

2600, 2700, and 2800 Series Marvell $\mathsf{QLogic}^{\texttt{®}}$ Fibre Channel Adapters 41000 and 45000 Series Marvell Converged Network Adapters

User's Guide

Doc No. SN0154529-00 Rev. Z June 10, 2024



THIS DOCUMENT AND THE INFORMATION FURNISHED IN THIS DOCUMENT ARE PROVIDED "AS IS" WITHOUT ANY WARRANTY. MARVELL AND ITS AFFILIATES EXPRESSLY DISCLAIM AND MAKE NO WARRANTIES OR GUARANTEES, WHETHER EXPRESS, ORAL, IMPLIED, STATUTORY, ARISING BY OPERATION OF LAW, OR AS A RESULT OF USAGE OF TRADE, COURSE OF DEALING, OR COURSE OF PERFORMANCE, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT.

This document, including any software or firmware referenced in this document, is owned by Marvell or Marvell's licensors, and is protected by intellectual property laws. No license, express or implied, to any Marvell intellectual property rights is granted by this document. The information furnished in this document is provided for reference purposes only for use with Marvell products. It is the user's own responsibility to design or build products with this information. Marvell products are not authorized for use as critical components in medical devices, military systems, life or critical support devices, or related systems. Marvell is not liable, in whole or in part, and the user will indemnify and hold Marvell harmless for any claim, damage, or other liability related to any such use of Marvell products.

Marvell assumes no responsibility for the consequences of use of such information or for any infringement of patents or other rights of third parties that may result from its use. You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning the Marvell products disclosed herein. Marvell and the Marvell logo are registered trademarks of Marvell or its affiliates. Please visit www.marvell.com for a complete list of Marvell trademarks and guidelines for use of such trademarks. Other names and brands may be claimed as the property of others.

Copyright

Copyright © 2024. Marvell and/or its affiliates. All rights reserved.

Table of Contents

Preface

Supported Products	vii
Supported Operating Systems	viii
Intended Audience	viii
What Is In This Guide	viii
Related Materials	ix
Documentation Conventions	X
Technical Support	xiii
Downloading Updates and Documentation	xiii
Installation	
Supported Features	1
Installing the Hardware	2
Downloading the Latest Driver	2
Installing Management Tools	3
Installing the QConvergeConsole ESXi Client Plug-in	4
Configuring Fibre Channel Adapters 2600, 2700, 2800	
Identifying Marvell Adapters on VMware ESXi 7.0 and 8.0	6
Identifying Marvell Adapters Using the QConvergeConsole	0
ESAI Cilefil Plug-III	0
	10
	10
	10
	10
	12
	12
vSphere Client Plug-in	12
Configuring Driver Parameters Using ESXCLI	13
Dynamic Enable and Disable Driver Functionalities	20
Deploying NPIV for VMware ESXi 7.0 and 8.0.	21
Monitoring NPIV Virtual Ports from Management Tools	22
Deploying VMID for VMware ESXi 7.0 and 8.0	24

1

2

	Enabling USCM FPIN-LI/MPIO	24
	Configuring Boot from SAN	24
	Support for FC-NVMe	27
	Support for Universal SAN Congestion Mitigation (USCM)	27
3	Configuring 41000/45000 Converged Network Adapters	
	Software Components	29
	Driver Installation and Verification	29
	iSCSI Driver qedi	30
	Module Parameters for qedi	30
	Configuring the qedi Driver	31
	Verifying iSCSI Driver Presence	31
	Target Configuration	33
	Static Target Configuration	33
	Dynamic Target Configuration	35
	Differences from bnx2i	37
	Differences from bnx2fc	37
	FCoE Driver qedf	37
	Module Parameters for qedf	37
	Verifying FCoE Driver Presence	39
	FCoE Switch and Storage Configuration	40
	Configuring Boot from SAN	41
Α	Troubleshooting	
	Troubleshooting with LEDs.	46
	Dynamic Extended Logging	46
	Troubleshooting with Driver Logs	48
	Debugging Using vmkmgmt keyval Information	51
	Collecting System Logs for Troubleshooting	52
	Frequently Seen Issues	52
В	Revision History	

Glossary

List of Figures

Figure		Page
2-1	vmkmgmt_keyval Output	7
2-2	Installing the Asynchronous Driver.	11
2-3	Removing the Asynchronous Driver	11
2-4	NPIV Example	21
2-5	NPIV in vCenter Server Plug-in	23
2-6	Virtual Port Information in vCenter Server Plug-in	23
2-7	Starting Fast!UTIL from the BIOS Screen	25
3-1	Driver Installation Package Command.	29
3-2	Storage Adapters	32
3-3	Network Port Binding Page	33
3-4	Add Static Target Server	34
3-5	Static Configuration Added.	35
3-6	Add Send Target Server	36
3-7	Dynamic Target Configuration Added.	36
3-8	VC Client: Storage Adapter Page.	40
3-9	Starting Fast!UTIL from the BIOS Screen	43
A-1	Snapshot of Driver Log: ESXi 7.0 and 8.0	48

List of Tables

Table Page 2-1 Configurable Driver Parameters on 14 2-2 Dynamic Configurable Parameters: 20 Dynamic Configurable Parameters for 2-3 Each Adapter: ESXi 7.0 and 8.0 21 Common Driver Log Messages: ESXi 7.0 and 8.0. A-1 49 A-2 Output Fields for vmkmgmt keyval 51 Frequently Seen Issues A-3 52

Preface

This guide provides detailed instructions on the configuration and troubleshooting of Marvell[®] QLogic[®] Fibre Channel Adapters and Marvell FastLinQ[®] Converged Network Adapters on VMware[®] ESXi 7.0 and 8.0. It also provides details on the use of Marvell adapter features to enhance the value of server virtualization using VMware ESXi 7.0 and 8.0. Such features include virtual adapter configuration using N_Port ID virtualization (NPIV) and "boot from SAN" configuration along with Fibre Channel-Non-volatile Memory Express (FC-NVMe[®]) for lower latency storage connections.

Supported Products

Marvell provides a portfolio of Fibre Channel and Converged Network Adapters that can be used on the VMware ESXi 7.0 and 8.0 platforms:

- Marvell QLogic Fibre Channel Adapters
 - 2670 Series 16GFC Adapters in PCI Express and PCI Express Gen3 ExpressModule form factors (ESXi 7.0 only)
 - 2690 Series Enhanced 16GFC Adapters in PCI Express Gen3 form factor
 - □ 2740/2760 Series 32GFC Adapters in PCI Express Gen3 form factor
 - 2770 Series Enhanced 32GFC Adapters in PCI Express Gen4 form factor
 - □ 2870 Series 64GFC Adapters in PCI Express Gen4 form factor
- Marvell FastLinQ Converged Network Adapters
 - 41000 Series Adapters in PCI Express Gen3 form factors is covered in its respective user's guide.
 - 45000 Series Adapters in PCI Express Gen3 form factors is covered in its respective user's guide.

