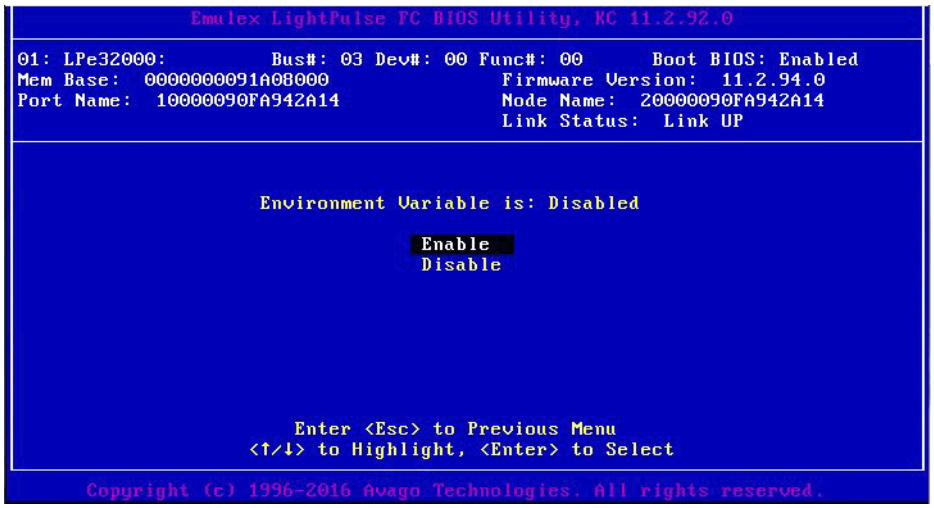

3. Select the desired setting.
4. Press **Enter** to accept the new setting.
5. Press **Esc** until you exit the BIOS utility.
6. Reboot the system.

4.4.9 Enabling or Disabling Auto Boot Sector

This option automatically defines the boot sector of the target disk for the migration boot process, which applies only to HP MSA1000 arrays. If there is no partition on the target, the default boot sector format is 63 sectors. The default setting is disabled.

To enable or disable auto sector format select, perform these steps:

1. From the main configuration menu (Figure 2), select **Configure Advanced Adapter Parameters** and press **Enter**.
   The adapter configuration menu is displayed (Figure 10).
2. Select **Enable or Disable Auto Boot Sector** and press **Enter**.
   The Enable or Disable Auto Boot Sector Format Select screen (Figure 20) is displayed.

Figure 20 Enable or Disable Auto Boot Sector Format Select Screen

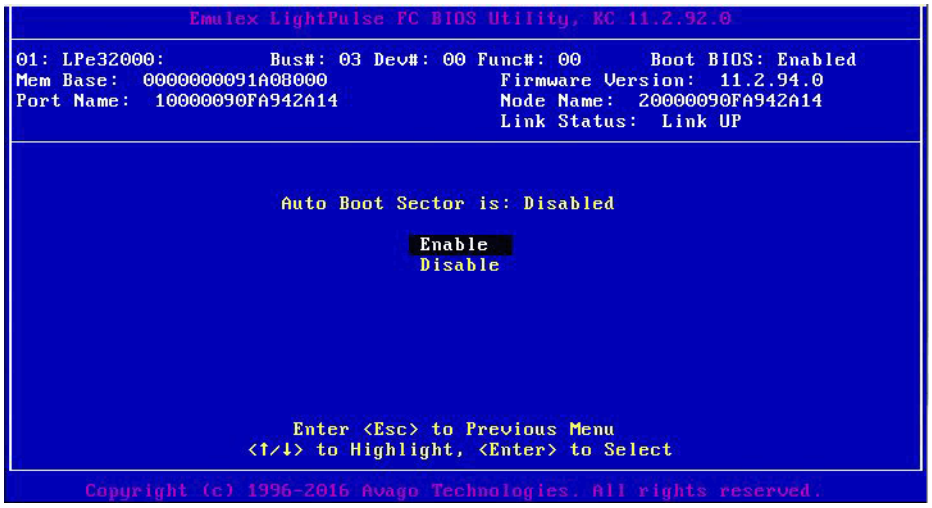

3. Select the desired setting.
4. Press **Enter** to accept the new setting.
5. Press **Esc** until you exit the BIOS utility.
6. Reboot the system.

4.4.10 Enabling or Disabling Brocade FA-PWWN

This option enables or disables the Fabric Assigned Port Word Wide Name (FA-PWWN). The default setting is disabled.

To enable or disable FA-PWWN select, perform these steps:

1. From the main configuration menu (Figure 2), select **Configure Advanced Adapter Parameters** and press **Enter**.
   The adapter configuration menu is displayed (Figure 10).
2. Select **Enable or Disable Brocade FA-PWWN** and press **Enter**.
   The Enable or Disable Brocade FA-PWWN screen (Figure 20) is displayed.

Figure 21 Enable or Disable Brocade FA-PWWN Screen

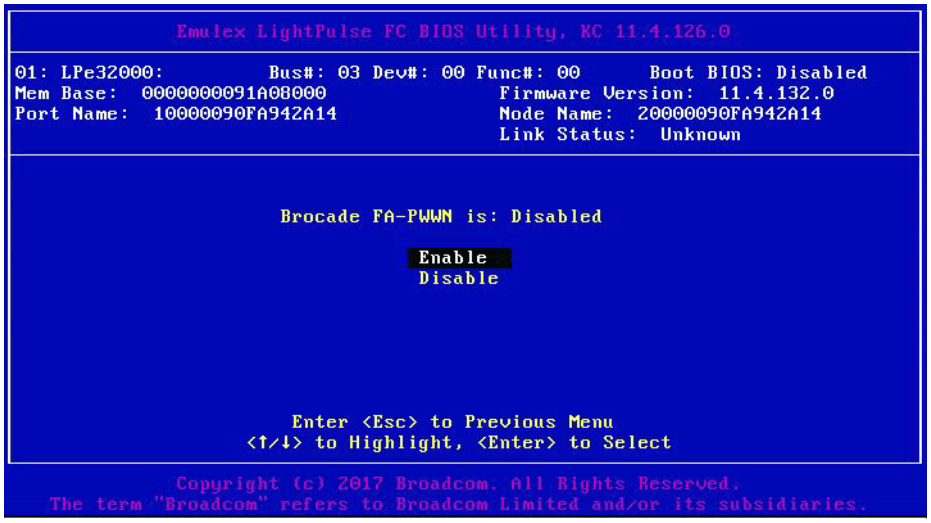

3. Select the desired setting.
4. Press **Enter** to accept the new setting.
5. Press **Esc** until you exit the `BIOS` utility.
6. Reboot the system.

4.4.11 Changing the Link Speed

NOTE Supported link speeds are specific to the adapter model and the optics installed. Examples of the LPe12000-series andLPe16000-series adapters are shown in Figure 22. Examples of LPe31000-series and LPe32000-series adapters are shown in Figure 23.

The default link speed is Auto Select, which automatically selects the link speed based on the adapter model.

NOTE In LPe12000-series adapters, the Link Speed setting will be overwritten by the FC driver and will not take effect during a regular boot into the local operating system. Therefore, the link speed BIOS setting will work for the actual boot, but will be overwritten once the operating system loads.

Possible link speeds (depending upon your adapter model):

- Auto Select – default
- 1Gb/s (if available)
- 2Gb/s (if available)
- 4Gb/s (if available)
- 8Gb/s (if available)
- 16Gb/s (if available)
- 32Gb/s (if available)

NOTE Some older devices do not support auto link speed detection; therefore, the link speed must be forced to match the speed of those devices.

To change the adapter's link speed, perform these steps:

1. From the main configuration menu (Figure 2), select **Configure Advanced Adapter Parameters** and press **Enter**.
   The adapter configuration menu is displayed (Figure 10).
2. Select **Link Speed Selection** and press **Enter**.
   The Link Speed Selection menu (Figure 22) is displayed.

Figure 22 Link Speed Selection Menu in LPe16000-Series Adapters

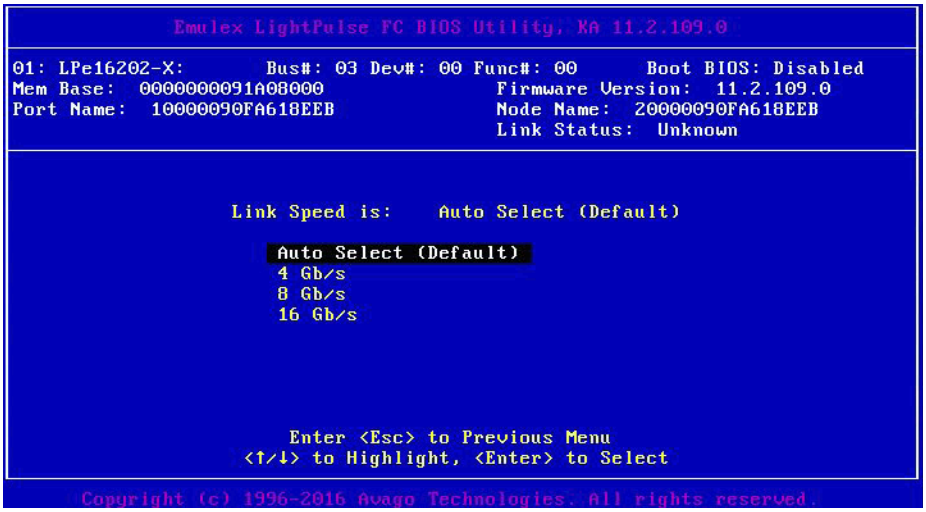

Figure 23 Link Speed Selection Menu in LPe32000-Series Adapters

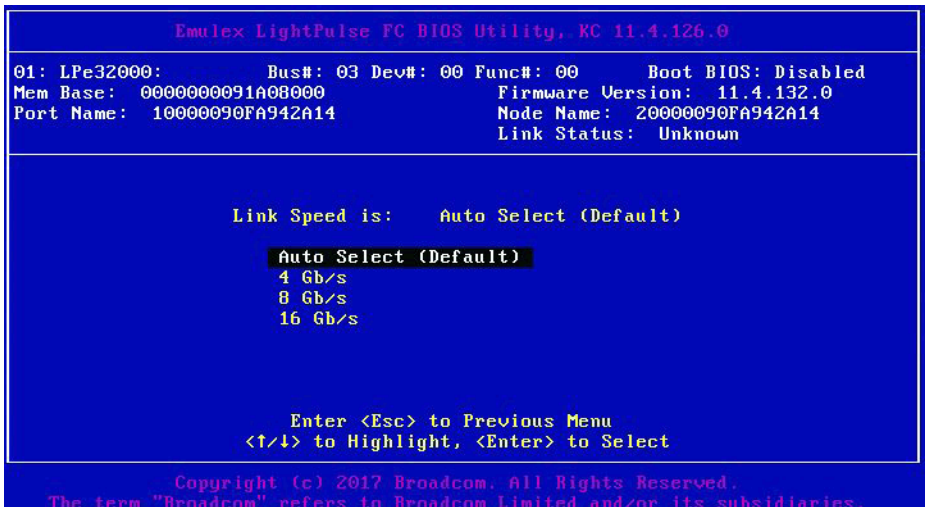

3. Use the up arrow and down arrow to select the desired FC link speed you would like to use.
4. Press **Enter** to accept the new link speed.
5. Press **Esc** until you exit the `BIOS` utility.
6. Reboot the system.

4.4.12 Resetting to Default Values

NOTE Resetting the adapter to default setting will clear all entries made while configuring the boot device, as detailed in Section 4.3, Configuring Boot Devices).

The BIOS utility enables you to reset BIOS boot parameters to their factory default settings. These defaults are listed in Table 4 on page 33.

Table 4 Adapter Defaults

Parameter	Default	Valid Values
Boot from SAN	Disabled	Enabled Disabled
ALPA Value	0x00 Fibre	See the ALPA reference table (Section 4.4.1, Changing the Default AL_PA.
EDD 3.0	EDD 3.0	Disabled (EDD 2.1) Enabled (EDD 3.0)
PLOGI Retry Timer	Disabled	Disabled 50 ms 100 ms 200 ms
Spinup Delay	Disabled	Enabled Disabled
Auto Scan	Disabled	Enabled Disabled
Start Unit	Disabled	Enabled Disabled
Environmental Variable	Disabled	Enabled Disabled
Auto Boot Sector	Disabled	Enabled Disabled
Topology	Auto Topology: Loop First (Default) **NOTE** On LPe31000-series and LPe32000-series adapters, only Point-to-Point is available.	Auto Topology: Loop First (Default) Auto Topology: Point-to-Point First FC-AL Point-to-Point
Link Speed	Auto (highest successfully negotiated speed)	

To reset parameters to their factory default settings, perform these steps:

1. On the main configuration menu (Figure 2) select **Reset Adapter Defaults** and press **Enter**.
   A screen similar to Figure 24 is displayed asking if you want to reset to the default settings.

Figure 24 Reset Adapter Configuration to Defaults

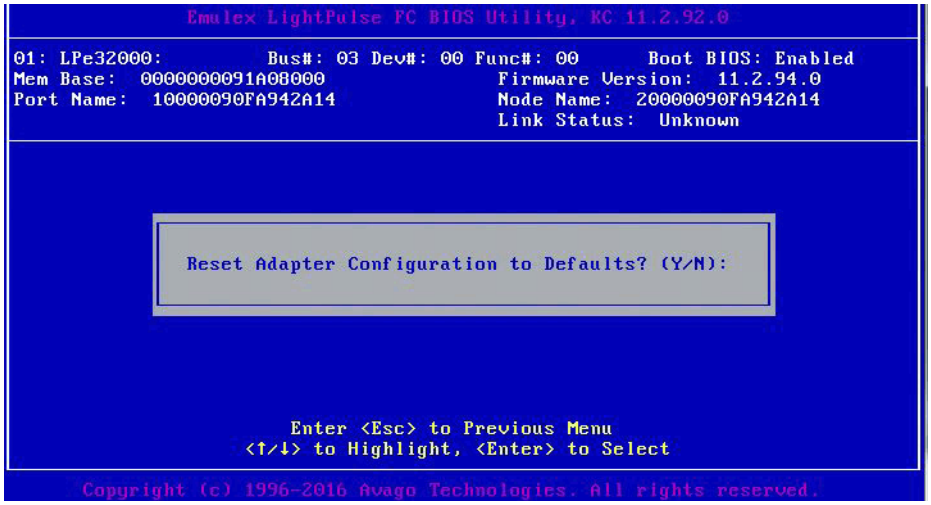

2. Press **Y**.

   All settings revert to their factory default values.

4.4.13 Using Multipath Boot from SAN

Multi-boot BIOS is in compliance with the BIOS Boot Specification (BBS). The system must have a Multi-boot system BIOS to use this feature. Multi-boot system BIOS allows you to select any boot disk in the system BIOS setup menu. The boot disk can be an FC drive, a SCSI drive, an IDE drive, a USB device, or drive. The Emulex BIOS supplies the first eight drives to the system BIOS menu. The Multi-boot system BIOS can override the FC drive that is selected in the `BIOS` utility.

For example, the system has only eight FC disks. The boot disk has ALPA 02. However, you can select ALPA 23 in the system BIOS setup menu. The boot device is the FC disk with ALPA 23 instead of ALPA 02, as is set in the `BIOS` utility.

If your system supports Multi-boot BBS, the local boot disk (drive C) is the first entry in Multi-boot on the system BIOS setup menu. The list of entries is determined by the list of configured boot entries in the `BIOS` utility. For example:

```
Adapter 1: boot_entry0, boot_entry1
Adapter 2: boot_entry2, boot_entry3
```

The order of boot entries exported to Multi-boot (BBS) is:

```
boot_entry0, boot_entry1, boot_entry2, and boot_entry3.
```

However, Multi-boot allows changing the boot order in the server BIOS, which allows any disk to become the C drive.

Chapter 5: OpenBoot

OpenBoot commands are supported on SPARC systems.

5.1 Attribute Commands

The following commands show the boot device ID, boot list, OpenBoot versions, and so on.

5.1.1 .boot-id

Syntax	`.boot-id`
Description	Shows the current boot device ID.
Parameters	None

5.1.2 devalias

Syntax	`.devalias`
Description	Shows the boot list.
Parameters	None

5.1.3 .fcode

Syntax	`.fcode`
Description	Shows the current version of OpenBoot.
Parameters	None

5.1.4 host-did

Syntax	`.host-did`
Description	Shows the actual current AL_PA of the adapter.
Parameters	None

5.1.5 .nvram

Syntax	`.nvram`
Description	Shows the current flags for OpenBoot. `show-devs "/pci@1f,0/pci@1/lpfc@1" select-dev/* select lpfc@1` (for example) `*/.nvram`
Parameters	None

5.1.6 probe-scsi-all

Syntax	`.probe-scsi-all`
Description	Shows the current SCSI and FC devices.
Parameters	None

5.1.7 show-devs

Syntax	`.show-devs`
Description	Shows a list of the devices found.
Parameters	None

5.1.8 .topology

Syntax	`.topology`
Description	Shows the current topology.
Parameters	None

5.2 Functional Commands

The following commands remove boot IDs, resets default values, sets link speed, and so on.

