Emulex® Drivers for VMware ESXi

User Guide
Release 12.4

Table of Contents

Chapter 1: Introduction

This product supports the following Emulex® HBAs:

- LPe12000-series adapters
- LPe16000-series adapters
- LPe31000-series adapters
- LPe32000-series adapters
- LPe35000-series adapters

The VMware ESXi drivers support the FC protocol.

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 FC parameters:

- Adding LUNs and targets
- Configuring driver parameters
- Creating an FC remote boot disk
- Managing devices through the CIM interface
- Working with vPorts
- Configuring VVols
- Troubleshooting FC issues

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

1.1 ESXi Compatibility

For compatible firmware versions, go to the Documentation and Downloads page at 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 and subsequent releases
- RHEL 7.1 and subsequent releases
- RHEL 8.0
- SLES 11 and updates, SLES 12 and updates, and SLES 15 and updates
- XenServer Citrix 7.0
- OVM 3.3.3
- OVM 3.4.1
- Oracle UEK R4 and R5
- Oracle S10 and S11

- Windows Server 2012 and R2
- Windows Server 2016
- Windows Server 2019

1.3 Abbreviations

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

Table 1: Acronyms and Abbreviations

Acronym/Abbreviation	Description
CDB	Command Descriptor Block
CS_CTL	Class-Specific Control
CT	common transport
DA_ID	destination address identifier
DHCHAP	Diffie-Hellman Challenge Handshake Authentication Protocol
DID	Device ID
DIF	data integrity field
ELS	Extended Link Service
FA-PWWN	Fabric Assigned Port WWN
FC-GS	Fiber Channel Generic Services
FC-SP	Fibre Channel Security Protocols
FCF	Fibre Channel Forwarder
FCID	Fiber Channel Identifier
FCP	Fibre Channel Protocol
FDMI	Fabric-Device Management Interface
FLOGI	fabric login
I/O	input/output
IOCB	input/output control block
IOCTL	input/output control
IODM	I/O Device Management
IP CSUM	IP checksum
KV	key value
LBA	logical block address
LPFC	Light Pulse Fibre Channel
MB	megabyte
MSI	message signaled interrupt
MSI-X	message signaled interrupt-extended
N_Port	Node port
NFA	Network Flow Analysis
NPIV	N_Port ID virtualization
OB	Open Boot
OXID	Fibre Channel originator exchange
PE	protocol endpoint
PLOGI	port login
RSCN	registered state change notification

Table 1: Acronyms and Abbreviations (Continued)

Acronym/Abbreviation	Description
RSP	route/switch processor
VASA	vSphere APIs for Storage Awareness
VIB	vSphere Installation Bundle
VMID	Virtual Machine Identifier
VPD	vital product data
vPort	virtual port
VVols	virtual volumes
XRI	Extensible Resource Indicator

Chapter 2: Installation

This section provides information for installing the ESXi driver for the FC protocol.

2.1 Installing the FC Driver

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

2.1.1 Installing the FC 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 6.5 and subsequent operating systems, 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 User Guide* and the *Emulex OneCommand Manager Application for VMware vCenter 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 Installation Guide*.

NOTE: Before installing the management application, you must install the FC driver from the VMware website and then reboot the server.

Before installing the FC 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.

2.1.2 Uninstalling the FC Driver

Go to the VMware Patch Download page on the VMware website for instructions.

2.1.3 Installing the Native Mode FC Driver esxcli Plug-In

This section describes the installation and usage of the `esxcli` plug-in for the native mode FC (lpfc) driver on ESXi systems. This diagnostic tool is for the FC driver module.

NOTE: You can download the `esxcli` plug-in from the Broadcom® website.

To install the `esxcli elxfc` plug-in for ESXi 6.5 and ESXi 6.7, perform these steps:

1. Unzip and copy the following file to `/var/log/vmware/` on the ESXi host:

   `EMU-esx-6.x.x-elxlpfccli-<VIB version>-offline_bundle-<OS version>.zip`

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

   `esxcli software vib install -d /<pathname>/EMU-esx-6.x.x-elxlpfccli-<VIB version>-offline_bundle-<OS version>.zip`

3. Restart the `hostd` using the following command:

   `/etc/init.d/hostd restart`

   After the driver has restarted, the `elxfc` namespace appears under `esxcli`.

Supported commands:
`esxcli elxfc`

Usage:
`esxcli elxfc <cmd> <cmd options>`

Table 2: Available Namespaces

Namespace	Description
`adapter`	Adapter information of an Emulex FC HBA
`bginjerr`	Blockguard error injection page information of an Emulex FC HBA
`dhchap`	Authentication information of an Emulex FC adapter
`event`	Events information of an Emulex FC HBA
`eventsnap`	Events information with snap after display of an Emulex FC HBA
`fcf`	FCF table information of an Emulex FC HBA
`heap`	Heap usage statistics and information of an Emulex FC HBA
`iostat`	I/O statistics information of an Emulex FC HBA
`kv`	KV page information of an Emulex FC HBA
`logmessage`	Log message of an Emulex FC HBA in `vmkernel.log`
`lun`	SAN LUN information of an Emulex FC HBA
`memlog`	Memlog information of an Emulex FC HBA
`memory`	Memory statistics and information of an Emulex FC HBA
`mgmt`	Management statistics and information of an Emulex FC HBA
`modinfo`	Module information of an Emulex FC/FCoE HBA
`node`	Nodes in SAN information of an Emulex FC HBA
`param`	Dynamic configuration parameter information of an Emulex FC HBA
`port`	Port information and statistics of an Emulex FC HBA
`rsnn`	Symbolic node name information of an Emulex FC HBA
`sli4q`	sli4q information of an Emulex FC HBA
`target`	SAN targets information of an Emulex FC HBA
`vmid`	VMID page information of an Emulex FC HBA

Table 3: Available Commands

Command	Description
`list`	Lists Emulex FC adapter instance names.

2.2 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 to update the FC driver.

   `esxcli software vib install --maintenance-mode -d <pathname>/VMW-ESX-6.x.x-lpfc -<driver-version>-offline-bundle-<OS version>.zip`

   where <*driver-version*> represents the FC driver.

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

2.3 Installing the esxcli Management Tool

The esxcli management tool is delivered as an offline-bundle for the `esxcli` plug-in.

To install the esxcli management tool, perform the following steps:

1. Copy the offline-bundle to the ESXi host.
2. On the ESXi host, install the offline-bundle as follows:

   `esxcli software vib install -d /<pathname>/EMU-esx-6.x.x-elxmgmt-<VIB version> -offline_bundle-<OS version>.zip`

3. Restart the `hostd` using the following command:

   `/etc/init.d/hostd restart`

Chapter 3: Configuration

This chapter provides configuration information for the ESXi driver.

3.1 ESXi Command Line Tool Transition

For the ESXi 6.5 release and all subsequent ESXi 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 6.5 and ESXi 6.7 releases.

3.1.1 ESXi 6.5 and 6.7 Implementation

For ESXi 6.5 and subsequent ESXi operating systems, Broadcom uses the `esxcli` version of the command sequence, which has a different command syntax than the earlier releases. Some examples follow.

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 Driver Example

To set extended logging for the Emulex FC 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.

3.2 FC 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 section describes how to configure parameters using native ESXi tools. For a more comprehensive description of ESXi tools, go to the VMware public website. If you have further questions, contact a VMware technical support representative.

Refer to the *Emulex OneCommand Manager Application User Guide* and the *Emulex OneCommand Manager Command Line Interface User Guide* for information about the OneCommand Manager application.

Refer to the *Emulex OneCommand Manager for VMware vCenter User Guide* for information about the OneCommand Manager for VMware vCenter Server application.