Supported Operating Systems

- For 2600, 2700, and 2800 Series Marvell QLogic® Fibre Channel Adapters:
 - ESXi 7.0 U3, ESXi 8.0, ESXi 8.0 U1, ESXi 8.0 U2, ESXi 8.0 U3
- For 41000 and 45000 Series Marvell FastLinQ Converged Network Adapters:
 - ESXi 7.0U3, ESXi 8.0 U3

Intended Audience

This guide is intended for users deploying Marvell Fibre Channel and Converged Network Adapters on VMware ESXi 7.0 and 8.0. These users range from end users, such as data center managers and system administrators, to the test and development community.

What Is In This Guide

This user's guide provides a brief introduction to the Fibre Channel and Converged Network Adapter products from Marvell. The primary focus of this guide is to explain the adapter driver features and management tools available on VMware ESXi 7.0 and 8.0, and the supported features.

NOTE

For Converged Network Adapters, this guide covers only the Fibre Channel over Ethernet (FCoE) functionality. To download FCoE drivers, see the instructions in "Downloading Updates and Documentation" on page xiii.

This preface covers the intended audience, related materials, document conventions used, license agreements, and technical support. The remainder of the guide is organized into the following chapters and appendices:

- Chapter 1 Installation details the supported features and covers installation of hardware, drivers, and management tools.
- Chapter 2 Configuring Fibre Channel Adapters 2600, 2700, 2800 covers how to identify Marvell adapters, start *Fast!*UTIL, verify that the driver is loaded, configure driver parameters (including VMID), deploy N_Port ID virtualization (NPIV), monitor NPIV from the management tools, and configure the boot from SAN functionality.

- Chapter 3 Configuring 41000/45000 Converged Network Adapters provides information about driver storage, including:
 - 41000/45000
 - qedi (native)
 - qedf (native)
- Appendix A Troubleshooting details troubleshooting methods, including LEDs and logs.
- Appendix B Revision History contains a list of changes made to this guide since the last revision.

At the end of this user's guide are a glossary of terms to help you quickly find information.

Related Materials

For information about downloading documentation from the Marvell Web site, see "Downloading Updates and Documentation" on page xiii):

- Marvell adapters:
 - □ User's Guide: Fibre Channel Adapter (2600 Series), part number FC0054609-00
 - □ User's Guide: Fibre Channel Adapter (2700 Series), part number 83270-546-00
 - □ User's Guide: Fibre Channel Adapter (2800 Series), part number MA2854601-00
 - □ User's Guide: Converged Network Adapters and Intelligent Ethernet Adapters (FastLinQ 41000 Series), part number AH0054601-00
 - □ User's Guide: Converged Network Adapters and Intelligent Ethernet Adapters (FastLinQ 45000 Series), part number BC0154501-00
- Management tools for Marvell adapters on VMware ESXi:
 - □ User's Guide: Marvell ESXCLI Plug-in for VMware vSphere (3400, 41000, 45000 Series Adapters), part number BC0151101-00
 - User's Guide: Marvell FastLinQ PowerKit (Windows PowerShell cmdlets for 3400, 8400, 41000, 45000 Adapters), part number BC0054518-00
- For information about using Marvell Fibre Channel (MRVLFC) PowerKit with Linux, VMware, and Windows, see the following document:
 - □ User's Guide: Marvell QLogic PowerKit for 2600, 2700, 2800 Series Fibre Channel Adapters, part number TD-001187

- For information about using Marvell QLogic FC QConvergeConsole Extension for Windows Admin Center, see the following document:
 - □ User's Guide: Marvell QLogic Fibre Channel QConvergeConsole Extension for Windows Admin Center, part number TD-001728
- For information about using QConvergeConsole vSphere Client Plug-in, see the following document:
 - □ User's Guide: Marvell QLogic QConvergeConsole Plug-ins for vSphere (2600, 2700, 2800, 3400, 8300, 8400, 41000, 45000 Series Adapters), part number SN0054677-00

The following document provides additional relevant information on VMware ESXi and can be downloaded from the VMware Web site:

 vSphere Storage Guide—Describes storage options available to VMware ESXi 8.0 and ESXi 7.0:

https://docs.vmware.com/en/VMware-vSphere/index.html

Documentation Conventions

This guide uses the following documentation conventions:

- **NOTE** provides additional information.
- Text in blue font indicates a hyperlink (jump) to a figure, table, or section in this guide, and links to Web sites are shown in <u>underlined blue</u>. For example:
 - □ Table 9-2 lists problems related to the user interface and remote agent.
 - See "Installation Checklist" on page 6.
 - For more information, visit <u>www.marvell.com</u>.
- Text in **bold** font indicates user interface elements such as a menu items, buttons, check boxes, or column headings. For example:
 - Click the Start button, point to Programs, point to Accessories, and then click Command Prompt.
 - Under Notification Options, select the Warning Alarms check box.
- Text in Courier font indicates a file name, directory path, or screen output. For example:
 - □ To return to the root directory from anywhere in the file structure, type cd /root, and then press the ENTER key.

Text in **Courier bold** font indicates a command. For example:

□ Issue the following command: sh ./install.bin

- Key names and key strokes are indicated with UPPERCASE:
 - Press CTRL+P.
 - Press the UP ARROW key.
- Text in *italics* indicates terms, emphasis, variables, or document titles. For example:
 - □ For a complete listing of license agreements, refer to the applicable *Software End User License Agreement*.
 - □ What are *shortcut keys*?
 - □ To enter the date type *mm/dd/yyyy* (where *mm* is the month, *dd* is the day, and *yyyy* is the year).
- Topic titles between quotation marks identify related topics either within this manual or in the online help, which is also referred to as *the help system* throughout this document.
- Command line interface (CLI) command syntax conventions include the following:
 - Plain text indicates items that you must type as shown. For example:
 - qaucli -pr nic -ei
 - (angle brackets) indicate a variable whose value you must specify. For example:
 - <serial_number>

NOTE

For CLI commands only, variable names are always indicated using angle brackets instead of *italics*.

- □ [] (square brackets) indicate an optional parameter. For example:
 - [<file_name>] means specify a file name, or omit it to select the default file name.
- (vertical bar) indicates mutually exclusive options; select one option only. For example:
 - on|off
 - 1|2|3|4

- ... (ellipsis) indicates that the preceding item may be repeated. For example:
 - x... means *one* or more instances of x.
 - [x...] means *zero* or more instances of x.
- () (parentheses) and { } (braces) are used to avoid logical ambiguity.For example:
 - a | b c is ambiguous { (a|b) c} means a or b, followed by c { a | (b c) } means either a, or b c

Technical Support

Customers should contact their authorized maintenance provider for technical support of their Marvell QLogic and FastLinQ products.