5.2.1 .remove-boot-id

Syntax	`.remove-boot-id`

Description	Removes the boot ID from the boot list.
	`"/pci@1f,0/pci@1/lpfc@1" select-dev remove-boot-id /* to clear boot id settings */unselect-dev`
	or
	`"/pci@1f,0/pci@1/lpfc@1" select-dev set-default-mode /* to clear boot id settings */unselect-dev`

Parameters	None

5.2.2 set-default-mode

Syntax	`set-default-mode`
Description	Resets to the default value mode.
Parameters	None

5.2.3 set-link-speed

NOTE In LPe12000-series adapters, the Link Speed setting will be overwritten by the FC driver and will not take effect during a regular boot into the local operating system. Therefore, the link speed BIOS setting will work for the actual boot, but will be overwritten after the operating system loads.

Syntax `set-link-speed`

Description Shows the current link-speed setting. Changes and sets the link speed. The default is `0` = Auto Select Link Speed.

Parameters
`0` = Auto Select Link Speed (Default)
`1` = 1Gb/s Link Speed – Only
`2` = 2Gb/s Link Speed – Only
`4` = 4Gb/s Link Speed – Only
`8` = 8Gb/s Link Speed – Only
`16` = 16Gb/s Link Speed – Only
`32` = 32Gb/s Link Speed – Only

5.2.4 set-max-lun

Syntax `set-max-lun`

Description Shows the current maximum LUN support. Changes and sets the support setting, the default maximum LUN is 256.

Parameters
0 = Set Max LUN to 255 (Default)
1 = Set Max LUN to 1023
2 = Set Max LUN to 2047
3 = Set Max LUN to 4095

5.2.5 set-post-linkup

Syntax `set-post-linkup`

Description Controls if a linkup is to occur during a POST. The default is off.

Parameters
0 = Set Linkup Switch OFF (Default)
1 = Set Linkup Switch ON

5.2.6 OpenBoot Signature

Target ID can be bound to either DID (destination ID) or WWPN and saved in an adapter NVRAM. It can also be saved in an environmental variable boot-device.

Table 5 OpenBoot Signature Table

Signature	OpenBoot Signature
`Valid_flag`	Internal flag for OpenBoot
`Host_did`	Shows host DID number
`Enable_flag`	Internal flag for OpenBoot
`Topology_flag`	Topology flag for OpenBoot

Table 5 OpenBoot Signature Table (Continued)

Signature	OpenBoot Signature
Link_Speed_Flag	Set link speed
Diag_Switch	Set fcode diag switch
Boot_id	Shows target ID number
Lnk_timer	Internal use for OpenBoot
Plogi_timer	Internal use for OpenBoot
LUN	Shows boot LUN in use
DID	Shows boot ID in use
WWPN	Shows boot WWPN in use

Chapter 6: Configuring Boot using UEFI HII in a UEFI 2.1 System

You can configure boot functions using the Emulex Configuration utility. If you have several adapters in your system, the UEFI system firmware or boot code uses the highest version driver that is on one of your adapters.

NOTE Adapters with older versions of EFIBoot are managed by the more recent version, but only as long as the adapter with the most recent version is in the system. The adapters must be updated to actually update and not just use the most recent version available. EFIBoot is not supported on legacy CNAs, such as the LP21000 and LP21002.

NOTE Dell is not supporting UEFI boot on LPe12000-series adapters.

6.1 Updating an Adapter's Firmware or Boot Code

The Emulex UEFI Boot driver supports firmware updates by implementing the UEFI Firmware Management Protocol.

NOTE If a secure version of firmware (version 11.0 or later) is installed on an LPe31000-series or LPe32000-series adapter and you want to update to an earlier unsecure version of firmware, you must remove the secure firmware jumper block before performing the update. Refer to the adapter installation guide for more information.

6.2 Starting the Emulex Configuration Utility

Depending on the OEM UEFI configuration, the Emulex Configuration utility might appear under different setup menus in the OEM system firmware or BIOS.

To start the Emulex Configuration utility using HII, perform these steps:

1. Exit the UEFI shell (Figure 25).

Figure 25 Exiting the UEFI Shell

A front page is displayed.

NOTE Depending on the vendor UEFI configuration, the menu item on the front page that allows you to configure devices can be called Storage, Device Manager, Device Settings, or something else; refer to the documentation that accompanied the server.

2. Select the menu item that allows you to configure devices. On the screen that appears (Figure 26), a list of all the adapters in the system is displayed. Your list might vary depending on the installed adapters.

Figure 26 Device Configuration Screen Example

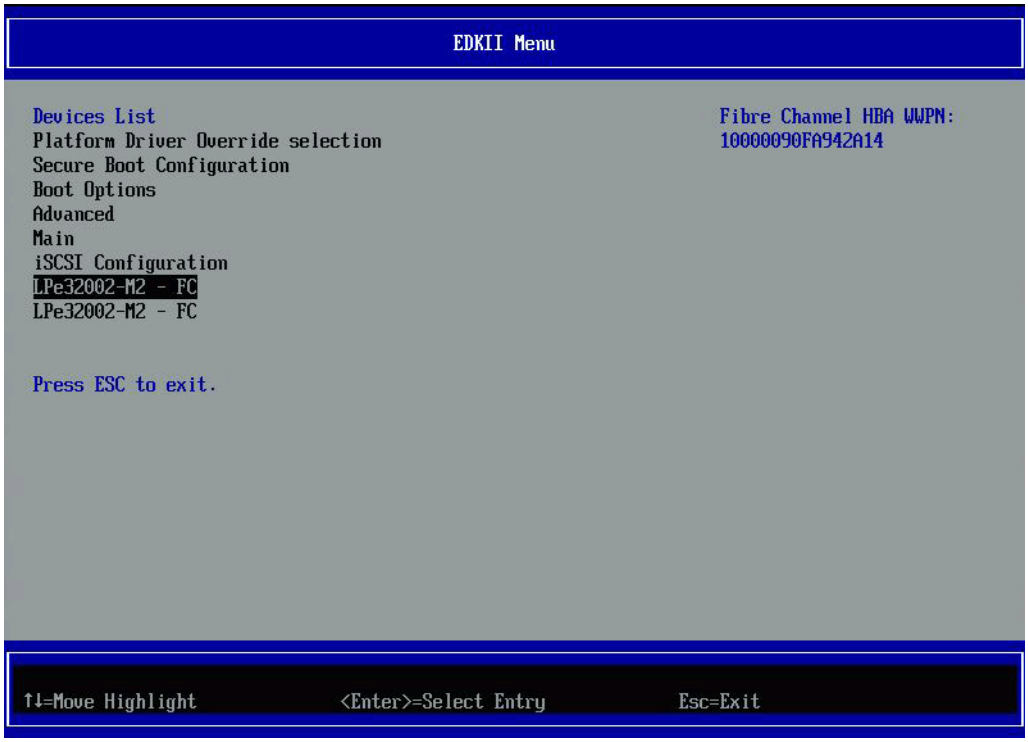

3. Select the adapter you want to configure and press **Enter**.
   The Emulex Adapter Configuration main menu screen (Figure 27) is displayed.

Figure 27 Emulex Adapter Configuration Main Menu Screen

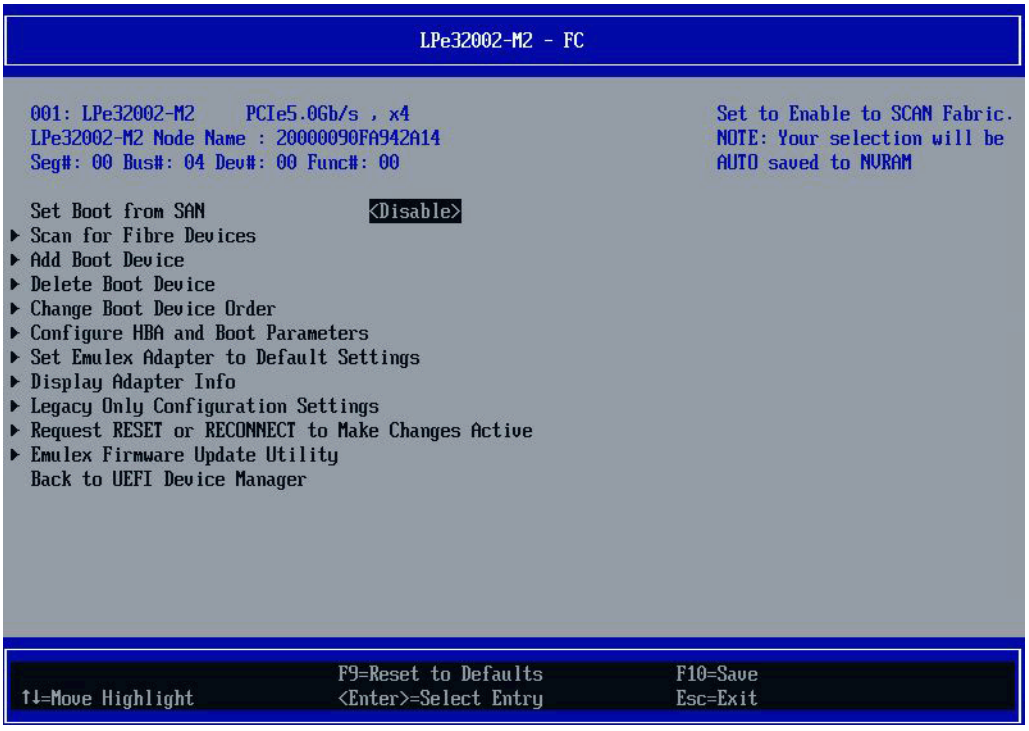

6.2.1 Emulex Configuration Utility Conventions in UEFI/HII

The Emulex Configuration utility has menus and configuration screens. Use the following methods to navigate them:

- Press the up and down arrows on your keyboard to move through and select menu options or configuration fields. When multiple adapters are listed, use the up and down arrows to scroll to the additional adapters.
- Press the **+**, **-**, or **Enter** keys to change numeric values.
- Press **Enter** to select a menu option, to select a row in a configuration screen, or to change a configuration default.
- Use the navigation entries on the page to move about the utility.
- Select **Commit** to save changes. Select **Discard** to not save changes.

6.3 Configuring Boot in UEFI/HII

The Emulex Configuration utility has numerous options that can be modified to provide for different behavior. Use the Emulex Configuration utility to do the following tasks:

- Set boot from SAN
- Scan for fibre devices
- Add and delete boot devices
- Change boot device order
- Configure boot settings that apply only to Emulex Legacy boot
- Configure HBA and boot parameters
- Set adapters to their default settings
- Display adapter information
- Legacy only configuration settings
- Request Reset or Reconnect to make changes active
- Emulex Firmware Update utility

NOTE Changes made to parameters common to UEFI and x86 drivers are changed in both driver configuration utilities.

6.4 Setting Boot from SAN

To set boot from SAN, perform these steps:

1. From the Device Configuration screen (Figure 26), select the adapter whose boot from SAN setting you want to change and press **Enter**.
2. From the Main menu (Figure 27), select **Set Boot from SAN**. The current boot setting is displayed. Press Enter.

   A Disable/Enable menu appears (Figure 28).

Figure 28 Main Menu, Boot from SAN Options Menu

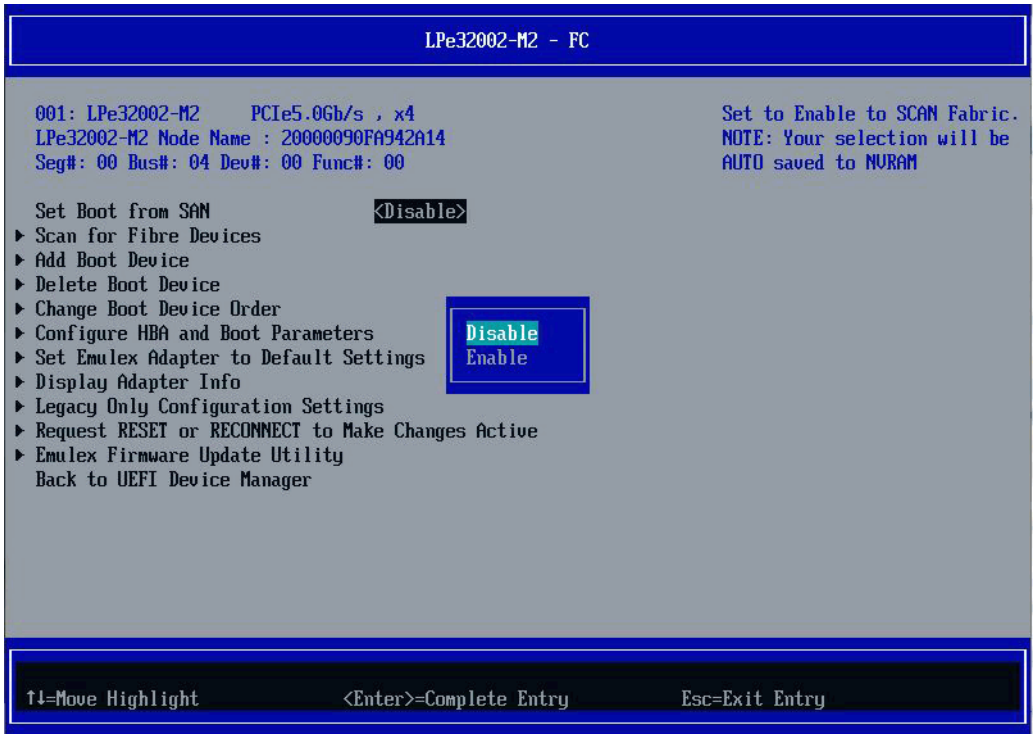

3. Make your selection and press **Enter**.
   The utility displays the new boot from SAN setting (Figure 29.

Figure 29 New Boot From SAN Setting

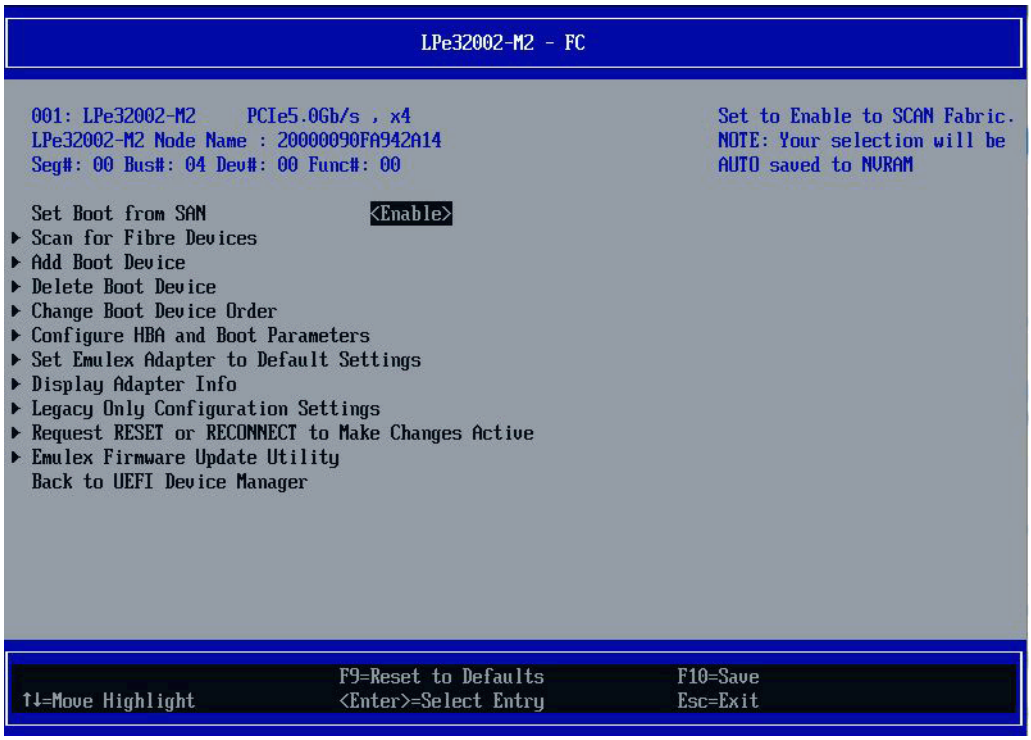

6.5 Scanning for Fibre Devices

To scan for fibre devices, perform these steps:

1. From the Device Configuration screen (Figure 26), select the adapter that you want to scan for Fibre devices and press **Enter**.
2. From the Main menu (Figure 27), select **Scan for Fibre Devices** and press **Enter**.