3.2.1 FC 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. Refer to the *Emulex OneCommand Manager Application User Guide* or the *Emulex OneCommand Manager for VMware vCenter 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 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 3.2.6, FC 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 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 a VMware technical support representative.

The 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 adapters in your system.

1. Locate the parameter `lpfc_lun_queue_depth` in Table 4.
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 FC-SP-2 Authentication (DHCHAP) Support

NOTE:

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

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

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

FC-SP-2 DHCHAP authentication is disabled by default. To enable FC-SP-2 DHCHAP authentication, you must pass the `lpfc_enable_auth` parameter to the driver by typing the following command:

```
# esxcli system module parameters set -p lpfc_enable_auth=1 -m lpfc
```

You must reboot the system to activate the parameter.

After DHCHAP has been activated and configured, you can manually initiate authentication per adapter. For more information on manually initiating authentication, refer to the *Emulex OneCommand Manager Application User Guide*.

You can get, set, update, or delete a password using the OneCommand Manager application. For details on how to perform these operations, refer to the *Emulex OneCommand Manager Application User Guide*.

You can also display authentication configuration using the `keyval` page and `esxcli`.

For example, in the `keyval` page, type:

```
/usr/lib/vmware/vmkmgmt_keyval/vmkmgmt_keyval  -i vmhba4/Emulex  -g -k   auth
```

An output similar to the following is shown:

```
Key 'auth':     Authentication info: Authentication enabled
local WWPN 10:00:00:90:fa:c7:c2:a8      remote WWPN ff:ff:ff:ff:ff:ff:ff:ff
```

`ff:ff:ff:ff:ff:ff:ff:ff` is the switch port name.

3.2.3 Trunking Support

NOTE:

- Before you enable or disable trunking on the Emulex adapter, follow the instructions from Brocade to enable or disable trunking on the switch.
- Dynamic D_Port cannot co-exist with the trunking feature (also called FC port aggregation) on LPe35000-series adapters. If trunking is enabled, the firmware automatically disables Dynamic D_Port.
- FA-PWWN and Dynamic D_Port cannot co-exist with the trunking feature on LPe35000-series adapters. If trunking is enabled, the firmware automatically disables FA-PWWN and Dynamic D_Port.
- Trunking is only supported on LPe35000-series adapters.
- Trunking is not supported at 8 Gb/s speeds, and the link will not come up at this speed.

Trunking enables you to combine multiple physical FC links to form a single logical link (aggregated port). The aggregated port's maximum link speed is the sum of the maximum link speeds of the individual physical links comprising the aggregated port. For example, an aggregated port comprised of two physical links running at 32 Gb/s each will have a potential logical (aggregate) link speed of 64 Gb/s. The actual link speed of the aggregated port depends on the states (active or non-active) of the individual physical links comprising the aggregated port.

The physical links comprising an aggregated port are referred to as lanes. Only 2-lane and 4-lane aggregated ports are supported. For LPe35002 adapters, only 2-lane port aggregation is possible. If 2-lane port aggregation is configured on an LPe35002 adapter, the two physical links are combined to form a single 2-lane aggregated port whose aggregate speed is potentially 64 Gb/s, assuming both physical links are active.

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

There are no driver parameters to enable trunking. Trunking is a one-time configuration, and the configuration is performed through the OneCommand Manager application. For more information on configuring trunking, refer to the *Emulex OneCommand Manager Application User Guide*.

When trunking is configured for an adapter, Internal loopback and External loopback tests are supported. For more information on running loopback tests, refer to the *Emulex OneCommand Manager Application User Guide*.

You can view the existing trunk links using `esxcli elxmgmt`.

An example of an LPe35002 adapter, showing 2-lane port aggregation with a single SCSI host follows:

```
# esxcli elxmgmt trunkinfo get -w 10:00:00:10:9b:41:73:8c

Trunking Mode: 2-lane trunking
Trunk Port 0: Logical Link Speed: 64 Gbs
  Physical Link 0: Active
  Physical Link 1: Active
```

An example of an LPe35004 adapter, showing 2-lane port aggregation with a single SCSI host follows.

```
# esxcli elxmgmt trunkinfo get -w 10:00:00:10:9b:65:dd:24

Trunking Mode: 2-lane trunking

Trunk Port 0: Logical Link Speed: 64 Gbs
  Physical Link 0: Active
  Physical Link 1: Active

Trunk Port 2: Logical Link Speed: 64 Gbs
  Physical Link 2: Active
  Physical Link 3: Active
```

An example of an LPe35004 adapter, showing 4-lane port aggregation with a single SCSI host follows.

```
# esxcli elxmgmt trunkinfo get -w 10:00:00:10:9b:65:dd:24

Trunking Mode: 4-lane trunking

Trunk Port 0: Logical Link Speed: 128 Gbs
  Physical Link 0: Active
  Physical Link 1: Active
  Physical Link 2: Active
  Physical Link 3: Active
```

3.2.4 Dynamically Adding LUNs

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

3.2.5 Dynamically Adding Targets

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

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

To force the ESXi server to rescan all devices, perform one of these actions:

- Run the following command:
  `esxcli storage core adapter rescan -A vmhbaX`
- From the vSphere Client, select **Configuration Tab** > **Storage Adapters**, and then click **Rescan All**.

3.2.6 FC Driver Configuration Parameters

Table 4 lists the FC 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 4: FC Driver Parameters

Module Parameter	Description	ESXi 6.5 and ESXi 6.7 Native Mode Driver Model Values	Comments
`lpfc_compression_log`	Defines how often the compression logs are written (in seconds). This parameter is dynamically updated.	Default = 300 Min. = 5 Max.= 86,400	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 the throttle stops the driver from spamming the logs on a very high frequency failure.
`lpfc_delay_discovery`	Delays N_Port discovery when the Clean Address bit is cleared. This parameter requires a system reboot.	Default = 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 port name is changed.
`lpfc_devloss_tmo`	The number of seconds the driver holds I/O waiting for a lost device to return. This parameter is dynamically updated.	Default = 10 Min. = 1 Max. = 255	—
`lpfc_discovery_threads`	The maximum number of ELS commands that can be outstanding during discovery. This parameter requires a system reboot.	Default = 32 Min. = 1 Max. = 64	—
`lpfc_enable_auth`	Enables or disables DHCHAP authentication. The possible values are: ■ 0 = Authentication is disabled (default). ■ 1 = Authentication is enabled. This parameter is dynamically updated.	Default 0 = Disable Min. 0 = Disable Max. 1 = Enable	**NOTE:** DHCHAP and Dynamic D_Port are mutually exclusive. If both features are enabled simultaneously, warning messages will result. You must either disable DHCHAP using the `lpfc_enable_auth` parameter, or disable Dynamic D_Port using the OneCommand Manager application.

Table 4: FC Driver Parameters (Continued)

Module Parameter	Description	ESXi 6.5 and ESXi 6.7 Native Mode Driver Model Values	Comments
`lpfc_enable_bb_credit_recovery`	Toggles the Buffer-to-Buffer Credit Recovery feature. This parameter requires a system reboot.	Default = 1 Min. = 0 Max. = 1	Applicable to SLI-4 adapters only.
`lpfc_enable_da_id`	Controls the DA_ID CT command. Unregisters objects with the fabric nameserver.	Default 1 = Enable Min. 0 = Disable Max. 1 = Enable	—
`lpfc_enable_fcp_priority`	Enables or disables FCP priority. This parameter requires a system reboot.	Default = 0 Min. = 0 Max. = 1	■ 0 = Do not use the nodelist priority table (default). ■ 1 = Use the nodelist priority table.
`lpfc_enable_mds_diags`	Enables MDS diagnostics. This parameter is dynamically updated.	0 = Disabled (default) 1 = Enabled	Enables or disables the Cisco Fibre Channel Link Diagnostics feature. **NOTE:** The parameter should be disabled (set to 0) after diagnostics are complete.
`lpfc_enable_qfull`	Enable the driver's `lun_queue_depth` ramp down and ramp 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 I/O successfully completes. ■ When disabled (value = 0), the driver takes no action when an I/O completes with a Task Set Full status. This parameter is dynamically updated.	Default = 1 Min. = 0 Max. = 1	Disabling this parameter has a potential impact on the overall VM performance. Consult your target vendor for guidance.