Downloading Updates and Documentation

To download firmware, software, and documentation:

- 1. Go to <u>www.marvell.com</u>.
- 2. Click **Support**, and then under **Tools & Resources**, click **Driver Downloads**.
- 3. In the Marvell Drivers window:
 - a. (MUST) Under CATEGORY, select either FIBRE CHANNEL ADAPTERS or CONVERGED NETWORK ADAPTERS.
 - b. (optional) Under PLATFORM/OS, select the platform/OS that matches your system.
 - c. (optional) Under PART NUMBER, select the part number for your adapter.
 - d. (optional) Under KEYWORDS, type a keyword describing what you are looking for.
- 4. Click **Apply**.

- 5. Locate the firmware (boot code), software (drivers, management tools), or document (documentation for user's guides) you need, and then do one of the following:
 - a. Click the blue text in the DESCRIPTION column.
 - b. Click the arrow in the DOWNLOAD column.

NOTE

Marvell recommends downloading the associated Read Me and Release Notes for more information. To find them, enter either **Read Me** or **Release Notes** in the KEYWORDS search box.

A message may appear asking you to review and accept the Marvell Limited Use License Agreement.

6. If applicable, read the agreement, select the check box, and then click **I ACCEPT** to accept the end license agreement and start the download.

1 Installation

This chapter provides the following installation information:

- Supported Features
- "Installing the Hardware" on page 2
- "Downloading the Latest Driver" on page 2
- "Installing Management Tools" on page 3

Supported Features

When properly installed in a VMware ESXi 7.0 and 8.0 environment, Marvell QLogic 16Gb, 32Gb, 64Gb Fibre Channel Adapters and Marvell FastLinQ 10/25/40/50/100Gigabit Ethernet (GbE) Converged Network Adapters support the following features:

- **Topologies** include FC-AL, FC-AL2, point-to-point, and switched fabric.
- **Protocols** include FCP-3-SCS, FC-tape (FCP-2), and FC-NVMe.
- Link Rates include 16Gb, 32Gb, and 64Gb for Fibre Channel Adapters, and 10/25/40/50/100Gb Enhanced Ethernet for Converged Network Adapters with auto-rate negotiation and support for MSI-X.
- Single driver support for
 - All Marvell QLogic Fibre Channel family of adapters
 - □ All Marvell FastLinQ 41000/45000 family of adapters
- SNIA Common Adapter API v1.x and v2.0 compliant.
- Marvell CIM Providers for Adapter Management for ESXi 7.0.
- DSDK based ESXCLI plugin and Powerkit provider support on ESXi8.0 and later versions
- N_Port ID virtualization (NPIV) inbox driver supports 255 virtual adapters per port.
- Boot from SAN support enables servers on a network to boot their operating system from a Fibre Channel redundant array of independent/inexpensive disks (RAID) unit on the SAN.

Page 1

Doc No. SN0154529-00 Rev. Z

Copyright © 2024 Marvell

- Support for MSI-X on 16Gb, 32Gb, and 64Gb Fibre Channel and 10/25/40/50/100GbE Converged Network Adapters.
- Support for VMware SCSI path failover for Active/Active and Active/Passive storage devices.
- **Support for VMware NVMe path failover** for Active/Active storage devices.
- Support for VMware VMotion[®] and VMware Storage VMotion, which enables live migration of running VMs from one physical server to another with zero downtime, continuous service availability, and complete transaction integrity.
- Remote SAN and adapter management uses the QConvergeConsole management tool plug-ins for VMware ESXi: the QConvergeConsole vSphere Client Plug-in (for information on installing these plug-ins, refer to the User's Guide: Marvell QLogic QConvergeConsole Plug-ins for vSphere).
- Support for virtual machine identifier (VMID) tags I/O frames so they can be associated with the virtual machine that issued each I/O for reporting or traffic management purposes as they traverse the SAN fabric.

Installing the Hardware

For detailed information on installing a Marvell Fibre Channel or Converged Network Adapter on a host, refer to the appropriate user's guide for that adapter (see "Related Materials" on page ix). For instructions on downloading documentation from the Marvell Web site, see "Downloading Updates and Documentation" on page xiii.

Downloading the Latest Driver

The Marvell QLogic driver for Fibre Channel ships "inbox" with VMware ESXi 7.0 and 8.0. However, in some specific cases, you *may* need to download the following from the VMware Website:

- An *out-of-box* driver CD
- VMware ESXi patches that might include specific fixes for adapters

Fibre Channel over Ethernet (FCoE) *requires* that you download the latest driver from the VMware Website:

https://my.vmware.com/web/vmware/downloads

Installing Management Tools

The management tools for Marvell 16Gb, 32Gb, and 64Gb Fibre Channel and 10/25/40/50/100GbE Converged Network Adapters on VMware ESXi platforms include the Marvell Adapter CIM providers, the QConvergeConsole plug-ins for VMware (the VMware vCenter Server Plug-in), and the Marvell QLogic Fibre Channel PowerKit.

For detailed information on installing the management tool plug-ins, refer to the *User's Guide: Marvell QLogic QConvergeConsole Plug-ins for vSphere*. For instructions on downloading documentation from the Marvell Web site, see "Downloading Updates and Documentation" on page xiii.

- Marvell Fibre Channel and Converged Network Adapter CIM Provider provides the back end for the adapter management capabilities of the QConvergeConsole vSphere Client Plug-in.
- QConvergeConsole ESXi Client Plug-in provides the following management capabilities:
 - □ Management for Fibre Channel and Converged Network Adapters
 - Storage and network maps that provide an end-to-end view of the Marvell adapter connections to the software and hardware components in the VMware ESX and ESXi environments
 - Downloading and updating the adapter boot code and firmware for all supported Marvell adapters
 - Querying and modifying driver parameters for all supported protocols.
 - □ Viewing and managing initiators, targets, and LUNs for Fibre Channel and FCoE ports
 - Querying statistics, running diagnostics, and obtaining transceiver information

For more information, see the *Marvell QLogic QConvergeConsole Plug-ins for vSphere User's Guide*, part number SN0054677-00.

- VMware Fibre Channel ESXCLI Plug-in provides the following management capabilities:
 - □ Management for Fibre Channel and FCoE Adapters
 - Updating the boot code and firmware for all supported Marvell adapters
 - Viewing targets and LUNs for Fibre Channel and FCoE ports
 - Querying statistics, running diagnostics, and obtaining transceiver information

For more information, see the User's Guide: Marvell QLogic QConvergeConsole Plug-ins for vSphere

Marvell QLogic Fibre Channel PowerKit (MRVLFC) provides custom PowerShell cmdlets to monitor and manage your Marvell QLogic Fibre Channel Adapters through the Windows PowerShell[®] interface.

For more information, see the *Marvell QLogic PowerKit User's Guide*, part number TD-001187.

Installing the QConvergeConsole ESXi Client Plug-in

For information on installing the QConvergeConsole ESXi Client Plug-in, refer to the *Marvell QLogic QConvergeConsole Plug-ins for vSphere User's Guide*. For information on downloading documentation from the Marvell Web site, see "Downloading Updates and Documentation" on page xiii.