   A list of the discovered targets is displayed (Figure 30). This is only a list of discovered target devices to determine SAN connectivity. To add or configure boot devices, see the following sections.

Figure 30 Discovered Targets Screen

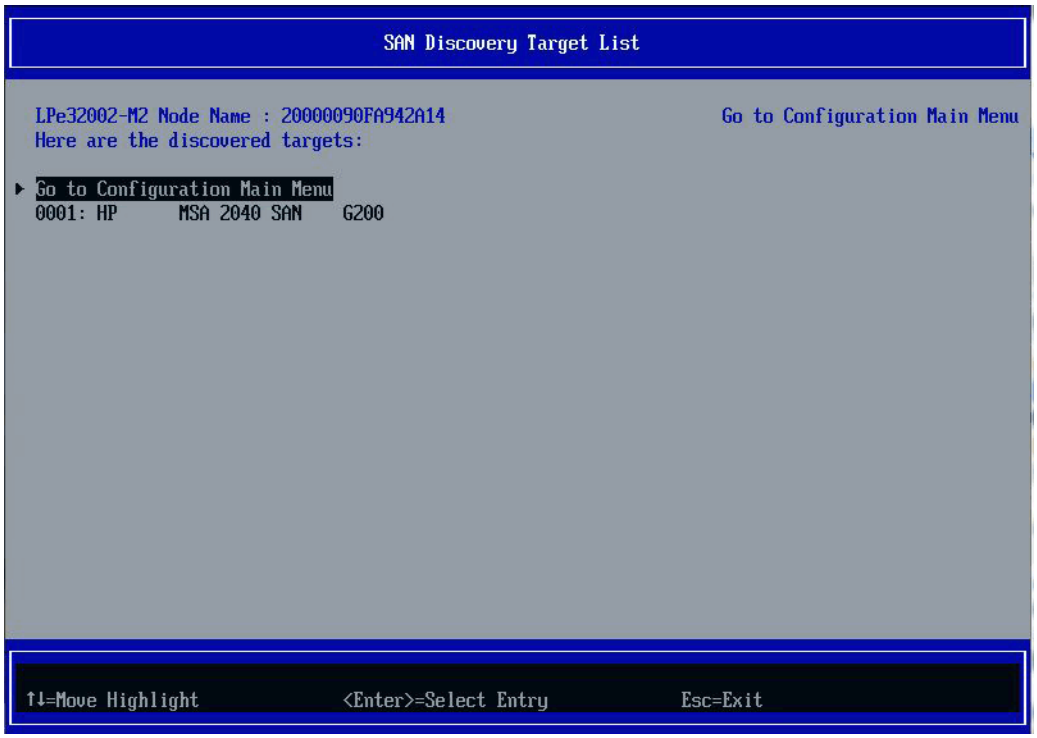

6.6 Adding Boot Devices

To add a boot device, perform these steps:

1. From the Device Configuration screen (Figure 26), select the adapter to which you want to add a boot device and press **Enter**.
2. From the Main menu (Figure 27), select **Add Boot Device** and press **Enter**.

   A screen appears displaying the discovered targets (Figure 31).

Figure 31 Discovered Targets Screen

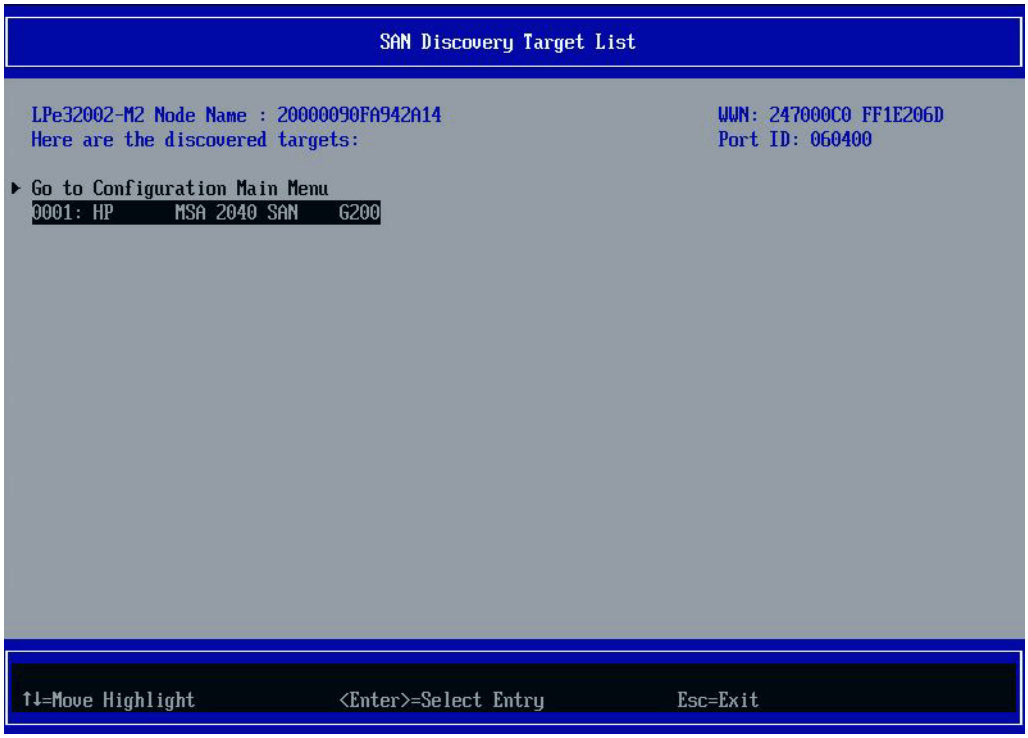

3. Select the target you want and press **Enter**.
   A list of bootable LUNS is displayed (Figure 32)

Figure 32 Boot Device Screen

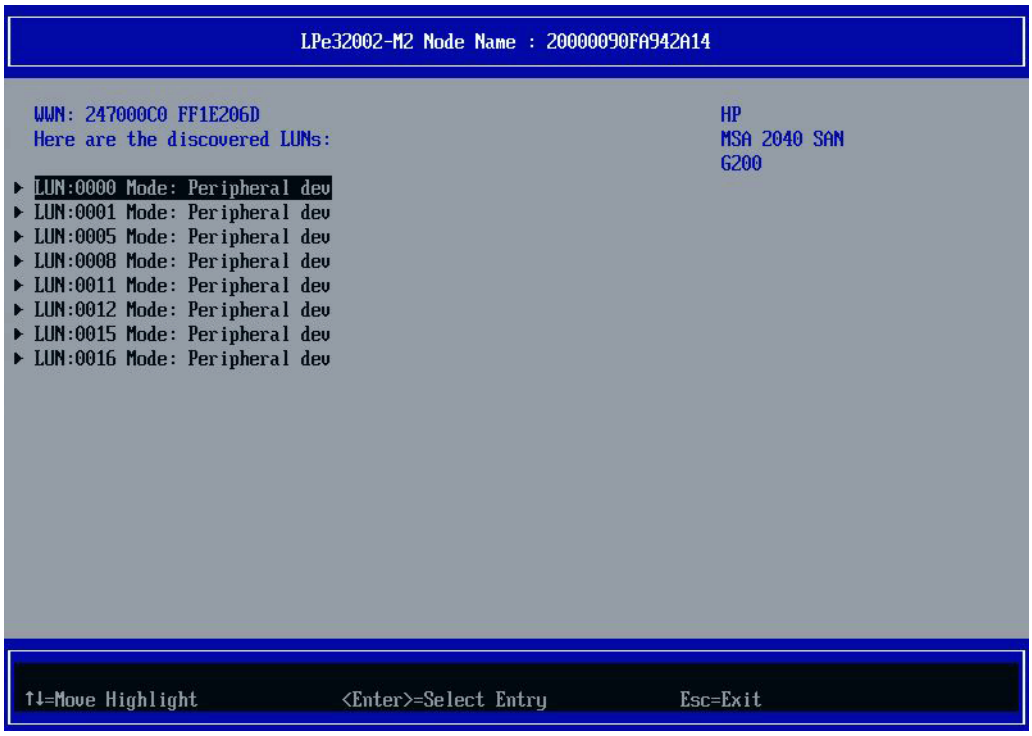

4. Select the boot device you want to add and press **Enter**.

   A menu appears that enables you to commit or discard your changes (Figure 33).

Figure 33 Commit/Discard Changes Menu

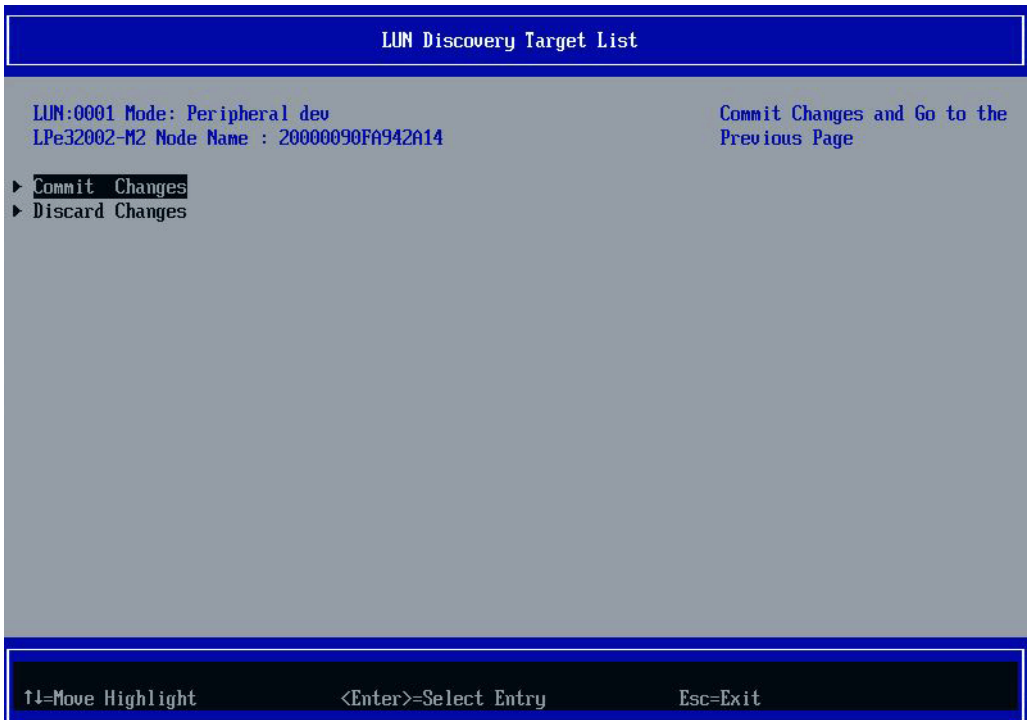

5. Select **Commit Changes** and press **Enter**.

6.7 Deleting Boot Devices

To delete boot devices, perform these steps:

1. From the Device Configuration screen (Figure 26), select the adapter from which you want to delete a boot devices and press **Enter**.
2. From the Main menu (Figure 27), select **Delete Boot Device** and press **Enter**.

   A list of boot devices is displayed (Figure 34).

Figure 34 Boot Device Screen

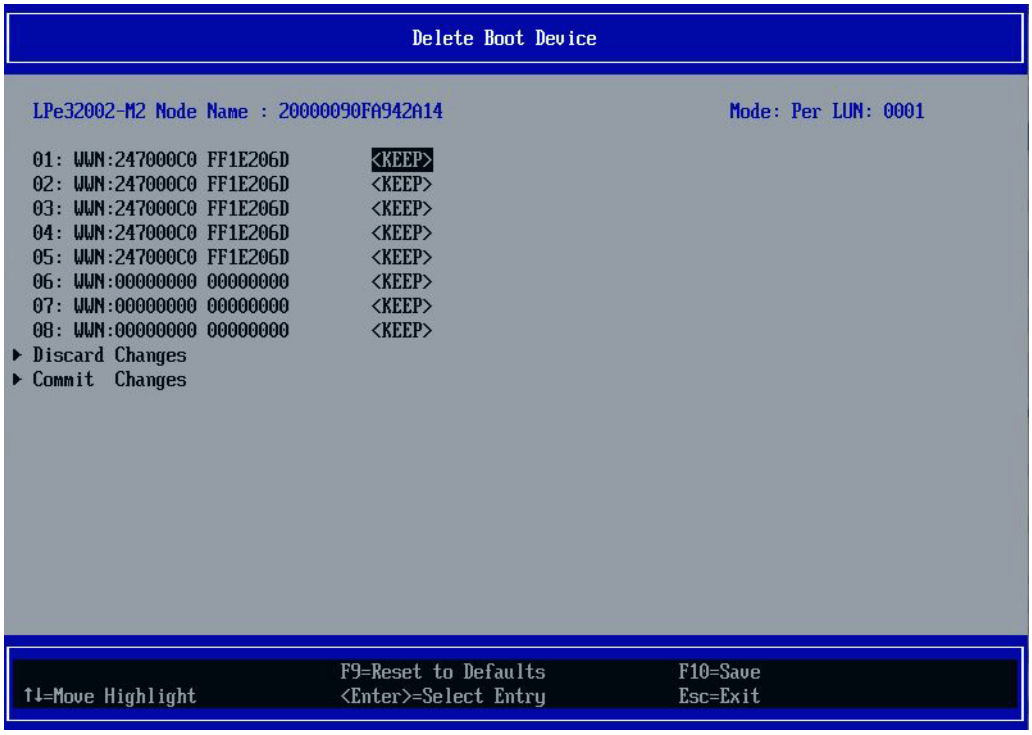

3. Select the boot device you want to delete and press the **space bar**.
   The device is highlighted and **<KEEP>** changes to **<DELETE>**. (Figure 35).

Figure 35 Selected Boot Device is highlighted

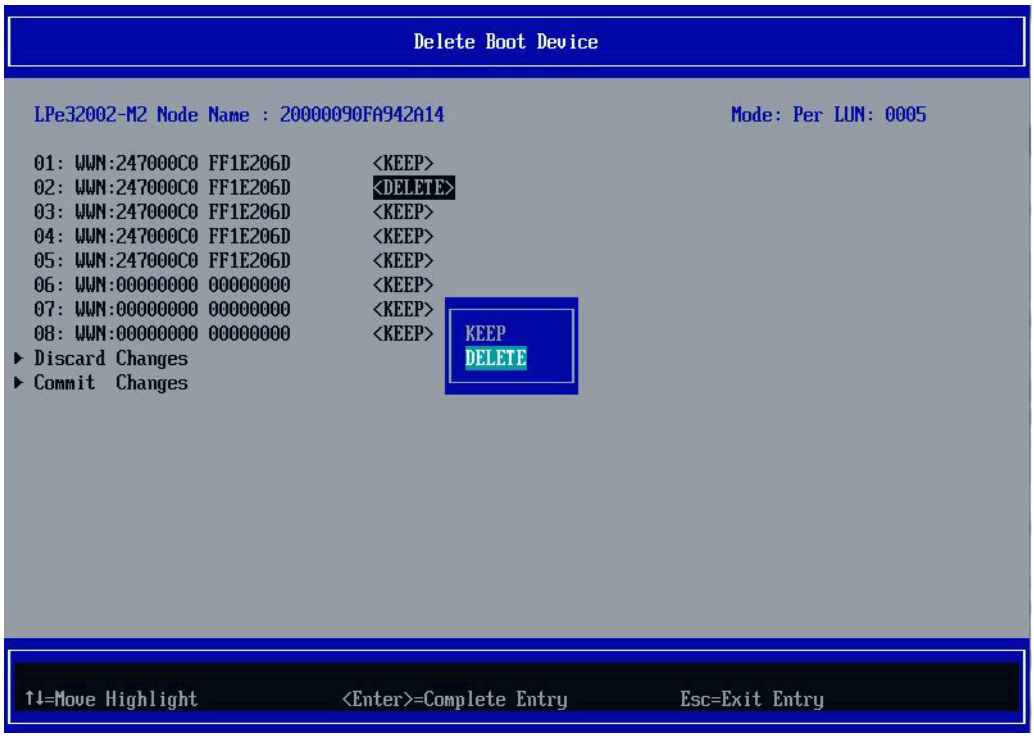

NOTE This screen is different for LPe16000, LPe31000, and LPe32000-series adapters, but the functionality is the same.

4. Select **Commit Changes** and press **Enter**.

Figure 36 Confirm Boot Device Delete

6.8 Changing Boot Device Order

NOTE Changing the boot device order only changes the order that the targets are discovered by the boot driver. The system BIOS controls the order that available boot devices in a system are prioritized.