Table 4: FC Driver Parameters (Continued)

Module Parameter	Description	ESXi 6.5 and ESXi 6.7 Native Mode Driver Model Values	Comments
`lpfc_enable_rrq`	Enables Reinstate Recovery Qualifier functionality. ■ 0x0 = Disabled, XRI/OXID use not tracked. ■ 0x1 = XRI/OXID reuse is timed with the 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. This parameter requires a system reboot.	Default = 2 Min. = 0 Max. = 2	—
`lpfc_enable_SmartSAN`	Enables SmartSAN functionality. The function works with FDMI-2 to provide enhanced fabric diagnostics. This parameter requires a system reboot.	Default = 0 disable Min. = 0 disable Max. =1 enable	—
`lpfc_external_dif`	Enables external DIF support on select devices. This parameter requires a system reboot.	0 = Disabled 1 = Enabled (default)	External DIF is not available for LPe12000-series adapters.
`lpfc_fcp_class`	Selects the FC class of service for FCP sequences. This parameter requires a system reboot.	Default = 3 Min. = 2 Max. = 3	—
`lpfc_fdmi_on`	Controls FDMI support. ■ 0 = FDMI support off. ■ 1 = FDMI support on. This parameter requires a system reboot.	Default = 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.
`lpfc_first_burst_size`	First burst size for targets that support first burst. This parameter is dynamically updated.	Default = 0 Min. = 0 Max. = 65,536	—

Table 4: FC Driver Parameters (Continued)

Module Parameter	Description	ESXi 6.5 and ESXi 6.7 Native Mode Driver Model Values	Comments
`lpfc_hba_queue_depth`	The maximum number of FCP commands queued to an FC adapter. The driver automatically adjusts the `hba_queue_depth` to match adapter capabilities. This setting may be overridden. This parameter requires a system reboot.	Default = 8,192 Min. = 32 Max. = 8,192	—
`lpfc_iocb_cnt`	IOCBs allocated for ELS, CT, and abort sequence in 1024 increments. This parameter requires a system reboot.	Default = 1 Min. = 1 Max. = 5	—
`lpfc_link_speed`	Selects link speed. Valid values are: ■ 0 = Auto select ■ 2 = 2 Gigabaud ■ 4 = 4 Gigabaud ■ 8 = 8 Gigabaud ■ 16 = 16 Gigabaud ■ 32 = 32 Gigabaud This parameter is dynamically updated.	Default = 0 Min. = 0 Max. = 32	For ESXi 6.5 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.
`lpfc_log_verbose`	Verbose logging bit-mask. This parameter is dynamically updated.	Default = 0 Min. = 0 Max. = 0x7fffffff	—
`lpfc_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. This parameter is dynamically updated.	Default = 30 Min. = 1 Max. = 254	—

Table 4: FC Driver Parameters (Continued)

Module Parameter	Description	ESXi 6.5 and ESXi 6.7 Native Mode Driver Model Values	Comments
`lpfc_max_heap_size`	Maximum allowable memory consumption per server for the LPFC module. This parameter requires a system reboot.	Default = 128 MB Min. = 64 MB Max. = 512 MB	—
`lpfc_max_luns`	The maximum number of LUNs allowed. This parameter requires a system reboot.	Default = 65,535 Min. = 1 Max. = 65,535	Setting in ESXi 6.5 and 6.7 allows for sparse LUN IDs above 256.
`lpfc_max_multiq`	Sets 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. This parameter requires a system reboot.	Default = 4 Min. = 0 Max. = 8	Sets the number of MultiQ FCP I/O 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_max_scsicmpl_time`	Uses the 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. This parameter is dynamically updated.	Default = 0 Min. = 0 Max. = 60,000	—
`lpfc_max_targets`	The maximum number of discovered targets allowed. This parameter requires a system reboot.	Default = 256 Min. = 0 Max. = 4096	A driver parameter to adjust the supported target count.
`lpfc_max_vmid`	Maximum number of VMs to be tagged. The range is 4 to 255. This parameter requires a system reboot.	Default = 8 Min. = 4 Max. = 255	This value indicates the number of VMIDs supported.
`lpfc_mem_tracker`	Turns memory tracking on or off. Valid values are: ■ 1 = On ■ 0 = Off This parameter requires a system reboot.	0 = Disabled (default) 1 = Enabled	Keeps track of driver heap and slab allocation/free. It is used only when debugging driver memory issues.

Table 4: FC Driver Parameters (Continued)

Module Parameter	Description	ESXi 6.5 and ESXi 6.7 Native Mode Driver Model Values	Comments
`lpfc_nlp_slab_cnt`	NLP slab entries. This parameter requires a system reboot.	Default = 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.
`lpfc_ras_fwlog_buffsize`	Buffer size to be allocated in host memory for firmware logging. This parameter requires a system reboot.	Default = 0 Min. = 0 Max. = 4 The possible values are: ■ 0 = Disable firmware log written to the host memory (default) ■ 1 = ¼ MB (256K) host memory to be allocated for the firmware log ■ 2 = ½ MB (512K) host memory to be allocated for the firmware log ■ 3 = ¾ MB (768K) host memory to be allocated for the firmware log ■ 4 = 1 MB host memory to be allocated for the firmware log	This parameter is supported only on LPe31000-series, LPe32000-series, and LPe35000-series adapters.
`lpfc_ras_fwlog_level`	Firmware logging level. Valid only if firmware logging is enabled. This parameter requires a system reboot.	Default = 0 Min. = 0 Max. = 4	This parameter is supported only on LPe31000-series, LPe32000-series, and LPe35000-series adapters.
`lpfc_rb_slab_cnt`	Receives buffer slab entries. This parameter requires a system reboot.	Default = 256 Min. = 32 Max. = 256	Controls the maximum number of receive buffers that will be posted to the adapter.
`lpfc_scan_down`	Start scanning for devices from the highest AL_PA to the lowest. This parameter requires a system reboot.	Default = 1 Min. = 0 Max. = 1	—

Table 4: FC Driver Parameters (Continued)

Module Parameter	Description	ESXi 6.5 and ESXi 6.7 Native Mode Driver Model Values	Comments
`lpfc_sg_seg_cnt`	The maximum scatter gather segment count for DMA. The maximum data allowed in one SG element is 0x80000000. This parameter requires a system reboot.	Default = 64 Min. = 64 Max. = 4096	—
`lpfc_suppress_link_up`	Suppresses link up at initialization: ■ 0x0 = Bring up link ■ 0x1 = Do not bring up link ■ 0x2 = Never bring up link This parameter requires a system reboot.	Default = 0 Min. = 0 Max. = 2	Enable this parameter to assist with SAN issues during ESXi boot.
`lpfc_suppress_rsp`	Negotiates whether to suppress the RSP from a target during a PLOGI. This parameter is dynamically updated.	Default = 1 Min. = 0 Max. = 1	Applicable to SLI-4 adapters only.
`lpfc_task_mgmt_tmo`	The maximum time to wait for task management commands to complete. This parameter is dynamically updated.	Default = 60 Min. = 5 Max. = 180	—
`lpfc_tgt_queue_depth`	The maximum number of FCP commands queued to a specific target port. This parameter is dynamically updated.	Default = 65,535 Min. = 10 Max. = 65,535	—
`lpfc_throttle_log_cnt`	Do not exceed this number of messages logged within `throttle_log_time`. This parameter is dynamically updated.	Default = 10 Minimum = 1 Maximum = 1000	Logging mechanism intended to speed up issue diagnosis by reducing the need to enable driver logging.
`lpfc_throttle_log_time`	Do not exceed `throttle_log_cnt` number of logs within this time limit (seconds). This parameter is dynamically updated.	Default = 1 Min. = 1 Max. = 60	Works with `throttle_log_cnt`.