2 Configuring Fibre Channel Adapters 2600, 2700, 2800

Marvell QLogic Fibre Channel Adapters on VMware ESXi 7.0 and 8.0 configuration procedures include the following:

- "Identifying Marvell Adapters on VMware ESXi 7.0 and 8.0" on page 6
- "Verifying that the Driver Is Installed and Loaded" on page 12
- Configuring the Driver Parameters" on page 12
- "Deploying NPIV for VMware ESXi 7.0 and 8.0" on page 21
- "Monitoring NPIV Virtual Ports from Management Tools" on page 22
- "Deploying VMID for VMware ESXi 7.0 and 8.0" on page 24
- "Configuring Boot from SAN" on page 24
- "Installing the Asynchronous Fibre Channel Driver" on page 10
- "Enabling USCM FPIN-LI/MPIO" on page 24

Identifying Marvell Adapters on VMware ESXi 7.0 and 8.0

On VMware ESXi 7.0 and 8.0, you can identify Marvell Fibre Channel and Converged Network Adapters in the following ways:

- Identifying Marvell Adapters Using the QConvergeConsole ESXi Client Plug-in
- Identifying Marvell Adapters Using the ESXi Shell

Identifying Marvell Adapters Using the QConvergeConsole ESXi Client Plug-in

You can use the VMware vSphere Client Plug-in to identify Marvell Fibre Channel and Converged Network Adapters on VMware ESXi 7.0 systems.

The procedure in this section applies only to Marvell adapters that use the qlnativefc driver. It does not apply to the 41000/45000 Series Adapters.

NOTE

For detailed information on the Plug-in, refer to the User's Guide: Marvell QLogic QConvergeConsole Plug-ins for vSphere.

To identify Marvell Fibre Channel and Converged Network Adapters using the ESXi QConvergeConsole vSphere Client Plug-in:

- 1. Open a Web browser and navigate to the vSphere QConvergeConsole vSphere Client Plug-in
- 2. Enter the credentials for the vSphere UI Server to log in.
- 1. In the VMware vSphere QConvergeConsole vSphere Client Plug-in, click **vCenter** in the left navigation pane.
- 2. On the vCenter Home page, click **Hosts** in the left navigation pane to view a list of the hosts that are connected to this vCenter Server.

If no hosts are connected to the vCenter Server, connect a host to the vCenter Server by following the instructions provided by VMware.

- 3. In the provided host list, click one of the hosts.
- 4. Click the **Configure** tab.
- 5. In the system tree on the left, click the **QConvergeConsole** link.

The QConvergeConsole page lists the Marvell adapters on the selected host.

Identifying Marvell Adapters Using the ESXi Shell

You can use the ESXi shell to identify Marvell Fibre Channel and Converged Network Adapters on VMware ESXi 7.0 and 8.0 systems.

To identify Marvell FC adapters and CNAs using the VMware ESXi 7.0 and 8.0 shell:

1. Discover the device using lspci:

lspci | grep QLogic

Following is a sample output:

0000:43:00.0 Serial bus controller: QLogic Corp QLE2742 Dual Port 32Gb Fibre Channel to PCIe Adapter [vmhba4]

0000:43:00.1 Serial bus controller: QLogic Corp QLE2742 Dual Port 32Gb Fibre Channel to PCIe Adapter [vmhba5]

2. To list all the keys for a specific instance, issue the following command:

/usr/lib/vmware/vmkmgmt_keyval/vmkmgmt_keyval -i vmhba<X>/qlogic -1

Where vmhba<x> indicates the name of the storage adapter. The name refers to the physical adapter on the host, not to the SCSI controller used by the virtual machines.

For example:

/usr/lib/vmware/vmkmgmt keyval/vmkmgmt keyval -i vmhba67/qlogic -l



Figure 2-1. vmkmgmt_keyval Output

1. To view adapter details for the corresponding vmhba, issue the following command:

```
# /usr/lib/vmware/vmkmgmt keyval/vmkmgmt keyval -i vmhba<x>/qlogic -k ADAPTER -g
```

For example:

/usr/lib/vmware/vmkmgmt_keyval/vmkmgmt_keyval -i vmhba65/qlogic -k ADAPTER -g
Key 'ADAPTER':
QLogic PCI-Express Dual Channel 16Gb Fibre Channel HBA for QLE8362:
 FC Firmware Version: 6.06.03 (d0d5), Driver version 1.1.29.0
Host Device Name vmhba65
BIOS version 3.19
FCODE version 4.04
EFI version 5.39
Flash FW version 6.06.03
ISP: ISP2031, Serial# AFF1141F01200
MSI-X enabled
Request Queue = 0x4102da035000, Response Queue = 0x4102da056000

2. To view VMID details for the corresponding vmhba, issue the following command. The information dump includes whether VMID (application header and priority tagging) is supported on the adapter port, the number of VMIDs registered, and additional VMID information.

\$ /usr/lib/vmware/vmkmgmt_keyval/vmkmgmt_keyval -i vmhbaN/qlogic -k ADAPTER -g

The following example output shows the Application Identifier information with five VMs running:

```
Application Header (VMID) supported
Number of VMIDs registered = 5
Application Tag 0x209: UUID 50 0d ea e6 e0 83 63 1a-73 2e 70 6d 92 b6 f0 76
        Time since last I/0 (in ms): 0
Application Tag 0x10: UUID 50 0d al ff 73 fe 6d e9-6a 5d 85 18 f9 87 c3 03
        Time since last I/0 (in ms): 0
Application Tag 0x20a: UUID 50 0d ab 33 f7 5f 39 a2-07 ab 5c fe 07 3d 01 ef
        Time since last I/0 (in ms): 0
Application Tag 0x11: UUID 50 0d f8 1e e3 98 37 2c-fa 49 99 2d 64 07 79 4c
        Time since last I/0 (in ms): 0
Application Tag 0x20b: UUID 50 0d 15 bc 72 a5 b1 25-30 50 3c 93 00 52 de 82
        Time since last I/0 (in ms): 0
```

3. To view target details for the corresponding vmhba, issue the following command:

```
# /usr/lib/vmware/vmkmgmt_keyval/vmkmgmt_keyval -i vmhba<x>/qlogic -k TARGET -g
```

For example:

/usr/lib/vmware/vmkmgmt_keyval/vmkmgmt_keyval -i vmhba65/qlogic -k TARGET -g
Key 'TARGET':

Driver version 1.1.29.0

Host Device Name vmhba65

FC Target-Port List: scsi-gla1-target-0=5006016141e0f364:010300:2:0nline;

- 4. To view NPIV port details for the corresponding vmhba, issue the following command:
- # /usr/lib/vmware/vmkmgmt_keyval/vmkmgmt_keyval -i vmhba<x>/qlogic -k NPIV -g

For example:

```
# /usr/lib/vmware/vmkmgmt_keyval/vmkmgmt_keyval -i vmhba65/qlogic -k NPIV -g
Key 'NPIV':
Driver version 1.1.29.0
Host Device Name vmhba65
NPIV Supported : Yes
Number of Virtual Ports in Use = 4
Virtual Port Information:
    Device Name for Virtual Port 1: vmhba69
    Device Name for Virtual Port 2: vmhba68
Device Name for Virtual Port 3: vmhba70
Device Name for Virtual Port 4: vmhba71
```

The commands in the previous steps list the driver's supported features and its current state, in addition to its firmware, driver, and boot code versions.

vmkmgmt_keyval includes the following:

- □ Name shows the key value (keyval) number for corresponding adapter.
- □ Host adapter: loop state specifies the current state of the adapter port, including:
 - DOWN—Before 30 seconds when no connection to a fabric or target.
 - UP—Driver is updating its target database because of a change in the fabric.