To change boot device order, perform these steps:

1. From the Device Configuration screen (Figure 26), select the adapter whose boot device order you want to change and press **Enter**.
2. From the Main menu (Figure 27), select **Change Boot Device Order** and press **Enter**.

   A screen that displays the discovered targets appears (Figure 37).

Figure 37 Discovered Targets Screen

3. Select **Boot Device Order** and press **Enter**.

   A screen appears that displays the boot device order (Figure 38).

Figure 38 Boot Device Order Screen

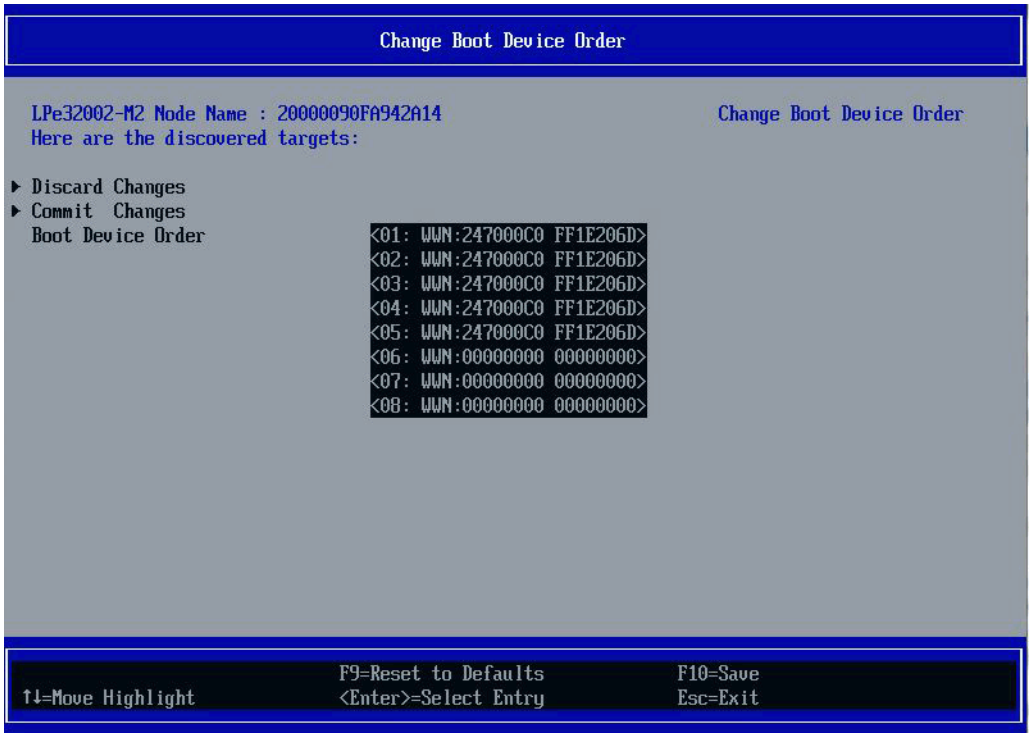

4. Press **Enter**.

   The Boot Device Order menu screen appears (Figure 39).

Figure 39 Boot Device Order Menu

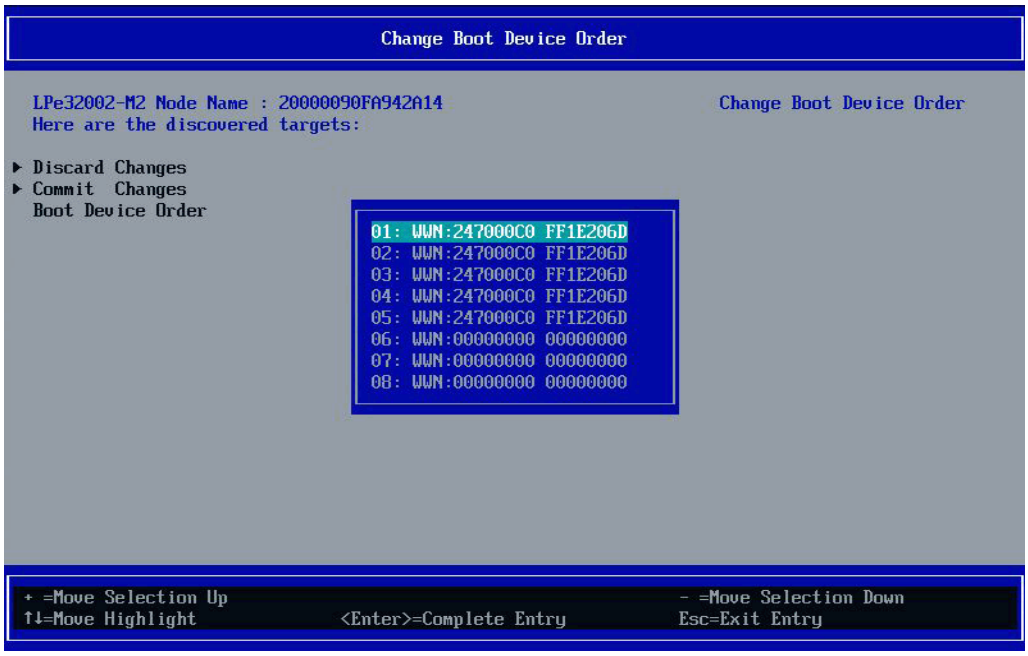

5. From the menu, select the device whose boot order you want to change. Use the **+** or **-** keys to change the order of the selected device and press **Enter**.

   A screen appears showing the new boot device order.

6. Press Enter to confirm the changes.

Figure 40 LPe12000-Series Adapters with Delay Device Discovery

Figure 41 LPe32000-Series Adapters with no Delay Device Discovery

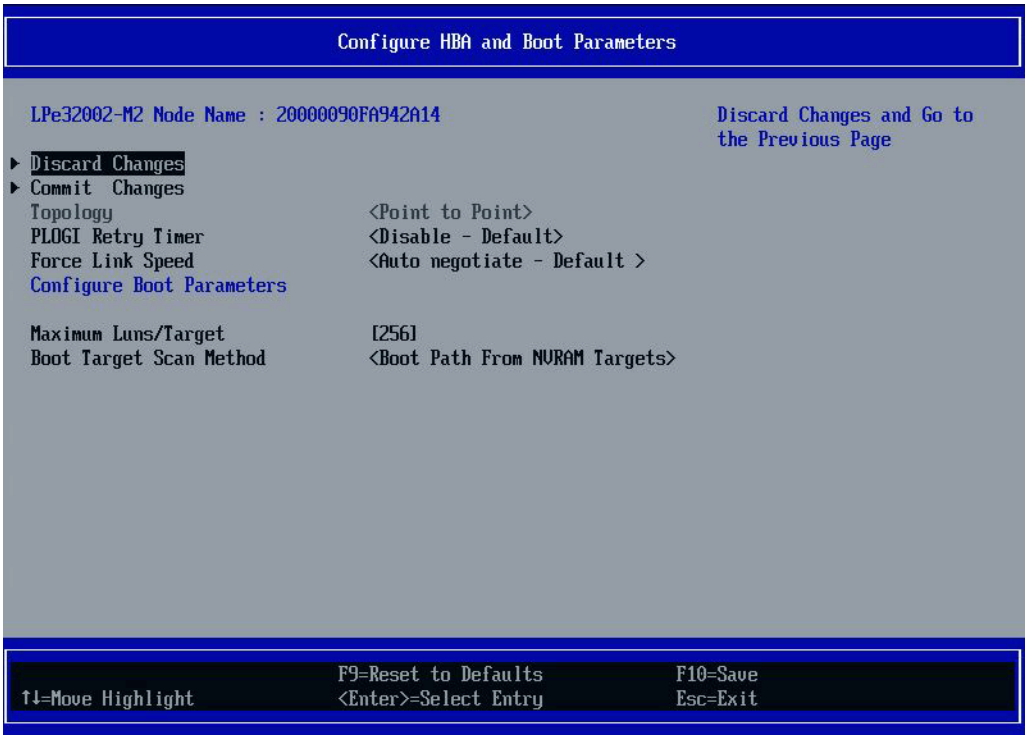

6.9 Configuring Adapter Parameters

The Emulex Configuration utility enables you to configure the following adapter parameters:

- Topology
- Port login (PLOGI) retry timer
- Link speed

6.9.1 Changing the Topology

NOTE Point to point is the only topology supported on LPe31000-series and LPe32000-series adapters.

By default, the LPe12000-series and LPe16000-series adapters use auto topology with loop first. If you want to use a topology other than the default, you must change the topology setting before configuring boot devices. For FC-AL, each adapter has a default AL_PA of 01 (hexadecimal).

The topology options available in LPe12000-series and LPe16000-series adapters are shown in Figure 42.

You can configure these items:

- Auto topology with loop first – default
- Auto topology with point-to-point first
- FC-AL
- Fabric point-to-point

To select the adapter topology, perform these steps:

1. From the Device Configuration screen (Figure 26), select the adapter whose topology you want to change and press **Enter**.
2. From the Main menu (Figure 27), select **Configure HBA and Boot Parameters** and press **Enter**.
3. From the Configure HBA and Boot Parameters menu, navigate to **Topology Selection** and press **Enter**. Information similar to Figure 42 is displayed.

Figure 42 Topology Menu in LPe12000-Series and LPe16000-Series Adapters

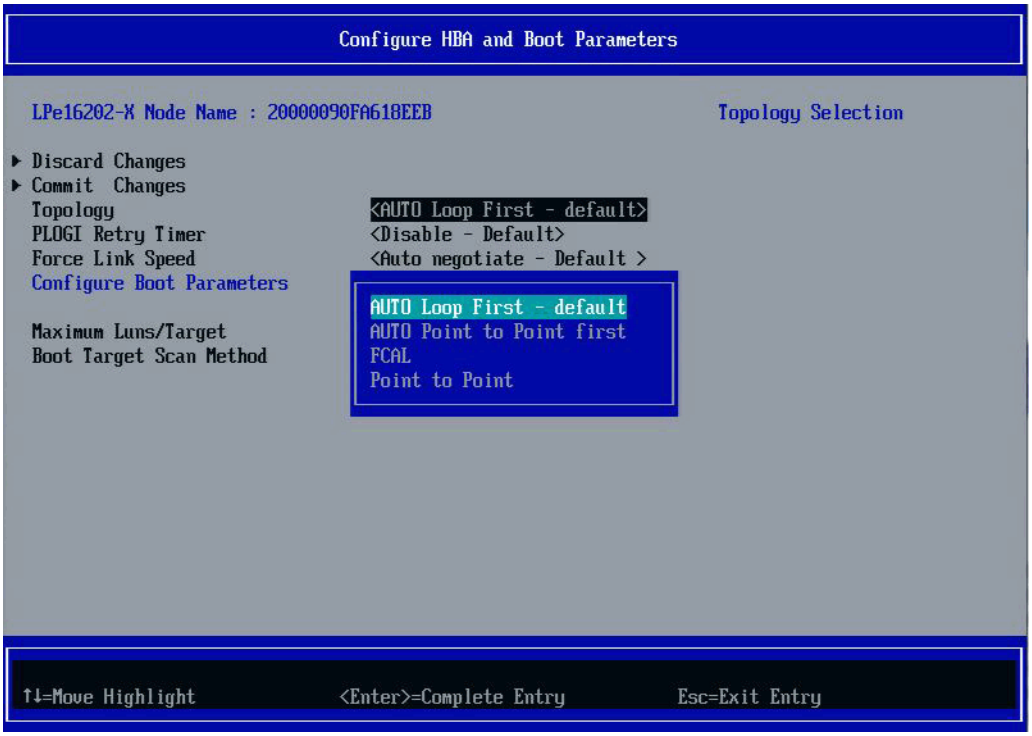

4. Select a topology and press **Enter**.
   The screen is refreshed with the new value.

   NOTE The presence of a fabric is detected automatically.
5. Press Esc to return to the Emulex Configuration Utility menu.
6. Select **Commit Changes** and press **Enter**.

The LPe31000-series and LPe32000-series adapters use point-to-point topology and cannot be changed. The topology options available in LPe31000-series and LPe32000-series adapters is shown in Figure 43.

Figure 43 Topology Menu in LPe31000-Series and LPe32000-Series Adapters

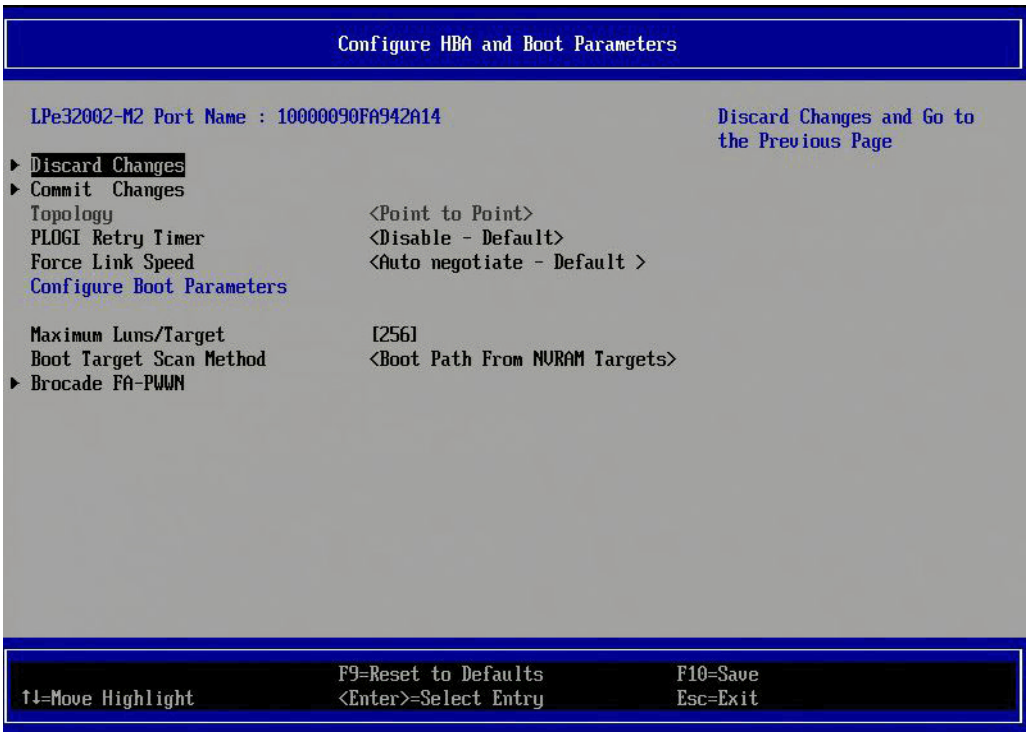

6.9.2 Changing the PLOGI Retry Timer

This option allows you to set the interval for the PLOGI retry timer. This option is especially useful for Tachyon-based RAID arrays. Under very rare occasions, a Tachyon-based RAID array resets itself, and the port goes offline temporarily in the loop. When the port comes to life, the PLOGI retry interval scans the loop to discover this device.

You can choose one of these options:

- Disable – Default
- 50 msec
- 100 msec
- 200 msec

To change timer values, perform these steps:

1. From the Device Configuration screen (Figure 26), select the adapter whose PLOGI retry timer information you want to change and press **Enter**.
2. From the Main menu (Figure 27), select **Configure HBA and Boot Parameters** and press **Enter**.
3. From the Configure HBA and Boot Parameters menu (Figure 42), navigate to **PLOGI Retry Timer** and press **Enter**. The PLOGI Retry Timer menu appears (Figure 44).

Figure 44 PLOGI Retry Timer Menu Screen

4. Select a retry timer option and press **Enter**.
   The screen is refreshed with the new value.
5. Press Esc to return to the Emulex Configuration Utility menu.
6. Select **Commit Changes** and press **Enter**.

6.9.3 Changing the Link Speed

Use this feature to change, or force, the link speed between ports instead of auto-negotiating. Supported link speeds are specific to the adapter model and the optics installed. The menu only displays options that are valid for the selected adapter.

NOTE In LPe12000-series adapters the Link Speed setting will be overwritten by the FC driver and will not take effect during a regular boot into the local operating system. Therefore, the link speed BIOS setting will work for the actual boot, but will be overwritten once the operating system loads.

Possible link speed choices:

- Auto negotiate – Default
- 1Gb/s (if available)
- 2Gb/s (if available)
- 4Gb/s (if available)
- 8Gb/s (if available)
- 16Gb/s (if available)
- 32Gb/s (if available)

To change the link speed, perform these steps:

1. From the Device Configuration screen (Figure 26), select the adapter whose link speed you want to change and press **Enter**.
2. From the Main menu (Figure 27), select **Configure HBA and Boot Parameters** and press **Enter**.
3. From the Configure HBA and Boot Parameters menu (Figure 42), navigate to **Force Link Speed** and press **Enter**.
   The Force Link Speed menu appears (Figure 45).