Table 4: FC Driver Parameters (Continued)

Module Parameter	Description	ESXi 6.5 and ESXi 6.7 Native Mode Driver Model Values	Comments
`lpfc_topology`	Selects 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 This parameter is dynamically updated.	Default = 0 Min. = 0 Max. = 6	For LPe31000-series, LPe32000-series, and LPe35000-series adapters' topology cannot be changed and is set to point-to-point.
`lpfc_use_adisc`	Use address discovery on rediscovery, initiated by RSCN, to authenticate FCP devices instead of port login. This parameter is dynamically updated.	Default = 0 Min. = 0 Max. = 1	—
`lpfc_use_msi`	Uses preferred MSI-X interrupt mode, if possible. ■ 0 = MSI disabled (INTx mode) ■ 1 = MSI enabled ■ 2 = MSI-X enabled This parameter requires a system reboot.	Default = 2 Min. = 0 Max. = 2	—
`lpfc_vmid_app_header`	VMID Application ID tagging. This parameter requires a system reboot.	0 = Disabled (default) 1 = Enabled	Brocade® VMID uses the application service header field to communicate the capability.
`lpfc_vmid_inactivity_timeout`	Inactivity timeout duration in hours. The range is 0 to 24. This parameter is dynamically updated.	Default = 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.

Table 4: FC Driver Parameters (Continued)

Module Parameter	Description	ESXi 6.5 and ESXi 6.7 Native Mode Driver Model Values	Comments
`lpfc_vmid_priority_tagging`	VMID CS_CTL tagging. This parameter requires a system reboot.	Default = 0 Min. = 0 Max. = 2	Cisco VMID uses the priority tagging field to communicate the capability. ■ 0 = Disables priority tagging. Priority tagging is also disabled when the parameter is set to anything other than 1 or 2. ■ 1 = Enables priority tagging only for targets that support it in their PLOGI LS_ACC response. ■ 2 = Enables priority tagging for all targets, whether they support it in their PLOGI LS_ACC response.

NOTE: The values in Table 5 and Table 6 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 5: 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 6: Adapter Attributes for fdmi_on Parameter

Hexadecimal Value	Information Type
0x001	Node Name
0x002	Manufacturer
0x003	Serial Number
0x004	Model

Table 6: Adapter Attributes for fdmi_on Parameter (Continued)

Hexadecimal Value	Information Type
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

3.2.7 Creating an FC Remote Boot Disk

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

3.2.8 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.

3.2.8.1 Using the OneCommand Manager GUI

For VMware ESXi 6.5 and 6.7 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.

When Windows OneCommand Manager is used to manage HBAs installed on ESXi 6.7 Update 1, most of the features are unavailable. This is due to recent changes made by VMware to its SFCB daemon.

The following table lists the features that are functional and nonfunctional when Windows OneCommand Manager is used to manage HBAs discovered on ESXi 6.7 Update 1.

Functional	Nonfunctional
Discovery	SetPortSpeed
Port Attributes	SetPortEnabled
Hba Attributes	EnableBootCode
Port Statistics	Diagnostics
Server Attributes	WWN Management
DriverParams	Dump
GetXCVRdata	Get IPL File Name
Pcidata	Firmware Download
VPD	Firmware Parameters
	Trunking

Instead of OneCommand Manager for Windows, use the OneCommand Manager for VMware CLI (`elxvcpcmd`) or the esxcli elxmgmt utilities to manage HBAs installed on ESXi 6.7 Update 1.

3.2.8.2 Using the OneCommand Manager Application for VMware vCenter

The OneCommand Manager for VMware vCenter Server application uses the standard CIM interface to manage the adapters and supports CIM-based device and adapter management. The OneCommand Manager for VMware vCenter Server application also supports existing adapter management functionality based on its proprietary management stack and the standard HBA API interface. To manage adapters (including updating the firmware) on an ESXi 6.5 or 6.7 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 User Guide*.

NOTE:

- For the ESXi 6.7 vSphere Web client, flex is not supported, so you must use the vSphere client, HTML. Currently the vSphere client, HTML supports only limited functionality.
- For LPe35000-series adapters only:

  In some cases, a firmware update requires either a firmware reset or a basic PCIe reset, depending on the features available in the new firmware. A firmware reset is performed automatically if it is needed, regardless of the operating system. In some cases, a full reboot is required to activate new firmware or to enable a new feature. In that case, a message similar to one of the following messages appears after the firmware download is complete:

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

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

  For a list of features that require a reboot to be enabled, refer to the *Emulex LPe35000-Series HBA Firmware and Boot Code Release Notes*.

3.2.9 Installing the Emulex CIM Provider

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

3.2.10 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 vPorts link state is offline. VMware ESXi supports fabric mode only.
- You can create vPorts only on 8, 16, and 32 GFC adapters.
- The OneCommand Manager application sees all vPorts created by the driver, but the application has read-only access to them.

3.2.11 Configuring VVols in ESXi 6.5 and 6.7

The Emulex native mode FC driver supports the VVols feature released with ESXi 6.5 and 6.7. 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 ESXi to the array is managed through an access point or PE and the storage provider.

3.2.11.1 Storage Containers

Storage containers are a logical abstraction and hold groups of VVols that are physically provisioned 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.

3.2.11.2 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.

3.2.11.3 Storage Providers

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

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

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 Driver

This section provides troubleshooting information for the FC driver.

Table 7 identifies some of the common situations and their potential resolutions.

Table 7: Troubleshooting the FC Driver

Situation	Resolution
The 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
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, determine 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 Driver Configuration), make sure you spelled the driver parameters correctly. The ESXi module subsystem does 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, and whether it has found a fabric and the link speed.[a] 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 shows if the node status is changing or if the node is going offline.[b] 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.
The FC 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 www.broadcom.com.

Table 7: Troubleshooting the FC Driver (Continued)

Situation	Resolution
The system panics when booted with a failed adapter installed.	Remove the failed adapter and reboot.