- READY—Driver is ready to handle I/O commands.
- DEAD—No connection to a fabric or target.
- UNKNOWN—When Loop state is not specified.
- ZIO mode is a non-zero value that indicates that the driver is operating in the Interrupt Coalescing mode (zero interrupt operation), reducing the load on the CPU.
- NPIV Supported indicates whether the adapter port supports NPIV. If Yes, the maximum quantity of supported virtual ports on that specific adapter port is shown in the next line, Max Virtual Ports.
- SCSI Device Information lists the adapter's world wide names (WWNs) and FC IDs.
- □ FC Target-Port List lists the target devices currently attached to the adapter port.
- □ Non-Target FC Port Information lists all the devices that were connected to this port since power on. This field is used for debug purposes only, and may contain some redundant information.
- □ Virtual Port Information lists (if virtual ports were created on that specific adapter port) the virtual port's world wide names (WWNs) and state of corresponding virtual port.
- FC Port Information for Virtual Port 1 lists the virtual target devices currently attached to the virtual adapter port.

Installing the Asynchronous Fibre Channel Driver

Follow the instructions in this section to install or reinstall the asynchronous Fibre Channel driver. This section applies to 2600/2700/2800 Series Fibre Channel Adapters (does *not* apply to 41000/45000 Series Adapters).

Initial Driver Installation

To install the driver, issue the appropriate command:

```
# esxcli software vib install -n qlnativefc -d
/<directory of offline-bundle>
Of
esxcli software vib install -v /<path>/async-driver.vib
Of
esxcli software component apply -n /<path>/
driver-component.zip
```

For example:

```
# esxcli software vib install -v
/vmfs/volumes/storage2/qlnativefc-1.1.3.0-10EM.550.0.0.114902
4.x86_64.vib
```



Figure 2-2. Installing the Asynchronous Driver

Subsequent Driver Installation

1. Find the vSphere Installation Bundle (VIB) you are updating by issuing the following command:

```
# esxcli software vib list | grep qlnativefc
qlnativefc 1.1.2.0-10EM.550.0.0.1096149 QLogic
VMwareCertified 2013-05-21
```

2. Remove the old VIB by issuing the following command:

```
# esxcli software vib remove -f -n qlnativefc
```

```
~ # esxcli software vib remove -f -n qlnativefc
Removal Result
Message: The update completed successfully, but the system needs to be rebooted for the changes to be ef
Reboot Required: true
VIBs Installed:
VIBs Removed: QLogic_bootbank_qlnativefc_1.1.3.0-10EM.550.0.0.1149024
VIBs Skipped:
```

Figure 2-3. Removing the Asynchronous Driver

3. Install the new driver, following the steps in "Initial Driver Installation" on page 10.

Verifying that the Driver Is Installed and Loaded

Follow these steps to ensure that the Marvell Fibre Channel or Converged Network Adapter driver is installed and loaded on VMware ESXi 7.0 and 8.0 systems.

To verify that the Marvell Fibre Channel or Converged Network Adapter driver is installed and loaded:

1. Log in to the service console as *root* user.

NOTE

To log in to the service console, use an SSH (secure shell) client to connect to the ESXi host. Ensure that SSH access is enabled on the ESXi host by following the instructions from VMware.

2. Issue the following command for 2600, 2700, 2800 Fibre Channel Adapters to verify that the driver is installed:

```
# esxcli software vib list | grep qlnativefc
```

3. Issue the following command to verify that the driver is loaded:

```
# esxcfg-module -1 | grep qlnativefc
```

Configuring the Driver Parameters

You can configure driver parameters for the adapter with the following methods:

- Configuring Driver Parameters Using the QConvergeConsole vSphere Client Plug-in
- Configuring Driver Parameters Using ESXCLI
- Dynamic Enable and Disable Driver Functionalities

Configuring Driver Parameters Using the QConvergeConsole vSphere Client Plug-in

For detailed information on how to use the QConvergeConsole vSphere Client Plug-in to configure driver parameters, refer to the *User's Guide: Marvell QLogic QConvergeConsole Plug-ins for vSphere*, "Using the vCenter Server Plug-in" chapter, Host Management: Fibre Channel Parameters section.

For instructions on downloading documentation from the Marvell Web site, refer to "Downloading Updates and Documentation" on page xiii.

Configuring Driver Parameters Using ESXCLI

Use the <code>esxcfg-module</code> command to configure the Marvell Fibre Channel and Converged Network Adapter driver parameters. The procedure in this section applies only to Marvell adapters that use the qlnativefc driver. It does not apply to the 41000/45000 Series Adapters.

To configure driver parameters on VMware ESXi 7.0 and 8.0:

1. To verify that the Marvell Fibre Channel adapter driver is loaded on the ESXi system, issue the following command.

```
# esxcfg-module -1 | grep qlnativefc
qlnativefc 14 1604
```

2. To get the options for a specific module and verify whether it is loaded on boot, issue the following command:

```
# esxcfg-module -g qlnativefc
qlnativefc enabled = 1 options = 'ql2xextended error logging=1'
```

3. To list the driver parameters, issue the following command:

```
# esxcfg-module -i qlnativefc
```

4. To set the driver parameters, issue the following command:

```
# esxcfg-module -s <options> qlnativefc
```

Where <options> is one of the configurable parameters listed in Table 2-1 on page 14.

For example, issue the following command to set extended error logging in the driver **and** to set the maximum device queue depth reported by the adapter to a value of 64:

```
# esxcfg-module -s "ql2xextended_error_logging=1,
ql2xmaxqdepth=64" qlnativefc
```

- 5. For the preceding changes to take effect, reboot the server as follows:
 - # reboot

NOTE

All configuration changes made using the <code>esxcfg-module -s</code> command are *persistent*, meaning that the configuration changes will be saved across system reboots.

Table 2-1 provides an alphabetical list and descriptions of all the Marvell driver configurable parameters for VMware ESXi 7.0 and 8.0.