Figure 45 Force Link Speed Menu in LPe16000-Series Adapters

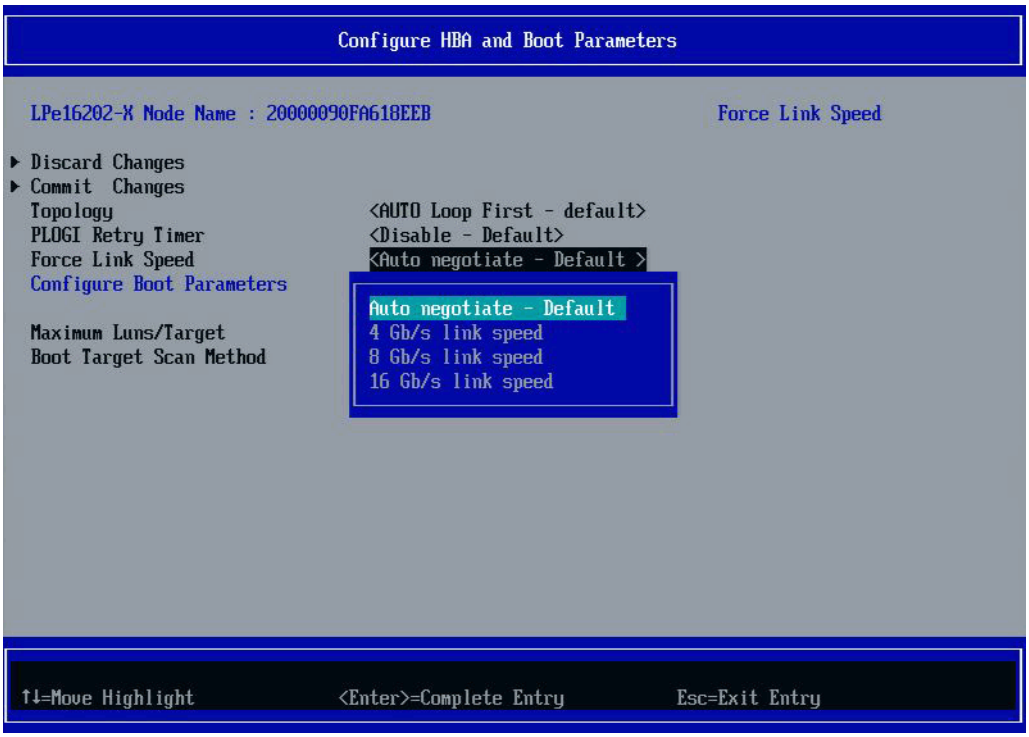

Figure 46 Force Link Speed Menu in LPe31000-Series and LPe32000-Series Adapters

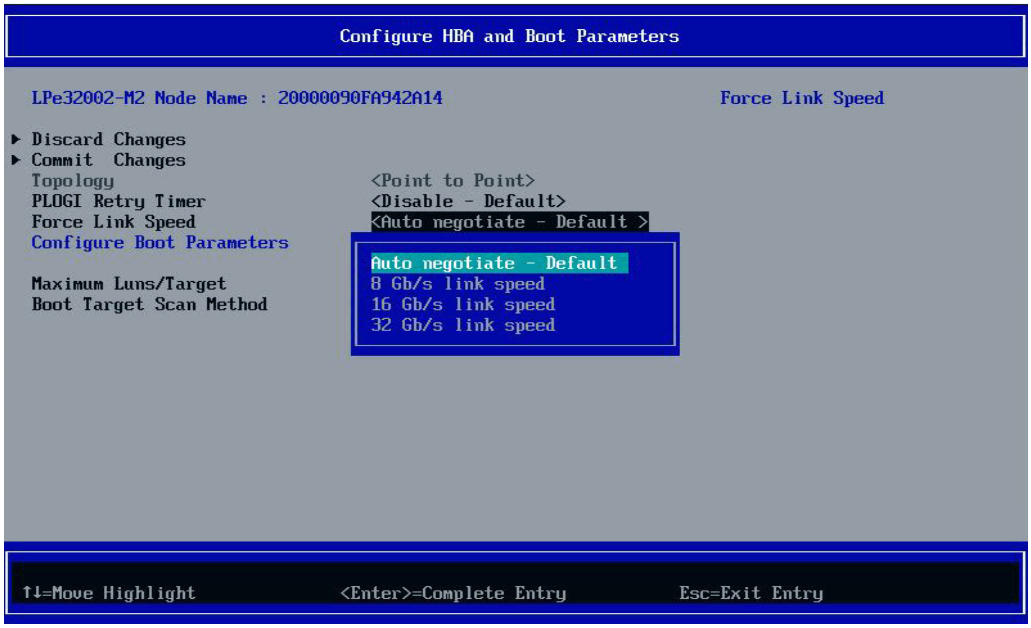

4. Select a link speed and press **Enter**.

   The screen is refreshed with the new value.
5. Press Esc to return to the Emulex Configuration Utility menu.
6. Select **Commit Changes** and press **Enter**.

6.10 Configuring Boot Parameters

You can change the following parameters:

- Maximum LUNS per target
- Boot target scan method
- Device discovery delay

6.10.1 Changing the Maximum LUNs per Target

The maximum number of LUNs represents the maximum number of LUNs that are polled during device discovery. The minimum value is 1, the maximum value is 4096. The default is 256.

To change the maximum number of LUNs, perform these steps:

1. From the Device Configuration screen (Figure 26), select the adapter whose maximum LUNs per target information you want to change and press **Enter**.
2. From the Main menu (Figure 27), select **Configure HBA and Boot Parameters** and press **Enter**.
3. From the Configure HBA and Boot Parameters menu (Figure 42), navigate to **Maximum Luns/Target** and press **Enter**.

   The screen becomes editable (Figure 47).

Figure 47 Maximum LUNs/Target Screen

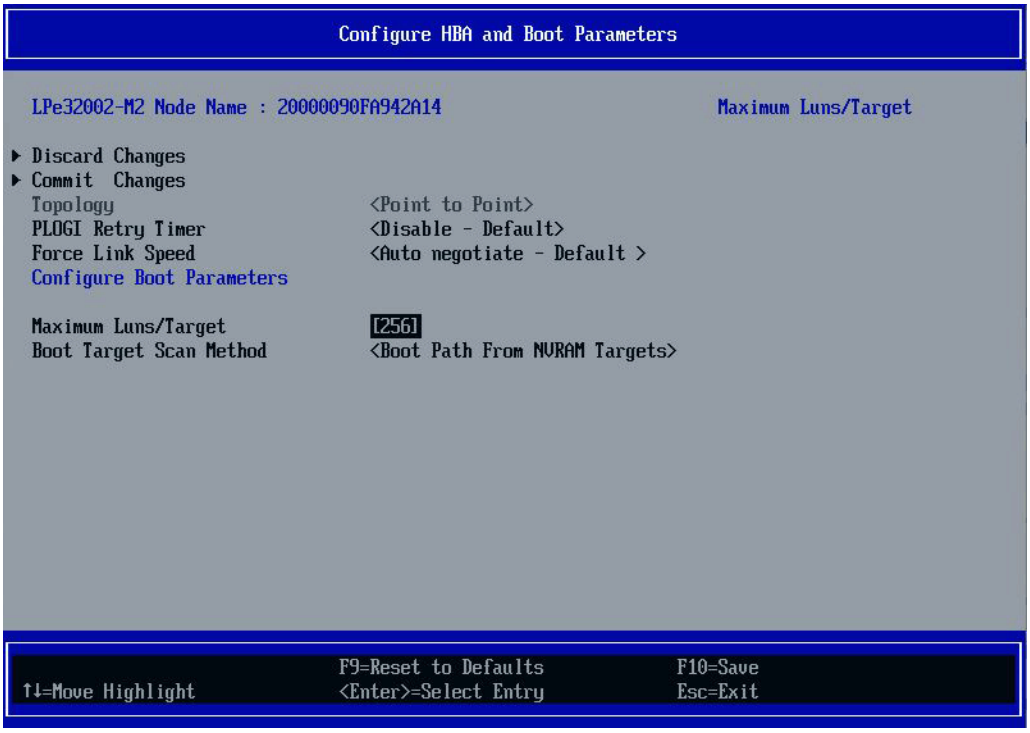

4. Type the maximum LUN value (between 1 and 4096) and press **Enter**.
   The screen is refreshed with the new value.

   NOTE 256 is the default, and typical, maximum number of LUNs in a target device. A higher number of maximum LUNs causes the discovery to take more time.
5. Press Esc to return to the Emulex Configuration Utility menu.
6. Select **Commit Changes** and press **Enter**.

6.10.2 Changing Boot Target Scan Method

This option is available only if none of the eight boot entries are configured to boot from DID or WWPN. The Configure Boot Devices menu configures up to eight boot entries for fabric point-to-point, public loop, or private loop configurations.

The boot scan options follow:

- Boot Path from NVRAM Targets – Boot scan discovers only LUNs that are saved to the adapter's NVRAM. Select up to eight attached devices to use as potential boot devices. Limiting discovery to a set of eight selected targets can greatly reduce the time it takes for the EFIBoot driver to complete discovery.
- Boot Path from Discovered Targets – Boot scan discovers all devices that are attached to the FC port. Discovery can take a long time on large SANs if this option is used.
- Do not create a boot path.
- Boot Scan from EFIFCScanLevel – Allows third-party software to toggle between Boot Path from NVRAM and Boot Path from Discovered Targets by manipulating an UEFI system NVRAM variable. After the scan is set to EFIFCScanLevel, the scan method can be changed without entering the UEFI Boot configuration utility.

If EFIFCScanLevel is selected, the scan is determined by the value of the EFIFCScanLevel variable maintained by the UEFI system firmware or boot code. The value of this variable can be changed either by using the menu in the EFIBoot Configuration utility, or by using third-party software.

To change the boot target scan method, perform these steps:

1. From the Device Configuration screen (Figure 26), select the adapter whose boot target scan method you want to change and press **Enter**.
2. From the Main menu (Figure 27), select **Configure HBA and Boot Parameters** and press **Enter**.
3. From the Configure HBA and Boot Parameters menu (Figure 42), navigate to **Boot Target Scan Method** and press **Enter**.

   The Boot Target Scan Method menu appears (Figure 48).

Figure 48 Boot Target Scan Method Menu Screen

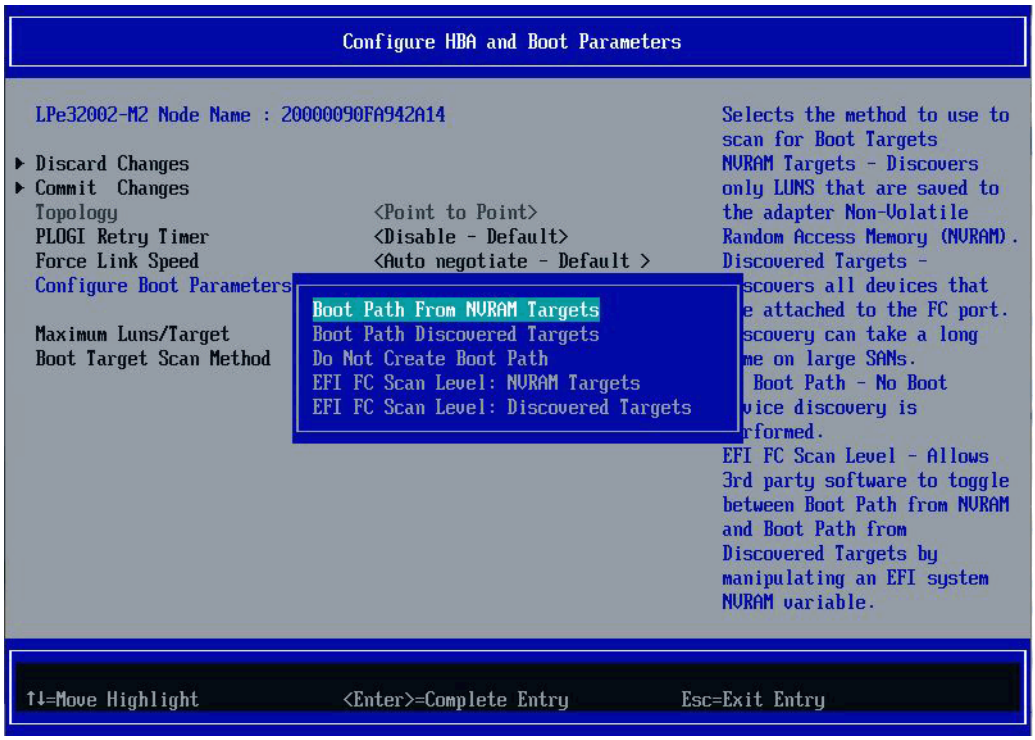

4. Select a boot scan setting and press **Enter**.

   The screen is refreshed with the new value. If you have a large SAN and set the boot path to **Boot Path Discovered Targets**, discovery takes a long time.
5. Press **Esc** to return to the Emulex Configuration Utility menu.
6. Select **Commit Changes** and press **Enter**.

6.10.3 Changing Device Discovery Delay

NOTE This option is available only in LPe12000-series adapters and is not available in LPe16000-series, LPe31000-series, and LPe32000-series adapters.

This parameter sets a delay to occur after an loop initialization and before a scan is initiated to discover the target. The default is off or 0 seconds.

Change the default if you have an HP MSA1000 or HP MSA1500 RAID array and if both of the following conditions exist:

- The MSA array is direct-connected or is part of an arbitrated loop (for example, daisy-chained with a JBOD).
- The boot LUN is not reliably discovered. In this case, a delay might be necessary to allow the array to complete a reset.

CAUTION Do not change the delay device discovery time if your MSA array is connected to a fabric switch. Setting it to any other time guarantees that the maximum delay time is seen on every loop initialization.

If both of the conditions exist, typically set this parameter to 20 seconds. However, the delay only needs to be long enough for the array to be reliably discovered after a reset. Your value might be different.

To change the delay device discovery value, perform these steps:

1. From the Device Configuration screen (Figure 26), select the adapter whose device discovery delay settings you want to change and press **Enter**.
2. From the Main menu (Figure 27), select **Configure HBA and Boot Parameters** and press **Enter**.
3. From the Configure HBA and Boot Parameters menu (Figure 42), select **Delay Device Discovery** and press **Enter**.

   The screen becomes editable (Figure 49).

Figure 49 Delay Device Discovery Screen

4. Use the **+/-** keys to change the delay device discovery value in increments of 10 seconds and press **Enter**.
   The screen is refreshed with the new value.
5. Press Esc to return to the Emulex Configuration Utility menu.
6. Select **Commit Changes** and press **Enter**.

6.11 Configuring Brocade FA-PWWN

NOTE This feature is available in LPe16000-series, LPe31000-series, and LPe32000-series adapters only.

This option enables or disables the Fabric Assigned Port Word Wide Name (FA-PWWN). The default setting is disabled.

To enable or disable the Brocade FA-PWWN, perform these steps:

1. From the Device Configuration screen (Figure 26), select the adapter whose Brocade FA-PWWN you want to enable or disable and press **Enter**.
2. From the Main menu (Figure 27), select **Configure HBA and Boot Parameters** and press **Enter**.
3. From the Configure HBA and Boot Parameters menu (Figure 42), navigate to **Brocade FA-PWWN** and press **Enter**.
   The Brocade FA-PWWN screen appears (Figure 50).