The FC 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 are there 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 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.

a. Data from KV command 1:

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


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


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

SLI Rev: 3
      MQ: Unavailable
      NPIV Supported: VPIs max 255  VPIs used 0
      RPIs max 4096  RPIs used 18   IOCBs inuse 0  IOCB max 16   txq cnt 0  txq max 0  txcmplq 0
      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
```

b. 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         xffffffd   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 Driver Log Messages

Log messages have traditionally been organized into logical groups based on code functionality in the FC driver. With the introduction of the latest 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 view 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 8 lists the groups and defines the associated number ranges.

Table 8: Message Log Table

LOG Message Verbose Mask Definition	Verbose Bit	Verbose Description
LOG_ELS	0x00000001	ELS events
LOG_DISCOVERY	0x00000002	Link discovery events
LOG_MBOX	0x00000004	Mailbox events
LOG_INIT	0x00000008	Initialization events
LOG_LINK_EVENT	0x00000010	Link events
LOG_MGMT_ERROR	0x00000020	IODM management error logging
LOG_FCP	0x00000040	FCP traffic history
LOG_NODE	0x00000080	Node table events
LOG_CGN_MGMT	0x00000100	Congestion management events
LOG_BG	0x00000200	BlockGuard events
LOG_MEM_HEAP	0x00000400	Memory tracker heap logging

Table 8: Message Log Table

LOG Message Verbose Mask Definition	Verbose Bit	Verbose Description
LOG_SLI	0x00000800	SLI events
LOG_FCP_ERROR	0x00001000	Log errors, not underruns
LOG_LIBDFC	0x00002000	Libdfc events
LOG_VPORT	0x00004000	NPIV events
LOG_MEM_SLAB	0x00008000	Memory tracker slab logging
LOG_EVENT	0x00010000	CT, TEMP, DUMP, logging
LOG_AUTH	0x00020000	DHCHAP logging
LOG_FCP_UNDER	0x00040000	FCP underrun errors
LOG_TASKMGMT	0x00100000	Task management events.
LOG_MGMT	0x00200000	IODM management trace logging
LOG_SCSI_CMD	0x00400000	ALL SCSI commands
LOG_EDIF	0x00800000	External DIF events
LOG_ALL_MSG	0x7fffffff	Log all messages

The following is an example of a LOG message:

```
lpfc_sli4_read_config:8989: 0:2003 cfg params Extents
```

In the preceding LOG message:

- `lpfc` – Identifies the driver binary
- `sli4_read_config` – Identifies the function generating the log
- `0` – Identifies the Emulex port number
- `2003` – 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`.

Appendix A: esxcli Management Tool

A.1 Usage

The esxcli management tool supports both local and remote management.

- Local management

  `esxcli elxmgmt <Command> <Parameters>`

- Remote management

  To manage the ESXi system directly, use the following command:

  `esxcli -s <server> -u <username> -p <password> -d <Thumbprint> elxmgmt <Command> <Parameters>`

  To manage the ESXi system using a vCenter server, use the following command:

  `esxcli -s <vCenter Server> -u <username> -p <password> -h <VI_HOST> -d <thumbprint> elxmgmt <Command> <Parameters>`

Table 9: Available Namespaces

Namespace	Description
`authconfig`	Authentication features of an Emulex FC HBA.
`driverparams`	Driver parameters of an Emulex FC HBA.
`hbaport`	Port associated features of an Emulex FC HBA.
`beacon`	Beacon state of an Emulex FC HBA.
`bootparams`	Boot parameters for a given port and given boot type for an Emulex FC HBA.
`congestparams`	Congestion management of an Emulex FC HBA.
`dumpdirectory`	Dump file directory for an Emulex FC HBA in the host.
`fwlog`	Firmware logging for a particular port of an Emulex FC HBA.
`fwparams`	Firmware parameters of an Emulex FC HBA.
`trunkinfo`	Trunking operations for an LPe35000-series adapter.
`xcvrdata`	Transceiver data of an Emulex FC HBA.

A.2 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.

The following table contains a list of esxcli management commands and their descriptions.

Table 10: esxcli Management Commands

Command	Description
`allnodeinfo -w <wwpn>`	Shows the target node information for each target accessible by the adapter.
`authconfig delete -w <initiator wwpn>`	Sets delete authentication for a specified HBA port. **NOTE:** This command replaces the `removeadapterauthconfig` command.

Table 10: esxcli Management Commands (Continued)

Command	Description
`authconfig get -w <initiator wwpn> -d <target wwpn>`	Shows the authentication configuration details. ■ d specifies the destination WWPN. It must be ff:ff:ff:ff:ff:ff:ff:ff for a switch or the actual WWPN for a target. **NOTE:** This command replaces the `getauthconfig` command.
`authconfig initiate -w <initiator wwpn> -d <target wwpn>`	Initiates authentication for a specified HBA port. ■ d specifies the destination WWPN. It must be ff:ff:ff:ff:ff:ff:ff:ff for a switch or the actual WWPN for a target. **NOTE:** This command replaces the `initiateauth` command.
`authconfig lst -w <wwpn>`	Shows the authentication configuration list. **NOTE:** This command replaces the `authconfiglist` command.
`authconfig remove -w <wwpn> -e <entity pair>`	Sets delete authentication from an entity pair for a specified HBA port. ■ e specifies the entity pair: (`localEntity`, `remoteEntity`) – `localEntity`: Source WWPN. – `remoteEntity`: Destination WWPN or all. – `all` deletes the entire authentication configuration. **NOTE:** This command is not supported on LPe12000-series adapters. **NOTE:** This command replaces the `removeauthconfig` command.
`authconfig secret set -w <initiator wwpn> -d <target wwpn> -f <flag> -t <new secret type> -v <new secret value>`	Sets the authentication configuration password for a specified HBA port. ■ d specifies the destination WWPN. It must be ff:ff:ff:ff:ff:ff:ff:ff for a switch or the actual WWPN for a target. ■ `f` specifies flag. Valid values are (1, 2). – 1 = Local; password used when the HBA initiates authentication to the switch and when using bidirectional authentication. – 2 = Remote; password used when the switch initiates authentication to the HBA and when using bidirectional authentication (required). ■ `t` specifies the new secret type. Valid values are (1, 2). – 1 = ASCII. – 2 = Hexadecimal (binary). ■ v specifies the new secret value. **NOTE:** This command replaces the `setauthconfigsecret` command.

Table 10: esxcli Management Commands (Continued)

Command	Description
`authconfig set -w <initiator wwpn> -d <target wwpn> -b <bi-directional> -p <dh-priority> -s <hash-priority> -m <mode> -r <re-authentication> -a <re-authentication-interval> -t <timeout>`	Sets the authentication configuration parameters for a specified HBA port. ■ d specifies the destination WWPN. It must be ff:ff:ff:ff:ff:ff:ff:ff for a switch or the actual WWPN for a target. ■ b specifies bidirectional. Valid values are (disabled, enabled). ■ p specifies dh-priority: (1, 2, 3, 4, 5), any combinations up to five digits. (that is 4321 specifies a priority of 1536:1280:1024:Null) 1 = Null 2 = 1024 3 = 1280 4 = 1536 5 = 2048 ■ s specifies hash-priority. Valid values are (md5, sha1). md5 = First md5, then sha1. sha1 = First sha1, then md5. ■ m specifies the mode. Valid values are (disabled, enabled, passive). ■ r specifies re-authentication mode. Valid values are (disabled, enabled). ■ a specifies re-authentication interval. Valid values are (0, 10 to 3,600), time in minutes. 0 = Disables re-authentication. ■ t specifies timeout value, time in seconds. **NOTE:** This command replaces the `setauthconfigparams` command.
`authconfig status get -w <initiator wwpn> -d <target wwpn>`	Shows the authentication configuration status details. ■ d specifies the destination WWPN. It must be ff:ff:ff:ff:ff:ff:ff:ff for a switch or the actual WWPN for a target. **NOTE:** This command replaces the `getauthstatus` command.
`beacon get -w <wwpn>`	Shows the beacon state of an Emulex FC HBA. **NOTE:** This command replaces the `getbeacon` command.
`beacon set -w <wwpn> -s <state> [-d <duration>]`	Sets the beacon state of an Emulex FC HBA. ■ s specifies the state of the beacon. Default is 1. Valid values are: – 1 = Beacon on (mandatory). – 0 = Beacon off. ■ d specifies the duration of beaconing. Maximum duration is 65,535 seconds. The default is 0 for infinite beaconing. This is an optional parameter. **NOTE:** This command replaces the `setbeacon` command.
`bootparams get -w <wwpn> -t <boot type>`	Shows the boot parameters for a given port and given boot type. ■ t specifies the boot type: {X86, EFI, OB}. **NOTE:** This command replaces the `getbootparams` command.

Table 10: esxcli Management Commands (Continued)

Command	Description
`bootparams set -w <wwpn> -p <parameter> -t <type> -v <value> - b <BootDev>`	Sets the boot parameter and boot type for the specified HBA port. ■ p specifies the parameter name. ■ t specifies the boot type: {X86, EFI, OB}. ■ v specifies the value to be set for the parameter. Parameter name, supported boot type, and valid values are as follows: – DefaultAlpa X86/OB { Value } – EnableAdapterBoot All { State } (0=Disable, 1=Enable) – EnableBootFromSan All { State } (0=Disable, 1=Enable) – LinkSpeed All { 0, 1, 2, 4, 8 } – 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)

Table 10: esxcli Management Commands (Continued)

Command	Description
	– 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, 3, 4 } – EnableFABL EFI { State } (0=Disable, 1=Enable) – MaxLunsPerTarget EFI { Value } – DelayDeviceDiscovery 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 } ■ b specifies the boot device entry number { 0 -7 } as shown in the output of the `bootparams get` command. This command replaces the `setbootparam` command.
`congestparams get -w <wwpn> -a <all>`	Shows the congestion management settings for a specified HBA port. ■ a specifies all (optional, advanced). Contact Broadcom Technical Support before you use this option.

Table 10: esxcli Management Commands (Continued)

Command	Description
`congestparams set -w <wwpn> -v <congestion parameter value> -p <congestion parameter name>[-s <burst parameter name>] [-r <burst size value>] [-b <bandwidth parameter name>] [-d <bandwidth value>]`	Sets the congestion parameters for a specified HBA port. ■ `v` specifies the congestion parameter value. ■ `p` specifies the congestion parameter name. For congestion parameter CGN-MODE, valid values are: – 0 = Off; congestion management is disabled. – 1 = Fixed; contact Broadcom Technical Support before you use this option. – 2 = Adaptive-Conservative; a minimum level of congestion management is performed. – 3 = Adaptive-Moderate; a medium level of congestion management is performed. – 4 = Adaptive-Aggressive; a strong level of congestion management is performed ■ `s` specifies the burst size parameter name CGN-BS (optional, advanced). Contact Broadcom Technical Support before you use this option. ■ `r` specifies the burst size value (optional, advanced). Contact Broadcom Technical Support before you use this option. ■ `b` specifies the bandwidth parameter name CGN-BW (optional, advanced), Contact Broadcom Technical Support before you use this option. ■ `d` specifies the bandwidth value (optional, advanced). Contact Broadcom Technical Support before you use this option.
`download -w <wwpn> -f <firmware pathname>`	Performs firmware download on to the adapter. ■ `f` specifies the fully qualified path to the firmware file.
`dporttest -w <wwpn>`	Sets of diagnostic tests that allow detection of physical cabling issues.
`driverparams defaults set -w <wwpn> -s [L,G] -t [P,T]`	Restores the driver parameter to the default value at the port or global level, either permanently or temporarily for the specified port. ■ `s` specifies the scope. – L for local (specified adapter). – G for global (all adapters on the host). ■ `t` specifies the nature of setting. – P for permanent (persists across reboot). – T for temporary. **NOTE:** This command replaces the `setdriverparamdefaults` command.
`driverparams set -w <wwpn> -s [L,G] -t [P,T]-p <parameter name> -v <value>`	Changes a driver parameter and designates the scope of the change. ■ `s` specifies the scope. – L for local (specified adapter). – G for global (all adapters on the host). ■ `t` specifies the nature of setting. – P for permanent (persists across reboot). – T for temporary. ■ `p` specifies the parameter name. ■ `v` specifies the value to be set for the parameter. Input must be a decimal value; otherwise, to input a hexadecimal value, add a prefix of 0x to the input. **NOTE:** This command replaces the `setdriverparam` command.

Table 10: esxcli Management Commands (Continued)

Command	Description
`dump -w <wwpn>`	Creates a dump file for a selected adapter. Dump files contain information, such as the 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 must 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. **NOTE:** If there are multiple dumps in the flash memory of an adapter, the command generates multiple `.txt` and `.bin` files.
`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. ■ d specifies the destination WWPN. It must be ff:ff:ff:ff:ff:ff:ff:ff for a switch or the actual WWPN for a target. ■ c specifies the count or number of cycles {1....99999}. ■ s specifies StopOnError: {0, 1}. – 0 = Do not stop test on error. – 1 = Stop test on error. **NOTE:** p specifies pattern with up to 8 hexadecimal characters (for example, 1a2b3c4d). This is an optional parameter. The number of cycles must be between 1 and 99,999. **NOTE:** The <*target wwpn*> is the WWPN of the target port and can be obtained by running the `targetmapping` command.
`enablebootcode -w <wwpn> -s <state>`	Enables or disables boot code on the adapter port. ■ s specifies the state of the boot code. Possible values are E and D. – E = Enable. – D = Disable.
`fctraceroute -w <initiator wwpn> -d <target wwpn>`	Shows the FC trace route details between the adapter port and the specified target port. ■ d specifies the destination WWPN. It must be ff:ff:ff:ff:ff:ff:ff:ff for a switch or the actual WWPN for a target.
`driverparams get -w <wwpn>`	Shows the name and values of each parameter. **NOTE:** This command replaces the `getdriverparams` command.
`driverparams global get -w <wwpn>`	Shows the name and global values of each parameter. **NOTE:** This command replaces the `getdriverparamsglobal` command.
`dumpdirectory get`	Displays the dump directory for the adapters in the host. **NOTE:** This command replaces the `getdumpdirectory` command.
`fwlog get -w <wwpn>`	Stops firmware log capture, retrieves the log to dump directory and restarts firmware log capture for a specified HBA port. **NOTE:** This command replaces the `getfwlog` command.
`fwparams get -w <wwpn>`	Shows the firmware parameters that can be set from the management layer. **NOTE:** This command replaces the `getfwparams` command.
`trunkinfo get -w <wwpn>`	Shows the complete trunking configuration of a physical port or trunked port specified by the WWPN on an LPe35000 adapter. **NOTE:** This command replaces the `gettrunkinfo` command.
`xcvrdata get -w <wwpn> [-t <type>]`	Shows the transceiver data in raw or formatted output. **NOTE:** This command replaces the `getxcvrdata` command.
`hbaattributes -w <wwpn>`	Lists HBA attributes of an Emulex FC HBA.

Table 10: esxcli Management Commands (Continued)

Command	Description
`hbaport attributes -w <wwpn>`	Lists port attributes of an Emulex FC HBA. **NOTE:** This command replaces the `portattributes` command.
`hbaport speed set -w <wwpn> -s<speed>`	Sets port speed of an Emulex FC HBA. ■ `s` specifies the supported link speed. Use the `hbaport attributes` command to get accurate list of supported speed values. Specify a value of 0 to configure Auto Detect mode. **NOTE:** A port reset is required to activate the new setting. **NOTE:** This command replaces the `setportspeed` command.
`hbaport statistics -w <wwpn>`	Lists port statistics of an Emulex FC HBA. **NOTE:** This command replaces the `portstatistics` command.
`hbaport disable -w <wwpn> -p <portstate>`	Disables the adapter port for a specified HBA port. ■ `p` specifies the port state as Disable 0. **NOTE:** Ensure that all I/O on the port is stopped before disabling the port. **NOTE:** A port reset may be required to activate the new setting.
`hbaport enable -w <wwpn> -p <portstate>`	Enables the adapter port for a specified HBA port. ■ `p` specifies the port state as Enable 1. **NOTE:** Ensure that all I/O on the port is stopped before enabling the port. **NOTE:** A port reset may be required to activate the new setting. **NOTE:** This command replaces the `setportenabled` command.
`listhbas`	Lists the Emulex FC HBAs.
`reset -w <wwpn>`	Resets the FC function.
`resetportstatistics -w <wwpn>`	Resets the FC port statistics.
`dumpdirectory set -d <DumpDirectory>`	Sets the dump directory. A dump directory must be set before a dump can be taken. **NOTE:** The dump directory must have a sub-directory under the `/vmfs/volumes` directory to store OneCommand Manager dumps *and* the directory must exist at the time of setting. The directory path must not contain spaces. **NOTE:** This command replaces the `setdumpdirectory` command.
`fwlog set -w <wwpn> -l <loglevel> -s <state>`	Starts or stops the firmware logging with the specified log level for a specified HBA port. ■ `l` specifies the logging level of firmware log. Valid values are 0 to 4. The log level determines the verbosity of the logs. 0 is the least verbose. ■ `s` specifies firmware logging. Valid values are: – 1 = Start firmware logging. – 0 = Stop firmware logging. **NOTE:** This command replaces the `setfwlog` command.

Table 10: esxcli Management Commands (Continued)

Command	Description
`fwparams set -w <wwpn> -p <parameter name> -v <parameter value>`	Sets the firmware parameters. The only firmware parameter that is currently supported is FA-PWWN. ■ p specifies the name of firmware parameter whose value is to be set. Valid values are: – FA-PWWN = {0=Disable, 1=Enable}. – FEC = {0=Disable, 1=Enable}. – DYNAMIC-DPORT = {0=Disable, 1=Enable} ■ v specifies the new value for the specified firmware parameter. **NOTE:** D_Port and FA-PWWN cannot be enabled simultaneously. If D_Port is enabled and you want to enable FA-PWWN, you must first disable D_Port. If FA-PWWN is enabled and you want to enable D_Port, you must first disable FA-PWWN. **NOTE:** This command replaces the `setfwparam` command.
`trunkinfo set -w <wwpn> -m <mode>`	Specifies the type of trunking to use for a specified HBA port on an LPe35000 adapter. ■ `m` specifies the trunking mode: (0, 1, 2) – 0 = Disables trunking on the adapter. – 1 = Enables two-lane trunking. – 2 = Enables four-lane trunking (only valid for quad-port adapters) **NOTE:** This command replaces the `settrunkmode` command.
`targetmapping -w <wwpn>`	Shows a list of mapped targets and the LUNs for the port.
`version`	Displays the version of different components.

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 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 LBA, an App Tag, and a Guard Tag (or CRC). A Type 1 DIF is defined as having a 2-byte Guard Tag, a 2-byte App tag, and a 4-byte Ref tag, which consist of the lower 32 bits of the LBA.

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

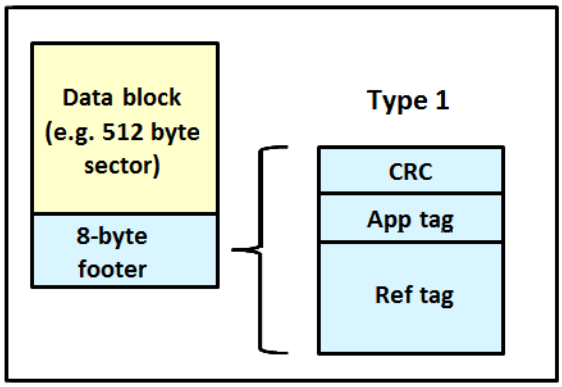

When data is written, the DIF is generated by the host, or by the adapter, based on the block data and the LBA. The DIF field is added to the end of each data block, and the data is sent through the SAN to the storage target. The storage target validates the CRC and Ref tag and, if correct, stores both the data block and DIF on the physical media. If the CRC does not match the data, the data was corrupted during the write. A Check Condition is returned 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 vice versa, on data being sent to or received from 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:

```
esxcli system module parameters set -m lpfc -p "lpfc_enable_bg=1"
```

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

B.3 SCSI Command Support – SCSI Operation Codes

When there are both `READ` and `WRITE` requests, the 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 the type of protection it has been 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 also includes information about the protection type that 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

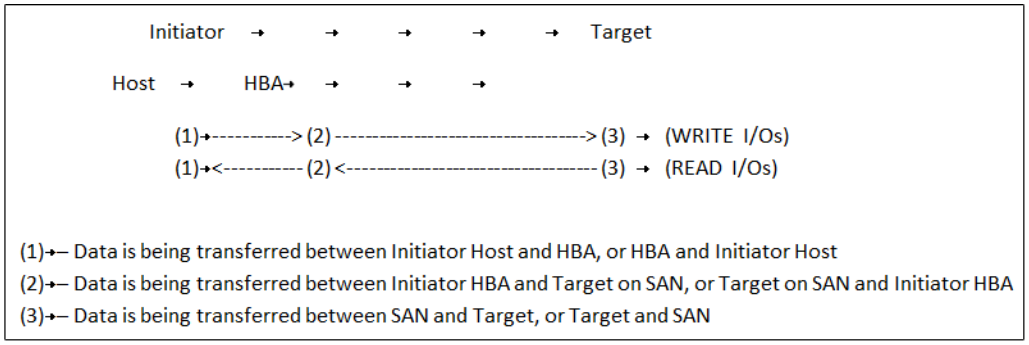

The initiator operations are listed in the following table.

Table 11: Initiator Operations

Initiator Operation	Initiator Host <-> Adapter	Adapter <-> Target	Comment
`NORMAL`	Unprotected	Unprotected	Used for unprotected I/O.
`READ_INSERT`	Protected	Unprotected	Reads the data from the target. The adapter then generates the protection data and transfers both data and protection data to the initiator host. No protection data is sent on the SAN. The adapter can insert the protection data guard tag as CRC or IP CSUM.
`READ_PASS`	Protected	Protected	Reads the data and protection data from the target on the SAN. The adapter verifies data integrity and transfers both data and protection data to the initiator host. The adapter can convert the protection data guard tag from CRC to IP CSUM.

Table 11: Initiator Operations (Continued)

Initiator Operation	Initiator Host <-> Adapter	Adapter <-> Target	Comment
READ_STRIP	Unprotected	Protected	Reads data and protection data from the target. The adapter verifies data integrity, discards 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 then generates protection data and writes 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 verifies protection data and writes both data and protection data to the target on the SAN. The adapter can convert the protection data guard tag from IP CSUM to CRC.
WRITE_STRIP	Protected	Unprotected	Transfers data and protection data from the initiator host. The adapter verifies data integrity, discards 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 ESXi 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 hypervisor. Hypervisor 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 hypervisor.

If the parameter is not passed to the driver, the default results in registering capabilities for all profiles.

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.

Table 12: lpfc_prot_mask Protection Types

Flag	Value	Indicates	Description
VMK_SCSI_TYPE1_PROT	1	Adapter supports T10 DIF Type 1	Adapter to target Type 1 protection
VMK_SCSI_DIX_TYPE0_PROT	8	Adapter supports DIX Type 0	Host-to-adapter protection only
VMK_SCSI_DIX_TYPE1_PROT	16	Adapter supports DIX Type 1	Host-to-adapter Type 1 protection

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

Table 13: Protection Data Placement for Supported Profiles

Flag	Value	BlockGuard Profile	Operation
VMK_SCSI_TYPE1_PROT	1	A1	READ_STRIP/WRITE_INSERT
VMK_SCSI_DIX_TYPE0_PROT	8	AST2	READ_INSERT/WRITE_STRIP
VMK_SCSI_DIX_TYPE1_PROT VMK_SCSI_TYPE1_PROT	17	AST1 / C1	READ_PASS/WRITE_PASS

B.4.2 lpfc_prot_guard

This parameter specifies the type of CRC that the ESXi hypervisor passes to the `lpfc` driver. The following table shows the three guard types: CRC, IP-CSUM, and TYPE_ALL with values of 0x1, 0x2, and 0x3 respectively.

Table 14: lpfc_prot_guard Guard Types

Flag	Value	Indicates
VMK_SCSI_GUARD_CRC	1	Adapter supports T10 DIF CRC
VMK_SCSI_GUARD_IP	2	Adapter supports both T10 DIF CRC and IP-CSUM
VMK_SCSI_GUARD_TYPE_ALL_VALID	3	Adapter supports both T10 DIF CRC and IP-CSUM

The default value for `lpfc_prot_guard` is VMK_SCSI_GUARD_TYPE_ALL_VALID, which results in registering capabilities for all guard types. 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.

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

Appendix C: Using the VMID Feature on a Brocade Switch

This appendix describes the setup required for the Broadcom ECD VMID feature to work on a Brocade switch.

NOTE: The Broadcom ECD VMID feature is not supported on the LPe12000-series adapters.

1. Set up the Brocade switch.
   a. You must install firmware version 8.0.1 or later.
   b. The Broadcom ECD VMID feature is enabled by default on Brocade switches.
2. Set up the target.
   a. Use an FC target that supports the VMID feature. VMID requires a VMID-supported target and initiator to work.
3. Set up the initiator.
   a. Set up the driver parameters to enable VMID on your ESXi host. For example:

   ```
   # esxcli system module parameters set -p "lpfc_max_vmid=8 lpfc_vmid_app_header=1" -m lpfc
   ```

   `lpfc_vmid_app_header` is the driver parameter that enables Brocade VMID support on ESXi.
   – The minimum value is 0 (default).
   – The maximum value is 1.
   `lpfc_max_vmid` is the driver parameter that indicates the number of VMID VMs supported.
   – The minimum value is 4.
   – The maximum value is 255.
   – The default value is 8.
4. After you have set the parameters, reboot the system for the changes to take effect.
5. Map a LUN from the FC target to a VM.
6. You can view the VMID of the VM whose LUN you have mapped in the earlier step using the following command:
   `# esxcli elxfc vmid get -n vmhba<X>`

   An output similar to the following is shown:

   ```
   Key 'vmid':
   lpfc VMID page: on
   ID00 READs:000000000001c611 WRITEs:0000000000000000
   UUID:
   35322065312062302036662038632035662036362036612d373920633320613920663620333220346420313920323 0
   String (52 e1 b0 6f 8c 5f 66 6a-79 c3 a9 f6 32 4d 19 20)
   Application id: 209
   Last access time: 2016-10-25T10:58:05
   Compressed: 52e1b06f8c5f666a79c3a9f6324d1920
   ```
7. On the Brocade switch, you can run the following command to verify the list of VMIDs:
   `> /fabos/cliexec/appserver --show -all`

An output similar to the following is shown:

```
------------------------------------------------------------
Displaying results for Fabric
------------------------------------------------------------
N_Port ID : 7e1200
Entity Name :
Entity ID (ASCII) : 52 e1 b0 6f 8c 5f 66 6a-79 c3 a9 f6 32 4d 19 20
Entity ID (Hex) :
0x3532206531206230203666203863203566203636203612d373920633320613920663620333220346420313920323000
0
Application ID : 0x00000209h (521)
------------------------------------------------------------
Application Server displays 1 entries
```

Appendix D: Using the VMID Feature on a Cisco Switch

This appendix describes the setup required for the Broadcom ECD VMID feature to work on a Cisco switch.

NOTE: The Broadcom ECD VMID feature is not supported on the LPe12000-series adapters.

1. Set up the Cisco switch.
   a. You must install firmware version 8.2 or later.
   b. The Broadcom ECD VMID feature is not enabled by default on Cisco switches. Run the following command once to configure VMIS globally on your switch:

      ```
      switch(config)# feature vmis
      ```

2. Set up the target.
   a. Use an FC target that supports the VMID feature. VMID requires a VMID-supported target and initiator to work.
3. Set up the initiator.
   a. Set up the driver parameters to enable VMID on your ESXi host.

   For example:

   ```
   # esxcli system module parameters set -p "lpfc_max_vmid=8, lpfc_vmid_priority_tagging=1" -m lpfc
   ```

   `lpfc_vmid_priority_tagging` is the driver parameter that enables Cisco VMID support on ESXi. The possible values are:
   – 0 = Disabled (default).
   – 1 = Priority tagging for targets that support it in their PLOGI `LS_ACC` response.
   – 2 = Priority tagging for all targets, whether they support `PLOGI LS_ACC`.

   `lpfc_max_vmid` is the driver parameter that indicates the number of VMID VMs supported.
   – The minimum value is 4.
   – The maximum value is 255.
   – The default value is 8.
4. After you have set the parameters, reboot the system for the changes to take effect.
5. Map a LUN from the FC target to a VM.
6. You can view the VMID of the VM whose LUN you have mapped in the earlier step using the following command:
   `# esxcli elxfc vmid get -n vmhba<X>`

   An output similar to the following is shown:

   ```
   lpfc VMID page: on
   VMID priority ranges:
           [x1 - xff], qos: x0
   VEM ID: 10:00:00:90:fa:c7:aa:b8:20:00:00:90:fa:c7:aa:b8
   ID00 READs:000000000005e853 WRITEs:000000000005e879
        UUID:
   353020303720323520373320333820366520323920366 12d3063203237203466203361203463203363203838206561
        String (50 07 25 73 38 6e 29 6a-0c 27 4f 3a 4c 3c 88 ea)
        CS_CTL VMID: x1
        Last access time: 2018-05-14T09:49:42
        Compressed: 50072573386e296a0c274f3a4c3c88ea
   ```

7. On the Cisco switch, you can run the following command to get the mapping between FCID and the WWPN:
   `(config-if)# show flogi database details`

An output similar to the following is shown:

```
--------------------------------------------------------------------------------------------
INTERFACE        VSAN    FCID           PORT NAME               NODE NAME            FLAGS
--------------------------------------------------------------------------------------------
fc3/4            102   0x010025  10:00:00:00:c9:d1:a4:ec 20:00:00:00:c9:d1:a4:ec PM
fc3/9            102   0x010045  10:00:e0:07:1b:ce:58:62 20:00:e0:07:1b:ce:58:62 P

Total number of flogi = 2.

FLAGS:
A area FCID allocation
L loop device
V FDISC
D the wwn matches the default OUI list
O the wwn matches the configured OUI list
P allocation was done based on the persistency table
M indicates a VMID capable FCID
```

On the Cisco switch you can run the following command to get the mapping between FCID and the VMID:

```
(config-if)# show vmis database
```

An output similar to the following is shown:

```
Total 1 entries
--------------------------------------------------------------------------------
INTERFACE         VSAN  FCID      LOCAL VEID           GLOBAL VEID
--------------------------------------------------------------------------------
fc3/4             102   0x010025  0x01        50072573-386e-296a-0c27-4f3a4c3c88ea
```

Appendix E: License Notices

E.1 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. See 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.]
 */
```


BROADCOM®