Driver Parameter (Type) ^a	Description	Default Value
ql2xallocfwdump (int)	Option to enable allocation of memory for a firmware dump during adapter initialization. Memory allocation requirements vary by ISP type. 0 = Memory not allocated 1 = Allocate memory	1
ql2xcmdtimeout (int)	Firmware time-out value in seconds for scsi com- mand.	0
ql2xdbwr (int)	Option to specify scheme for request queue posting. 0 = Regular doorbell 1 = Hardware-assisted doorbell (faster)	1
ql2xdevdiscgoldfw (int)	Option to enable device discovery with golden firm- ware. 0 = No discovery 1 = Discover device	0
ql2xdontresethba (int)	1 = Do not reset on failure0 = Reset on failure	0
ql2xenablemd (int)	 ISP82xx: Option to enable or disable the minidump capture capability. 0 = Disable minidump capture 1 = Enable minidump capture 	1
ql2xenablemsix (int)	Set to enable MSI or MSI-X interrupt mechanism. 0 = Enable traditional pin-based interrupt mechanism 1 = Enable MSI-X interrupt mechanism 2 = Enable MSI interrupt mechanism	1
ql2xextended_error_ logging (int)	Option to enable extended error logging. 0 = No logging 1 = Log errors	0

Driver Parameter (Type) ^a	Description	Default Value
ql2xfdmienable (int)	Enables FDMI registrations. 0 = No FDMI 1 = Perform FDMI	1
ql2xiidmaenable (int)	Enables iiDMA settings. 1 = Perform iiDMA 0 = No iiDMA	1
ql2xintrdelaytimer (int)	ZIO indicates the waiting time for firmware before it generates an interrupt to the host to notify completion of request.	1
ql2xioctltimeout (int)	ioctl time-out value in seconds for pass-through commands.	66 sec
ql2xloadfwbin (int)	Option to load firmware from the Flash or driver source for QLA82 <i>xx</i> . 0 = Load firmware bundled with driver source 1 = Load from Flash	1
ql2xloginretrycount (int)	Specify an alternate value for the NVRAM login retry count.	0
ql2xlogintimeout (int)	Login time-out value in seconds.	20
ql2xmaxlun (uint)	Defines the maximum LUNs to register with the SCSI midlayer. Maximum is 65535.	FFFFh
ql2xmaxqdepth (int)	Maximum queue depth to report for target devices.	64
ql2xmaxsgs (int)	Maximum scatter or gather entries per I/O request. The default is the maximum entries supported by the OS.	0
ql2xmdcapmask (int)	 ISP82xx: Option to set the driver capture mask for firmware minidump as follows: 00h = Use firmware recommended capture mask 03h = Capture mask 3h 0Fh = Capture mask 0Fh 1Fh = Capture mask 1Fh 7Fh = Capture mask 7Fh 	1Fh
ql2xoperationmode (int)	Option to disable ZIO mode: set to 0 to disable.	1

Driver Parameter (Type) ^a	Description	Default Value
ql2xplogiabsentdevice (int)	Port login (PLOGI) to devices that are not present after a fabric scan. This is needed for several broken switches. 0 = No PLOGI 1 = Perform PLOGI	0
ql2xshiftctondsd (int)	Set to control shifting of command type processing based on the total quantity of data segment descriptors (DSDs).	6
ql2xusedefmaxrdreq (int)	0 = Adjust PCIe maximum read request size 1 = Use system default	0
<pre>qlport_down_retry (int)</pre>	Maximum quantity of command retries to a port that returns a PORT-DOWN status.	10
ql2xbypass_log_throttle (int)	Option to bypass log throttling. 0 = Throttling enabled. 1 = Log all errors	1
ql2xattemptdumponpanic (int)	Attempt firmware dump for each function on the purple screen of death (PSOD). 0 = Do not attempt firmware dump 1 = Attempt firmware dump	0
ql2xdontresethba_83xx (int)	Option to reset the controller upon failure. 1 = Do not reset on failure 0 = Reset on failure	0
ql2xmqqos (int)	Enables multiple queue (MQ) settings. Set to the quan- tity of queues in MQ quality of service (QoS) mode.	1
ql2xmqcpuaffinity (int)	Enables CPU affinity settings for the driver. 1 = Turn on the CPU affinity 0 = No affinity of request and response I/O	1

Driver Parameter (Type) ^a	Description	Default Value
ql2xfwloadbin (int)	Option to specify location from which to load ISP firm- ware. 2 = Load firmware using the request_firmware() (hot- plug) interface. 1 = Load firmware from Flash 0 = Use default semantics	0
ql2xdisablenpiv (int)	Option to disable or enable the NPIV feature. 0 = NPIV enabled 1 = NPIV disabled	0
qla2xenablesmartsan (int)	 Enable HP Smart SAN management feature. 0 = Disable the Smart SAN feature 1 = Enable the Smart SAN feature 	0
ql2xfabricpriorityqos (int)	 Option to set the Fabric Priority quality of service (QoS) mode. 0 = Fabric Priority QoS is disabled 1 = Queue-based Fabric Priority QoS enabled 2 = IOCB-based Fabric Priority QoS enabled 	0
ql2xvmidsupport ^b (int)	Enable VMID support in the driver (VMware ESXi 7.0 or later). 0 = Disabled 1 = Enabled	1
ql2xvmidexpiration (int)	 Expiration time for VMID in minutes (VMware ESXi 7.0 or later). Lack of VM storage I/O activity (for example, if a VM shutdown occurs) leads to VMID de-registration after the expiration time. 0 = No expiration x = Expiration after x minutes 	240 (4 hours)
ql2xuseshadowregisters (int)	For ISP27 <i>xx</i> : Option to use shadow registers for request/response queues. 0 = Do not use shadow registers 1 = Use shadow registers	1

Driver Parameter (Type) ^a	Description	Default Value
ql2xnvmesupport ° (int)	Enable Fibre Channel-Non-volatile memory express (FC-NVMe) support in the drivers (VMware ESXi 7.0 and 8.0 only). 0 = Disabled 1 = Enabled	1
ql2xenhancedabort (int)	Use the enhanced abort feature on 2600, 2700, and 2800 Series Adapters for NVMe I/O. 0 = Disabled 1 = Enabled	1
ql2xrspq_follow_inptr	Follow the response queue in-pointer for response queue updates for 2700 Series and later adapters. 0 = Disabled 1 = Enabled	1
<pre>ql2x_scmr_driver_profile (int)</pre>	USCM profile to control the throttling policy during fabric congestion. The parameter is used only when a profile is not configured using a Marvell application. 0 = Monitor Only 1 = Conservative 2 = Moderate 3 = Aggressive	0
<pre>ql2x_scmr_drop_pct_low_wm (int)</pre>	Driver does not throttle requests if the throughput drops below this percentage value of the peak I/O bandwidth. The range of this parameter is from 1 to 99.	50
ql2x_scmr_use_slow_queue (int)	Queue I/O requests to slow-drain devices on slow queue to improve the firmware resource management. 0 = Off 1 = On	1
ql2xvirtuallane (int)	Negotiate for Virtual Lanes (VL) support and use it to route traffic to slow-drain devices when applicable. This parameter is used for 277X and later adapters only. 0 = Off 1 = On	0
ql2xuse_crc1_iocb (int)	Use CRC Type 1 IOCB for DIF IOs.	1