Figure 50 Brocade FA-PWWN Screen

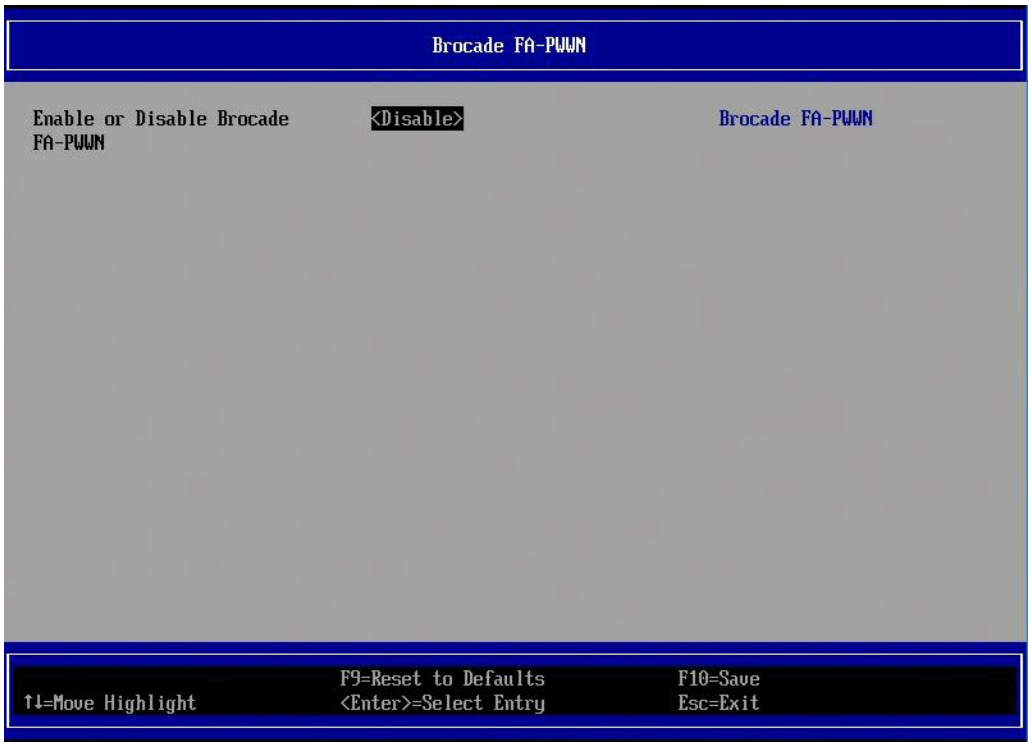

4. Select whether to enable or disable Brocade FA-PWWN.
5. Press **Enter** to accept the new value.
6. Press **Esc** to return to the Emulex Configuration Utility menu.

7. Select **Commit Changes** and press **Enter**.
8. Reboot the system.

After enabling the feature in the Configure HBA and Boot Parameters screen, the Port WWPN gets the configured FA-PWWN from the switch.

6.12 Resetting Emulex Adapters to Default Values

NOTE Resetting the adapter to default setting will clear all entries made while configuring the boot device, as detailed in Section 4.3, Configuring Boot Devices).

NOTE This option is available only in LPe12000-series and LPe16000-series adapters and is not available in LPe31000-series and LPe32000-series adapters.

The Emulex Configuration utility enables you to clear the NVRAM target list and set all boot device WWNNs back to 0.

These defaults are listed in Table 6.

Table 6 Adapter Defaults

Parameter	Default	Valid Values
Enable/Disable BIOS	Disabled	Enabled Disabled
PLOGI Retry Timer	Disabled	Disabled 50 msec 100 msec 200 msec
Boot Target Scan	Boot path from NVRAM targets	Boot path from NVRAM targets Boot path discovered targets Do not create boot path
Max LUNs Setting	256	0–4096
Topology	Auto Topology: Loop First (Default) **NOTE** On LPe31000-series and LPe32000-series adapters,only Point-to-Point is available.	Auto Topology: Loop First (Default) Auto Topology: Point-to-Point First FC-AL Point-to-Point
Delay Device Discovery	0000	0000–0255
Link Speed	0 (Auto-select)	0 = Auto Select (the adapter's speed is selected automatically based on its model). 1 = 1Gb/s 2 = 2Gb/s 4 = 4Gb/s 8 = 8Gb/s 6 = 16Gb/s 32 = 32Gb/s

To set Emulex adapters to their default settings, perform these steps:

1. From the Device Configuration screen (Figure 26), select the adapter whose default settings you want to change and press **Enter**.

2. From the Main menu (Figure 27), select **Set Emulex Adapters to Default Settings** and press **Enter**.
   A menu screen appears that enables you to set defaults or cancel default settings (Figure 51).
3. Select **Set Adapter Defaults** and press **Enter**.

Figure 51 Adapter Selection Screen

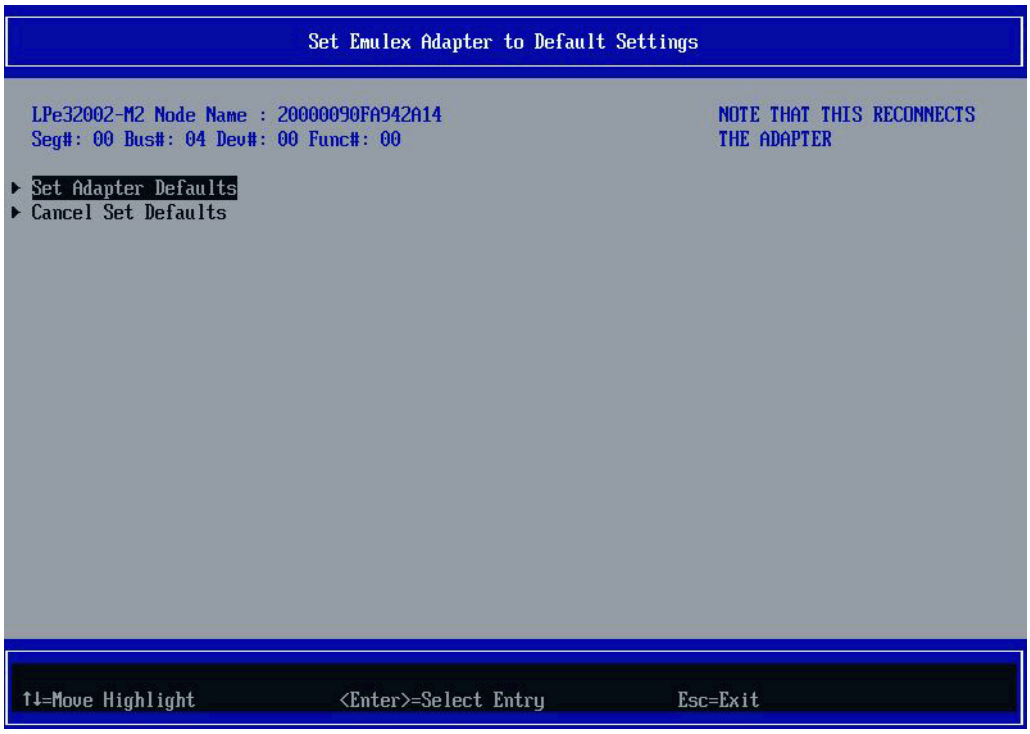

6.13 Displaying Adapter Information

The Adapter Information screen displays the following information about the selected adapter:

- HBA status
- Boot from SAN status
- Link speed
- Topology
- Firmware version
- Universal Boot version
- UEFI Boot version

To display adapter information, perform these steps:

1. From the Device Configuration screen (Figure 26), select the adapter whose information you want to view and press **Enter**.
2. From the Main menu (Figure 27), select **Display Adapter Info** and press **Enter**.
   A screen appears that displays information about the selected adapter (Figure 52).

Figure 52 Adapter Information Screen

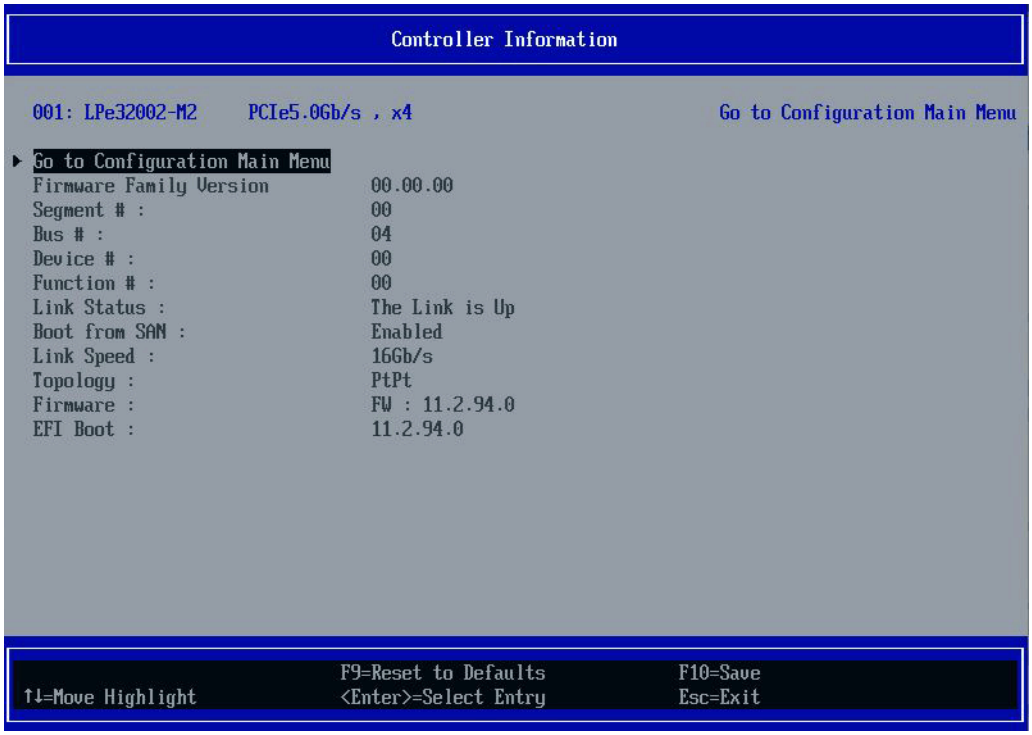

6.14 Legacy Only Configuration Settings

To change legacy only configuration settings, perform these steps:

1. From the Device Configuration screen (Figure 26), select the adapter whose Legacy settings you want to change and press **Enter**.
2. From the Main menu (Figure 27), select **Legacy Only Configuration Settings** and press **Enter**.

   A screen appears that displays settings that apply only to the Emulex Legacy boot driver (Figure 53).

Figure 53 Legacy Only Configuration Settings Screen

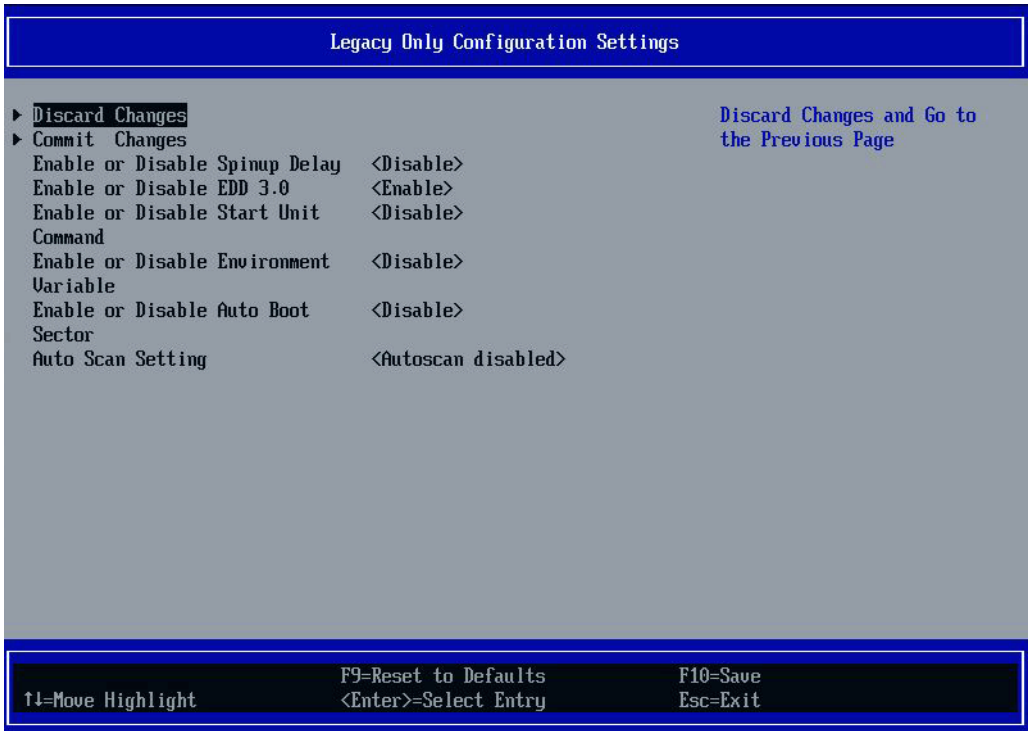

3. Make desired changes to any of these settings and select **Esc** or **Commit Changes**.

See Chapter 4: Emulex BIOS Utility for a description of the Legacy Only settings displayed in Figure 53.

6.15 Requesting a RESET or RECONNECT

NOTE This feature is available in LPe16000-series, LPe31000-series, and LPe32000-series adapters only.

After you have made changes to the UEFI BIOS, you might need to reset or reconnect the adapter for the changes to take effect. Reset or Reconnect is disabled by default, and you must enable it if you want the adapter to automatically reset or reconnect. Otherwise, the changes will not take effect until you manually reboot the server.

Reset forces the system to reboot after you exit the UEFIBoot (HII) utility.

Reconnect unloads and reloads the driver. This allows any changes to take effect immediately without requiring a reboot. The Reconnect feature also requires UEFI 2.5 compliant BIOS to enable this functionality. If the server BIOS is not UEFI 2.5 compliant, this option will be unavailable (grayed-out).

To request a reset or reconnect, perform these steps:

1. From the Device Configuration screen (Figure 26), select the adapter which you want to reset or reconnect and press **Enter**.
2. From the Main menu (Figure 27), select **Request RESET or RECONNECT to Make Changes Active**. The Request RESET or RECONNECT to Make Changes Active screen appears.

Figure 54 Request RESET or RECONNECT to Make Changes Active Screen - Reset

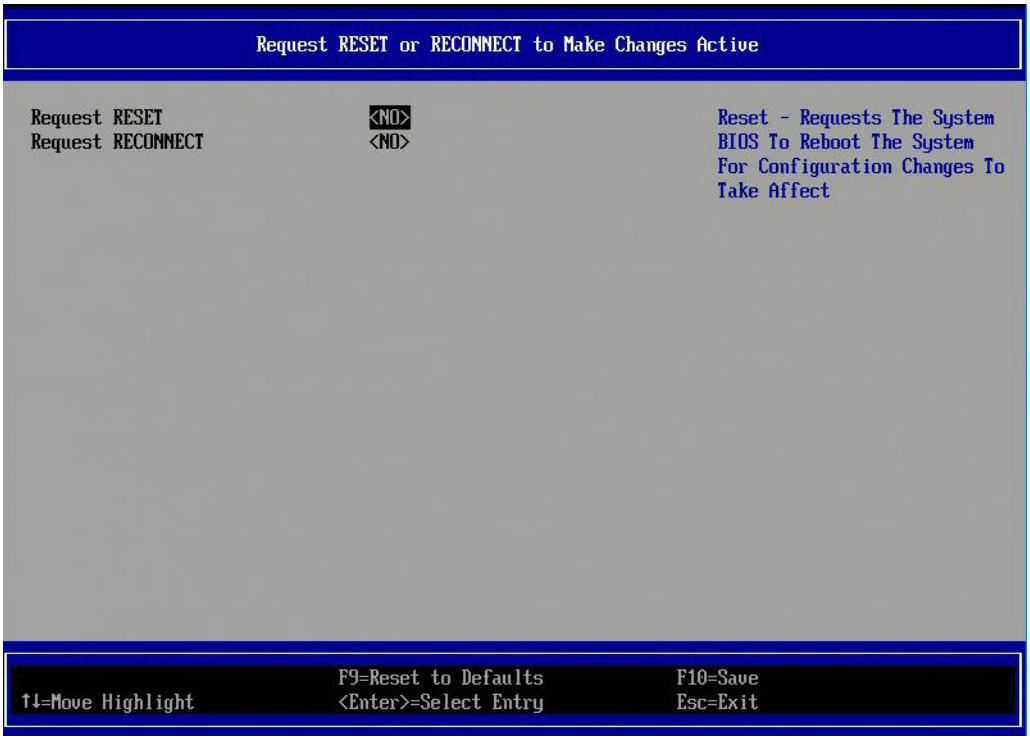

Figure 55 Request RESET or RECONNECT to Make Changes Active Screen - Reconnect

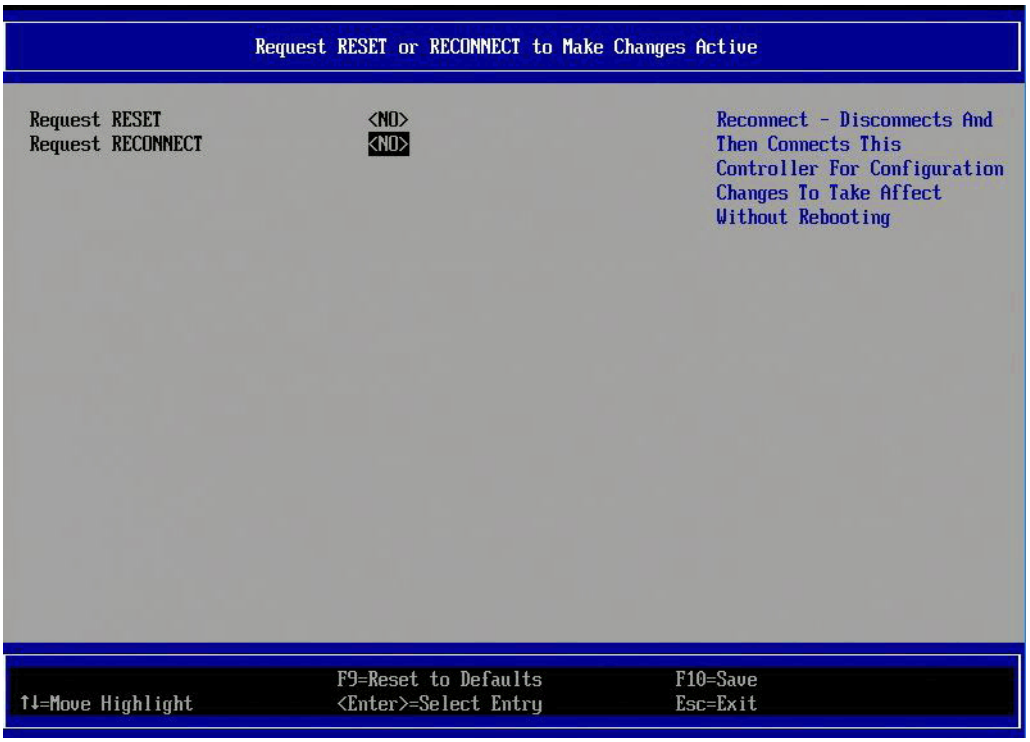

3. Do one of the following:

   To request a reset:

   a. Set **Request RESET** to yes to notify the system BIOS that you want to perform a system reset.
   b. Exit the UEFIBoot (HII) utility and return to the system BIOS setup screens.
   c. A pop up window prompts you to reset the system. Any changed Emulex adapter settings become active during the subsequent reboot.