Driver Parameter (Type) ^a	Description	Default Value
ql2xstoragepoll (int)	Enable the StoragePoll feature (ESX 8.0 only). 0 = Disabled 1 = Enabled	0
ql2xrspq_follow_inptr_ legacy	Follow response queue in-pointer for response queue updates for adapters earlier than 2700 Series. 0 = Disabled 1 = Enabled	0
ql2xcontrol_edc_rdf	Allows the driver to control sending register diagnostic function (RDF)/exchange diagnostic capabilities (EDC) commands for Universal SAN Congestion Mitigation (USCM) instead of the firmware (default). 0 = Disabled (firmware control) 1 = Enabled	1
ql2xautodetectsfp (int)	Detect SFP range and set appropriate distance. 0 = Disabled 1 = Enabled	1
ql2xmaxheapsize ^d (int)	Max heap size for driver memory allocations. Maximum value is 7FFF_FFFFh Default: 0x8000000 (128 MB)	0x800000 0 (128 MB)

^a int = integer; uint = unsigned integer

^b When ql2xvmidsupport is enabled, you must set the parameter ql2xfabricpriorityqos to 0.

c ql2xnvmesupport is supported only on 2690, 2740/2760, 2770, and 2870 Series Adapters.

dql2xmaxheapsize For configurations that require more memory, such as max 8 adapters in the system, modify this value to 512MB (2000000h).

Dynamic Enable and Disable Driver Functionalities

The Marvell driver provides a mechanism to dynamically enable or disable some functionalities in VMware ESXi 7.0 and 8.0. The procedure in this section applies only to Marvell adapters that use the glnativefc driver. It does not apply to the 41000/45000 Series Adapters.

To configure a parameter on all adapters, issue the following command:

```
# /usr/lib/vmware/vmkmgmt_keyval/vmkmgmt_keyval -i MOD_PARM/qlogic
-s scsi-qla<option> -k DRIVERINFO
```

Where *<option>* is one of the configurable parameters listed in Table 2-2. For example, the following command enables extended error logging:

```
# /usr/lib/vmware/vmkmgmt_keyval/vmkmgmt_keyval -i MOD_PARM/qlogic
-s scsi-qlaenable-log -k DRIVERINFO
```

Table 2-2. Dynamic Configurable Parameters:ESXi 7.0 and 8.0

Option	Function
enable-log	Enabling extended error logging
disable-log	Disabling extended error logging
donotreset	Clearing reset of failure
doreset	Setting reset of failure
disable-npiv	Disabling NPIV
enable-npiv	Enabling NPIV

To configure a parameter for a specific adapter, issue the following command:

/usr/lib/vmware/vmkmgmt_keyval/vmkmgmt_keyval -i vmhba<x>/qlogic -s scsi-<option> -k ADAPTER

Where <option> is one of the configurable parameters listed in Table 2-3 and <x> is one of the adapter's vmhba numbers.

For example, the following command schedules the loop initialization process (LIP) for keyval 3:

```
# /usr/lib/vmware/vmkmgmt_keyval/vmkmgmt_keyval -i vmhba65/qlogic
-s scsi-qlalip -k ADAPTER
```

Table 2-3. Dynamic Configurable Parameters forEach Adapter: ESXi 7.0 and 8.0

Option	Function
chip-reset	Requesting a chip resets
fwDump	Initiating fw_dump
enable-md	Enabling minidump capture
disable-md	Disabling minidump capture
lip	Scheduling LIP

Deploying NPIV for VMware ESXi 7.0 and 8.0

Marvell Fibre Channel and Converged Network Adapters support N_Port ID virtualization (NPIV) for Fibre Channel SANs. NPIV enables each virtual machine to have its own Fibre Channel worldwide port name (WWPN) by creating multiple virtual adapters on a single physical adapter.

VMware ESXi 7.0 and 8.0 enable NPIV support on Marvell Fibre Channel 16/32/64Gb Adapters and 10/25/40/50/100GbE Converged Network Adapters. Note that VMware ESXi 7.0 and 8.0 support NPIV only on raw device mapping (RDM) devices.

To configure NPIV, the Fibre Channel switch port must be configured to enable NPIV. The specific procedure to enable NPIV on the switch port depends on the switch itself. For details, refer to the Fibre Channel switch documentation.



Figure 2-4. NPIV Example

The VMware technical note, *Configuring and Troubleshooting N-Port ID Virtualization*, describes NPIV deployment for VMware ESXi in greater detail. This document is available on the VMware Website at:

http://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/techpaper/ vsp_4_vsp4_41_npivconfig-technical-note.pdf

Monitoring NPIV Virtual Ports from Management Tools

The Marvell management tool (QConvergeConsole vSphere Client Plug-in) also supports the NPIV functionality in the drivers. The plug-ins can display the virtual ports configured on the system and provide additional information on the virtual ports, including diagnostics and monitoring information. The plug-ins generate events and notify the user of any change (such as virtual port creation and deletion) at the back end to the virtual ports.

NOTE

For detailed information on configuring Marvell adapters using the plug-ins, refer to the User's Guide: Marvell QLogic QConvergeConsole Plug-ins for vSphere.

Figure 2-5 shows a typical virtual environment in the QConvergeConsole vSphere Client Plug-in with physical and virtual ports. The proc node provides information on the quantity of virtual ports and their individual states, as well as the target devices with which they can communicate (see also "Identifying Marvell Adapters Using the ESXi Shell" on page 7).




et							
KU0084 91557C New Obtacenter 172:255.11.73 Whene ESX, 4.1.0, 26024 172:25.10.216 Met With (2) With (2) Met Mith (2) Mit	Cressure Alexator Cremental Performance Concernal Product Identifier: Link Speed: Device Name: Link Speed: PCI Function Number: Eloide MAC Address: Transceiver Statistics Vendor: Vendor:	Certonavol Conformation Adapter Management Okine winiba2 10 Gbps 2 00:00:1E005-88-91 FINISAR CORP. SEP	Tada & Events Alarma	Permission Maps Storage Port Type: Port Name: Maximum Spreed: PCI Bus Number: Type: Set Identifier	Node Port 28-08-00 10 GBase-58 CHI / SP defined h	ConvergeConsole	
	 ■ 00-00-39-34 ■ Target_60-00-39-30-00- 00-39-34 	Part Number: Connector: ^e Revision:	FTLX8571D3BCL-QL LC		Speed: Serial Number: QLogic SFP installed:	10 Gbit/Sec AHJ0AF2 Yes	y seria 10 only
	E LUN_0	Concelling and	Temperature (C)	Voltage (V)	Tx Bias (mA)	Tx Power (mW)	Rx Power (mW)
	ELUN_1	Value	46.40	3.41	8.21	0.5528	0.4921
	ELUN_2	Status	Normal	Normal	Normal	Normal	Normal
	age (SCSI_1	High Alarm	78.00	3.70	11.80	0.8318	1.0000
	🖻 🔚 Port 2	High Warning	73.00	3.60	10.80	0.6607	0.7943
	194 FC_21-00-00-0E-1E-08-68-	Low Warning	-8.00	3.00	5.00	0.3162	0.0158
	HINI SC SI 2	Low Alarm	-13.00	2.90	4.00	0.2512	0.0100
					Refresh		

Figure 2-6. Virtual Port Information in vCenter Server Plug-in

Deploying VMID for VMware ESXi 7.0 and 8.0

Marvell Fibre Channel and Converged Network Adapters support VMID for Fibre Channel SANs. VMID support allows the association of a flow with a specific VM.