   To request a reconnect:

   a. Set **Request RECONNECT** to yes to notify the system BIOS that you want to perform a disconnect and reconnect on the adapter port.
   b. The changed Emulex adapter settings are activated without a system reset.

6.16 Emulex Firmware Update Utility

NOTE This feature is available in LPe16000-series, LPe31000-series, and LPe32000-series adapters only.

To update the firmware on an adapter, perform the following steps:

1. From the Device Configuration screen (Figure 26), select the adapter whose firmware you want to update and press **Enter**.
2. From the Main menu (Figure 27), select **Emulex Firmware Update Utility**. The Emulex Firmware Update Utility screen appears.

Figure 56 Emulex Firmware Update Utility Screen

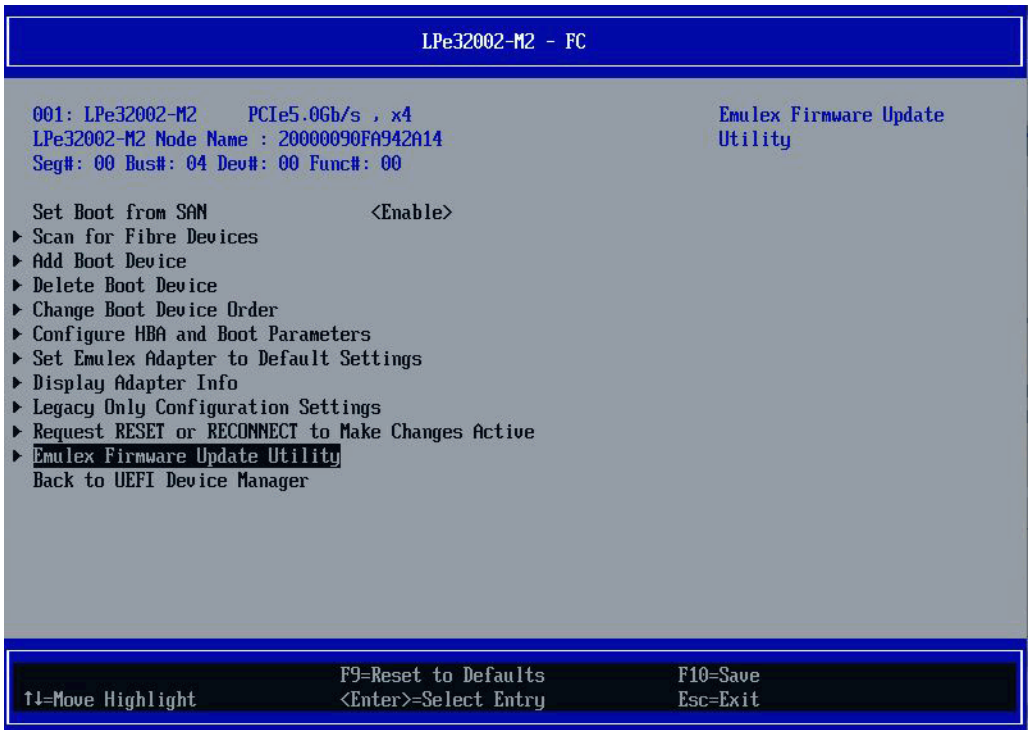

3. A list of devices appears. Select the storage device that contains the firmware file. The device can be a USB flash drive or a local disk.

Figure 57 Select Device Screen

4. After you select the storage device, press **Enter**. The **Select File or Folder** screen appears.

Figure 58 Select File or Folder Screen

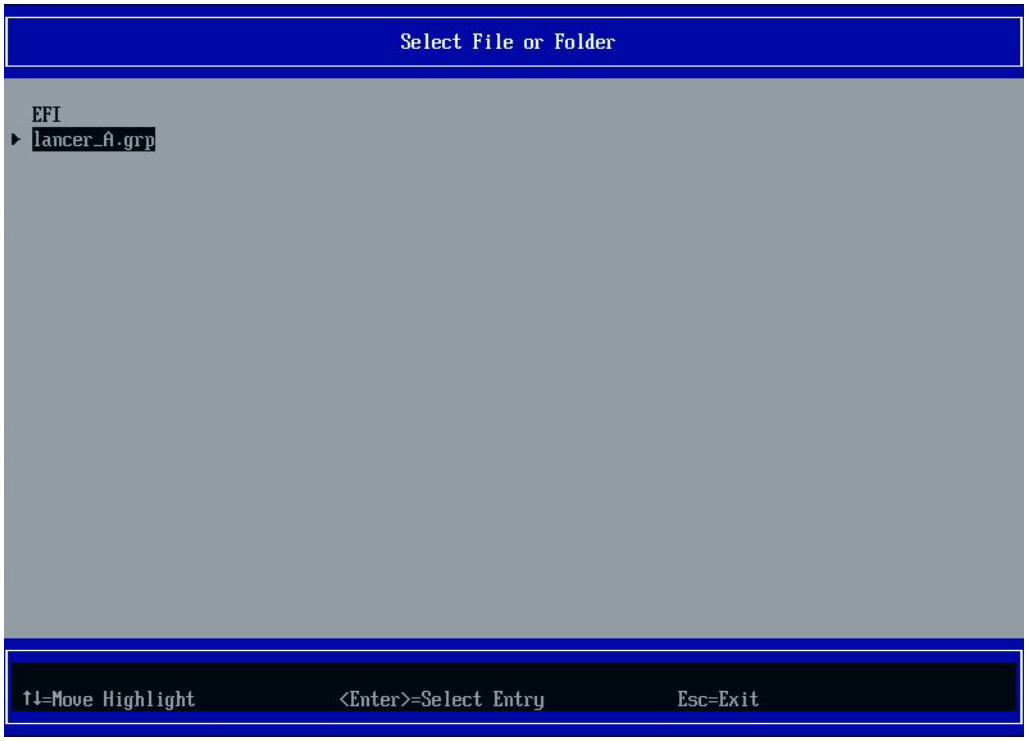

5. Browse and select the folder that contain the firmware image. Select the firmware file and press **Enter**. The **Start Update** screen appears.

Figure 59 Start Update Screen

6. Select **Start Downloading** and press **Enter**. When the operation is complete, the **Update Status** will indicate the completion status.

6.17 Enabling or Disabling OCSD Settings (HP Systems Only)

During a POST, this setting causes the EFIBoot driver to register (ENABLED), or not register (DISABLED), the device (that is, to add any entries) in the OCSD or OCBB buffers, as defined in the Option Card Sensor Data Reporting specification.

NOTE This menu is for HP systems only and does not appear on other systems.

To enable or disable OCSD, perform these steps:

1. Open the HP **System Configuration** menu (Figure 60).

Figure 60 HP System Configuration Menu Screen

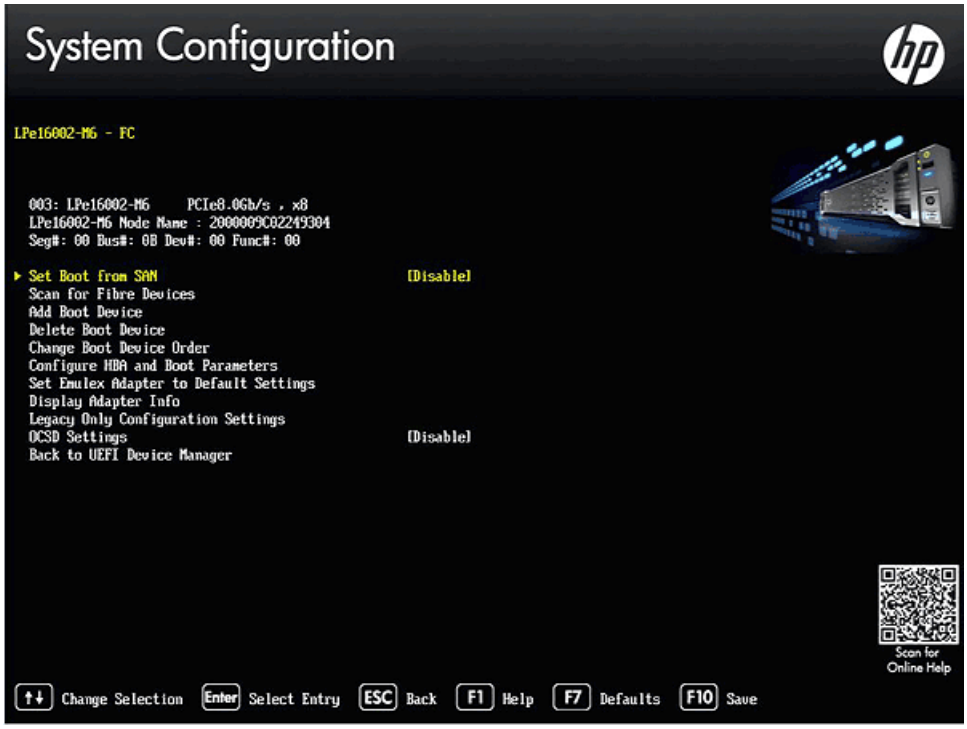

2. Navigate to **OCSD Settings** and press **Enter**.
   The Disable/Enable menu appears. (Figure 61).

Figure 61 OCSD Settings Menu Screen

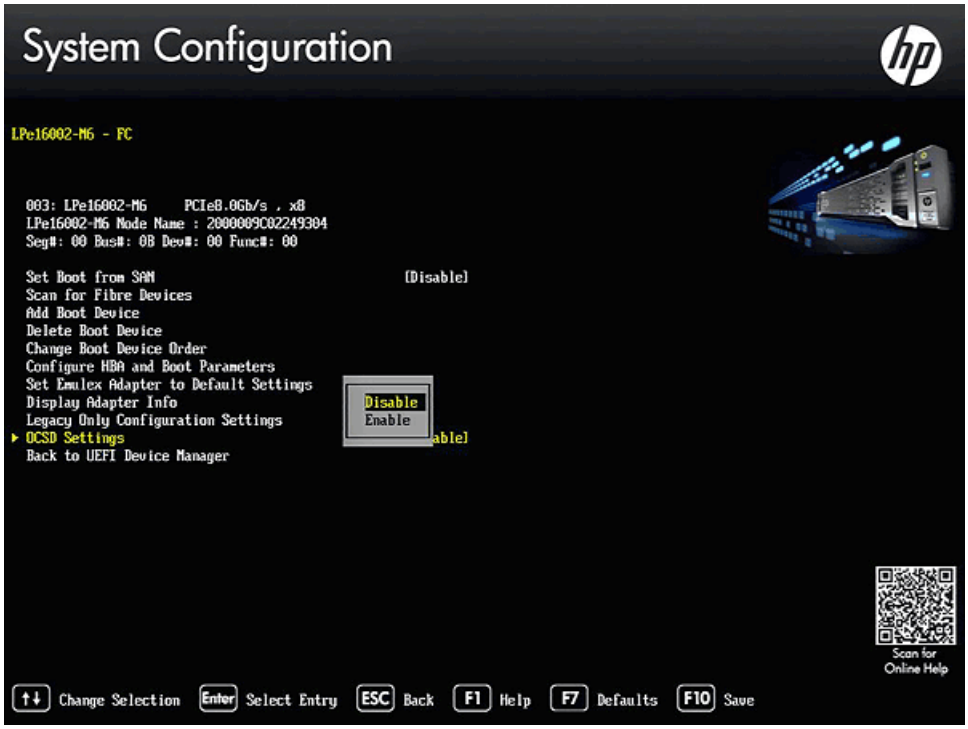

3. Make your selection and press **F10** to save it.

6.18 Returning to the UEFI Device Manager

To return to the UEFI Device Manager, select **Back to UEFI Device Manager** to return to the Device Configuration screen (Figure 26).

Chapter 7: Troubleshooting

Circumstances exist in which your system might operate in an unexpected manner. This section describes several of these circumstances and offers one or more workarounds for each situation.

7.1 x86 BootBIOS

The following table lists x86 BootBIOS issues and their resolutions.

Table 7 x86 BootBIOS Troubleshooting

Issue	Situation	Resolution
The Bootup Message Does Not Appear As the System Boots	You want to access the BIOS utility, but the bootup message does not appear.	Make sure that x86 BootBIOS has been loaded and enabled.
Retry This Adapter Message	The message `Retry This Adapter` appears during BIOS scanning.	Check the hardware configuration or reconfigure the adapter BIOS using the BIOS utility.
Cannot Mount Root File System Message (Solaris SFS Driver)	The message `Cannot Mount Root File System` appears during bootup.	Make sure the correct storage device is identified in the `scsi_vhci.conf` file.
Cannot Find UNIX Kernel Message (Solaris SFS Driver)	The message `Cannot Find UNIX Kernel` appears during bootup.	Set up the correct LUN to boot in the BIOS utility. The correct LUN can be seen at the end of the Device Address line when you issue a `luxadm display <device>` command. Refer to the luxadm documentation from Oracle for more information.
No Such Partition Message (Solaris SFS Driver)	The message `No Such Partition` appears during bootup.	Make sure the correct boot device is selected at the GRUB menu. Refer to the GRUB documentation from Oracle and the `/boot/grub/menu.lst` for more information.

`Link is down due to Unsupported Optic Installed`
or
`Unsupported Optic`
	These messages appear while scanning for targets or configuring boot devices in the Boot BIOS or during POST.
	`Link is down due to Unsupported Optic Installed`
	or
	In the Link Speed Selection menu
	`Unsupported Optic Installed`
	is displayed
	or
	In the Link Status field
	`Unsupported Optic`
	is displayed.
		Replace any unsupported optics with supported optics.

7.2 OpenBoot

The following table lists OpenBoot issues and their resolutions.

Table 8 OpenBoot Troubleshooting

Issue	Situation	Resolution
The System Cannot mount or fsck /etc/vfstab a FC Boot Disk (Solaris LPFC Driver)	During the boot process, the system cannot mount or `fsck /etc/vfstab` a FC boot disk.	Make sure that persistent binding is implemented correctly.
A Loaded File Is Not Executable (Solaris LPFC Driver)	After entering boot disk, a message states that the file that was loaded is not executable.	The boot block might not be installed correctly to the FC drive. See Section 2.3.1, Configuring Boot from SAN on Solaris (x86 and x64).
The System Hangs or Reboots After Displaying Driver Information (Solaris LPFC Driver)	The system hangs for a long time after displaying driver information, or it reboots after displaying driver information.	Possible incorrect topology set in the `/kernel/drv/lpfc.conf` file on the target disk.
FC Disk Is Not Found (Solaris LPFC Driver)	You have performed the setup tasks and the FC disk is not found when you reboot the system.	If the FC disk is not found when the system is rebooted, it might be necessary to do the following: 1. Type `cfgadm -a` to list the target. 2. Type `cfgadm -vc configure c1::c5t2200002037AE0091` to configure the FC target. 3. Type `cfgadm -c unconfigure c1` to remove the FC target. It also might be necessary to add an entry for the boot drive to the `sd.conf` file.

The Displayed List of Emulex Adapters Ends with "fibre-channel" (Solaris LPFC Driver)	After all Emulex adapters have been enabled to boot from SAN, the system has been rebooted, and you show all system devices, the path to an Emulex adapter ends with `fibre-channel`, for example:
	`/pci@1f,2000/fibre-channel`
	instead of `lpfc@#`, for example:
	`/pci@1f,4000/lpfc@2`
		The OpenBoot code is not loaded in the adapter firmware. You must install OpenBoot before you can set up boot from SAN on the adapter.

7.3 UEFI Boot

The following table lists UEFI Boot issues and their resolutions.

Table 9 UEFI Boot Troubleshooting

Issue	Situation	Resolution

`Link is down due to Unsupported Optic Installed`
Or
`Unsupported Optic`
	This message appears while scanning for targets or configuring boot devices.
	`Link is down due to Unsupported Optic Installed`
	In the Link Speed Selection menu
	`Unsupported Optic Installed`
	is displayed
	or
	In the Link Status field
	`Unsupported Optic`
	is displayed.
		Replace any unsupported optics with supported optics.

Not able to UEFI Secure Boot	UEFI Secure Boot fails	For the UEFI Secure Boot to function, the following requirements must be met: ■ Server hardware support must be enabled. ■ The operating system must support UEFI secure boot. ■ The UEFI driver must support UEFI secure boot (UEFI signed). ■ The operating system device driver must support UEFI secure boot (operating system specific). ■ The Common Information Model (CIM) provider must support UEFI secure boot (operating system specific).


BROADCOM®

BROADCOM®

Emulex® LightPulse® Adapter

Installation Guide

Version 11.2
December 30, 2016

Corporate Headquarters
San Jose, CA

Website
www.broadcom.com

Chapter 1: Introduction

This manual describes how to install an Emulex® LightPulse® adapter.

NOTE Illustrations in this manual are only examples. The actual hardware may vary.

Figure 1 Two-Port Adapter

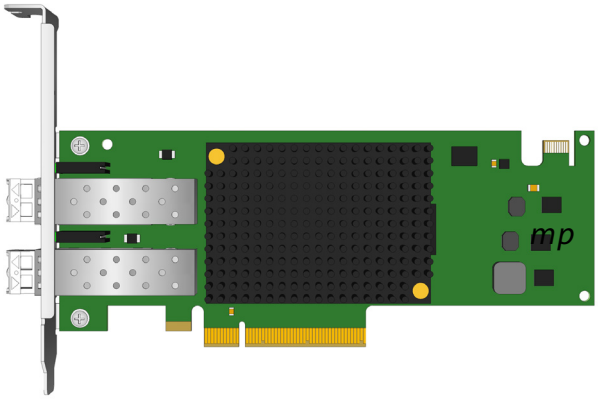

Figure 2 Single-Port Adapter

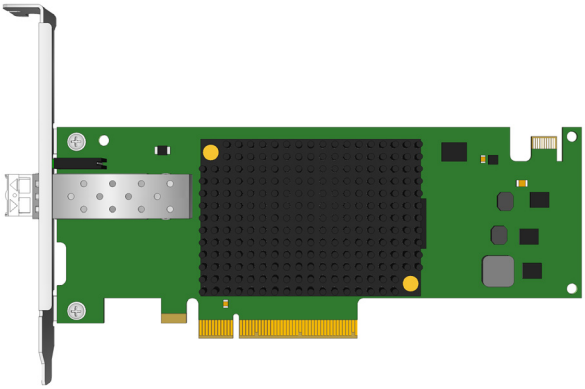

1.1 Adapter Identification

Each adapter ships with several numbers clearly marked on the board. Record these numbers before installation for later reference.

- **Institute of Electrical and Electronics Engineers, Inc. (IEEE) address** – a unique 64-bit identifier used for system configuration.
- **World Wide Name (WWN)** – derived from the IEEE address; the Fibre Channel (FC) industry uses the WWN for FC connectivity.
- **Serial number** – assigned by Broadcom® and used when communicating with Broadcom.

NOTE If the adapter has two ports, it has two IEEE addresses and two WWNs, one for each port.

Chapter 2: Installation: Preparing the Adapter for Installation

The Emulex LightPulse adapter can be connected to fiber optic cables (used with embedded optical transceivers). Additional information for these cable connections and other installation procedures are described in the following sections.

Before installing the adapter, perform one or both of the following procedures:

- Change the bracket to a low-profile version as explained in Section 2.1, Changing the Bracket.
- Disable or enable the Secure Firmware feature available on selected adapters using the instructions in Section 2.2, Enabling or Disabling the Secure Firmware Feature.

If changes to the adapter are not required, proceed to Section 2.3, Installing the Adapter.

2.1 Changing the Bracket

A standard, full-height Peripheral Component Interconnect Express (PCIe) bracket is factory-installed; however, a low-profile bracket is included in the box with any adapter that supports the low-profile bracket change. The low-profile mounting bracket is shorter than the standard bracket; approximately 3.11 in. (7.9 cm) compared to 4.75 in. (12.06 cm) long.

To change the adapter from a full-height to a low-profile bracket:

1. Some adapters come with optical transceivers embedded in their cage assemblies. These optical transceivers must be removed before the bracket can be removed. If the adapter contains optical transceivers, continue with the following steps; otherwise, proceed to step 4.

   CAUTION This is a delicate operation – take care not to damage the optical transceiver.

   The adapter uses different types of optical transceivers. Figure 3 shows an example of one type with the bail (handle) extended.

Figure 3 Typical Optical Transceiver

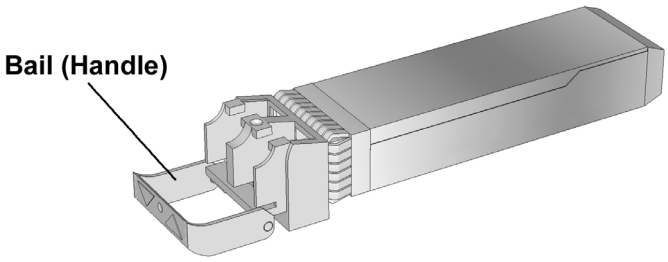

2. To remove a transceiver, pull the bail (handle) out and down to release the latch and gently pull the transceiver out. Do not force it. After the latch is released, the transceiver slides out easily.

   Figure 4 shows a transceiver with the latch released (bail extended) and another transceiver latched in place.

Figure 4 Releasing the Latch on an Optical Transceiver

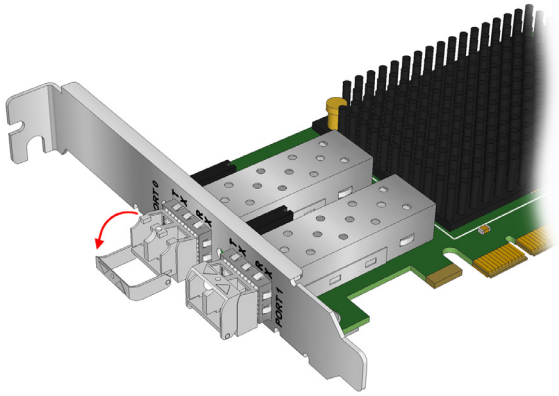

3. Using electrostatic discharge (ESD) precautions, store the transceiver in an ESD-safe place.
4. Remove the mounting bracket screws from the top of the adapter. Figure 5 shows the screws that are removed from the bracket.

Figure 5 Removing the Bracket

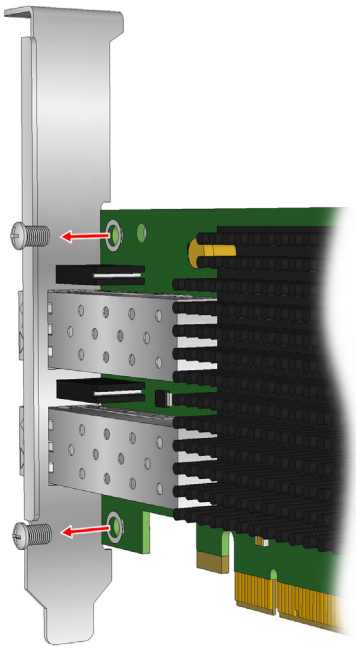

5. Remove the bracket and store it for future use.

6. Align the new mounting bracket tabs with the holes in the adapter.

   NOTE Be careful not to push the bracket past the transceiver or the RJ-45 housing's grounding tabs. Ensure the light emitting diodes (LEDs) are properly aligned with the holes in the bracket.

7. Reinstall the screws that attach the adapter to the bracket.
8. Reinstall the transceiver by sliding it into the housing. When the latch engages, it clicks.
9. Push the bail back into place.

2.2 Enabling or Disabling the Secure Firmware Feature

The LPe31000-series and LPe32000-series adapters have a jumper that enables or disables the Secure Firmware feature as required.

To update the firmware using a tool such as the OneCommand® Manager application, boot utilities, or Elxflash utilities, there is no need to remove the jumper; the Secure Firmware feature can remain enabled as long as the update is from a secure firmware version to another secure firmware version. To change the firmware version from a secure firmware version to an unsecured firmware version, you must remove the jumper before updating the firmware.

Replace the jumper after the firmware update is complete.

Figure 6 provides an example of the common location of the Secure Firmware jumper.

Figure 6 Secure Firmware Jumper Location J3

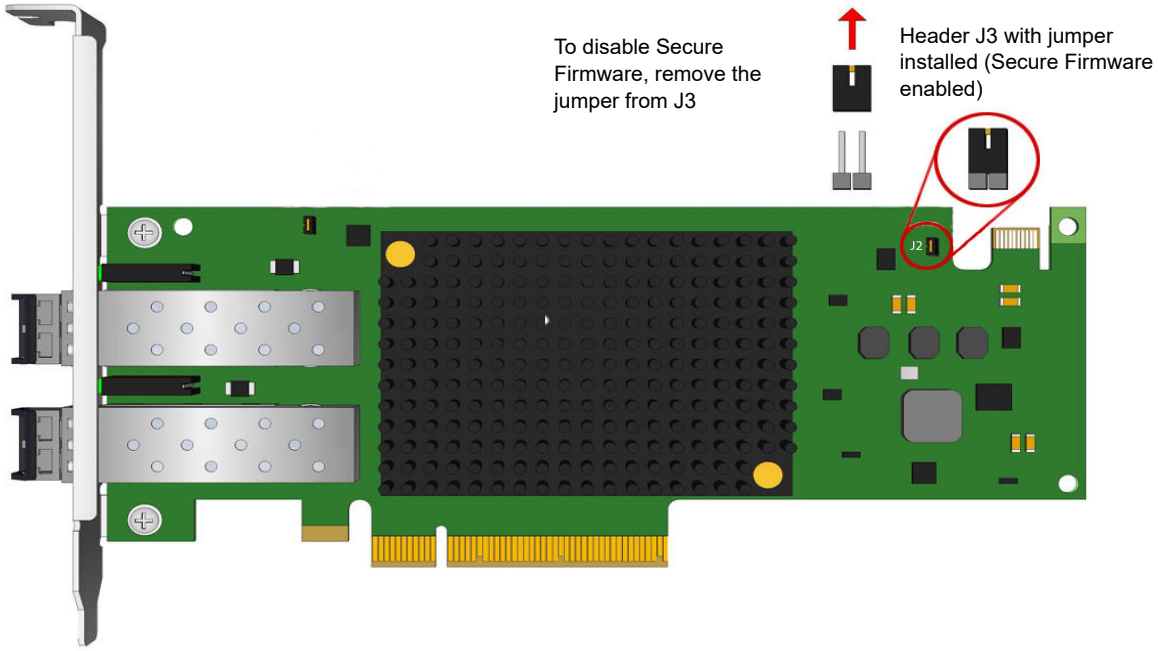

.

2.3 Installing the Adapter

To install the adapter:

1. Turn off and unplug the computer.
2. Remove the computer case.
3. For a stand-up adapter, remove the blank panel from an empty PCIe bus slot. See Section 2.1, Changing the Bracket to change the brackets if needed.

   NOTE To prevent damage to the adapter, select a slot that will not cause the adapter to interfere with the case.
4. Insert the adapter into the selected slot. Press firmly until the adapter is seated.
5. Secure the adapter's mounting bracket to the case with a panel screw or clip.
6. Replace the computer case and tighten the case screws.

The adapter is now installed in the server and is ready for device attachment.

2.4 Attaching Devices

You can connect fiber optic cables to the adapter.

An adapter does not allow normal data transmission on optical links unless it is connected to a compatible optical interface connection.

NOTE Use short range (SR), direct-connecting receptacle (DCR) multimode fiber cable for interface connections to a network.

To attach devices to the adapter:

1. Connect the appropriate cable to the adapter.
   When connecting an optical cable, ensure the cages have optical transceivers installed in them (see Figure 7).

Figure 7 Installing an Optical Transceiver

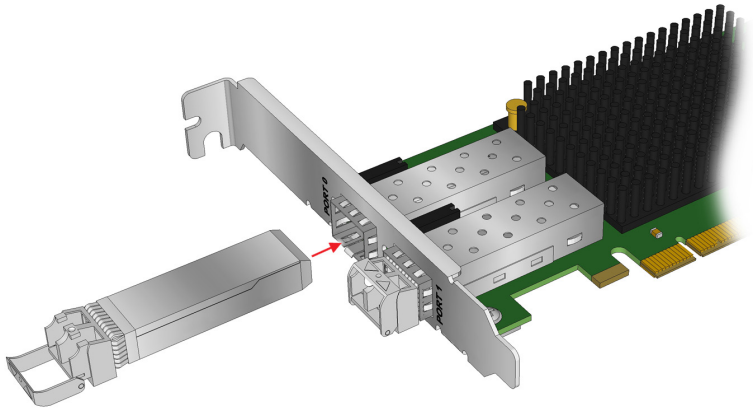

After the optical transceivers are installed, insert the optical cable into the LC connectors on the adapter (see Figure 8).

Figure 8 Connecting a Fiber Optic Cable

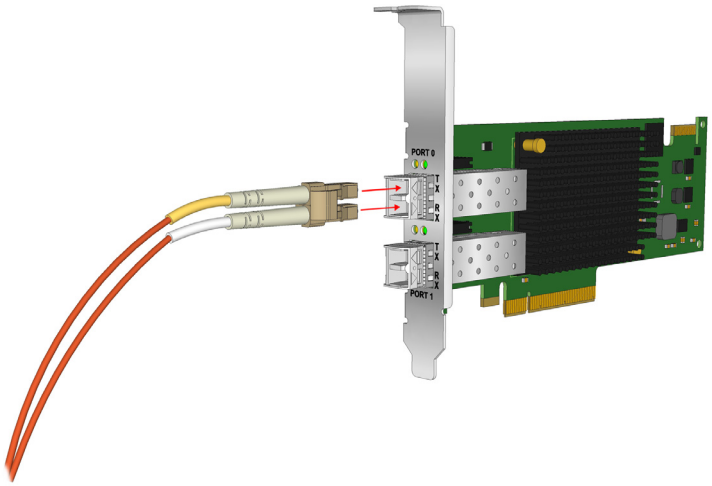

2.5 Applying Power

To apply power:

1. Verify that the adapter is securely installed in the computer.
2. Verify that the correct device is attached.
3. Plug in and turn on the computer.
4. Observe the boot banner for power-on self-test (POST) results.

2.6 Viewing the LEDs

You can view green and yellow LEDs through openings in the adapter's mounting bracket. The green LED indicates firmware operation and the yellow LED indicates port activity or link speed. Each port has a corresponding set of green and yellow LEDs.

Figure 9 LED Indicators

Table 1 summarizes LED activity.

NOTE For the Link Rate conditions, there is a 1 second pause when the LED is off between each group of fast blinks (2, 3, or 4). Observe the LED sequence for several seconds to be sure you have correctly identified the pattern.

Table 1 FC Adapter LED Activity

Green LED	Yellow LED	State
Off	Off	No SFP module installed or boot failure (dead board)
Off	On	POST failure (dead board)
Off	Slow blink	Boot failure after POST
Off	Flashing	POST processing in progress
On	Off	Failure in common code module
On	On	Failure in common code module
On	1 fast blink	Normal (link up at 2 gigabit Fibre Channel [GFC]) (legacy compatibility only)
On	2 fast blinks	Normal (link up at 4 GFC)
On	3 fast blinks 4 fast blinks (LPe12002-M8 adapters only)	Normal (Link up at 8 GFC)
On	4 fast blinks	Normal (link up at 16 GFC)
On	5 fast blinks	Normal (Link up at 32 GFC)
Slow blink	Off	Normal – link down
Fast blink	Fast blink	Beaconing


BROADCOM®