Marvell's drivers for VMware ESXi 7.0 and 8.0 offer support for VMID on Marvell's 269*x* and 2700 Series Fibre Channel Enhanced 16GFC and 32GFC Adapters, as well as the 2800 Series Fibre Channel 64GFC Adapters.

To configure VMID, the Fibre Channel switch port must be configured to enable VMID. The specific procedure to enable VMID on the switch port depends on the switch itself. For details, refer to the Fibre Channel switch documentation.

NOTE

ESXi 8.0 U3 now associates and displays VC UUIDs with VMIDs. For more information, see the VMware vSphere 8.0 U3 documentation.

Enabling USCM FPIN-LI/MPIO

USCM multipath switching based on marginal links with integrity issues can be enabled with ESX 8.0 U3 or later with Marvell QLogic QLE2700 and QLE2800 adapters using FC PowerKit or ESXCLI. See *Marvell*® *QLogic*® *PowerKit User Guide* for details.

NOTE

For more information about enabling FPINs for FPIN-LI/MPIO operation in Pluggable Storage Architecture (PSA), see the VMware vSphere 8.0 U3 documentation.

Configuring Boot from SAN

Boot from SAN refers to installing and booting a server's operating system from a storage device attached to the SAN instead of from a locally attached storage device. The OS is installed on one or more LUNs in the SAN array, and the Fibre Channel and Converged Network Adapters are configured to boot from the specified LUN. Boot from SAN provides several benefits including:

- Reducing server costs
- Ease of management
- Improved disaster tolerance
- Improved boot LUN availability

VMware ESXi 7.0 and 8.0 provide boot from SAN capability on Marvell Fibre Channel and Converged Network Adapters. Non-Volatile Memory Express (NVMe) boot from SAN is supported on VMware ESXi 7.0 and later. For more information, see the Boot from SAN Setup section in one of the following Marvell user's guides:

- User's Guide: Fibre Channel Adapter (2600 Series)
- User's Guide: Fibre Channel Adapter (2700 Series)
- User's Guide: Fibre Channel Adapter (2800 Series)

For storage array and Fibre Channel switch configuration, refer to the appropriate user manual for the device.

For VMware ESXi 7.0 and 8.0-specific configuration requirements, refer to the appropriate *VMware vSphere Storage* document:

https://docs.vmware.com/en/VMware-vSphere/index.html

To configure the Marvell Fibre Channel Adapter for boot from SAN using *Fast*/UTIL:

- 1. After installing the adapter, boot the server into BIOS.
- 2. To start the *Fast!*UTIL configuration utility, at the Marvell BIOS banner screen (Figure 2-7), press the CTRL+Q keys.



Figure 2-7. Starting Fast!UTIL from the BIOS Screen

- 3. In the *Fast*!UTIL utility, choose the appropriate adapter port to be configured.
- 4. On the *Fast*!UTIL Options page, select **Configuration Settings**, and then press ENTER.
- 5. On the Configuration Settings page, select **Host Adapter Settings**, and then press ENTER.
- 6. To set the BIOS to search for SCSI devices:
 - a. On the Host Adapter Settings page, select Host Adapter BIOS.
 - b. To toggle the value to **Enabled**, press ENTER.
 - c. To exit, press ESC.

7. After the **Host Adapter BIOS** is enabled, select the boot device as shown in the following table.

Enable Selectable Boot	WWPN or LUN Boot List	Device Boot
No	x	BIOS configures the first disk drive it finds as boot device.
Yes	None specified	BIOS configures the first disk drive it finds that is also a LUN 0 as boot device.
Yes	Specified	BIOS scans through the specified boot WWPN or LUN list until it finds a disk drive, and then configures it as the boot device.

- a. Choose Selectable Boot Settings, and then press ENTER.
- b. On the Selectable Boot Settings page, choose **Selectable Boot**.
- c. Toggle the value to **Enabled**, and then press ENTER.
- 8. After the Selectable Boot is enabled, select the boot device. Marvell Fibre Channel Adapters allow the specification of up to four WWPNs or LUNs as boot devices. To select up to four boot devices, follow these steps:
 - a. To move to the **Primary** location of the **Selectable Boot List** menu, press the DOWN ARROW key.
 - b. To see a list of accessible devices in the **Select Fibre Channel Device** menu, press ENTER.
 - c. Press the DOWN ARROW key to scroll down to the device you want to put into your **Selectable Boot** menu list.
 - d. To select the requested device and load it into the **Selectable Boot** menu list, press ENTER.
 - e. (Optional) Repeat this step to specify up to three alternate boot devices.

Now the ESX system is ready to boot from SAN from the specified boot devices.

NOTE

Marvell recommends that when both the adapter BIOS and the **Selectable Boot** option are enabled, you always select a device and put it in the **Primary** boot device location of the **Selectable Boot** menu list.

To configure the Marvell Fibre Channel Adapter for boot from SAN using the QConvergeConsole plug-in for VMware vCenter Server:

For detailed information on how to use the vCenter Server plug-in to configure boot from SAN, refer to the *User's Guide: Marvell QLogic QConvergeConsole Plug-ins for vSphere*, in the section Boot Configuration—Fibre Channel Port.

For instructions on downloading documentation from the Marvell Web site, refer to "Downloading Updates and Documentation" on page xiii.

Support for FC-NVMe

For information about using FC-NVMe, see the following documents:

- User's Guide: Fibre Channel Adapter (2600 Series)
- User's Guide: Fibre Channel Adapter (2700 Series)
- User's Guide: Fibre Channel Adapter (2800 Series)

Support for Universal SAN Congestion Mitigation (USCM)

For information about using VMware, see the following documents:

- User's Guide: Marvell ESXCLI Plug-in for VMware vSphere[®]
- User's Guide: Marvell QLogic QConvergeConsole Plug-ins for vSphere (2600, 2700, 2800, 3400, 41000, 45000 Series Adapters)

For information about using Marvell Fibre Channel (MRVLFC) PowerKit with Linux, VMware, and Windows, see the following document:

 User's Guide: Marvell QLogic PowerKit for 2600, 2700, 2800 Series Fibre Channel Adapters

For information about using Marvell QLogic FC QConvergeConsole Extension for Windows Admin Center, see the following document:

User's Guide: Marvell QLogic Fibre Channel QConvergeConsole Extension for Windows Admin Center

For information about using QConvergeConsole Plug-in for VMware vCenter Server and QConvergeConsole Plug-in, see the following document:

QConvergeConsole Plug-ins for VMware vSphere

3

Configuring 41000/45000 Converged Network Adapters

This chapter provides the following information about configuring 41000/45000 Series Converged Network Adapters for VMware storage driver:

- Software Components" on page 29
- "Driver Installation and Verification" on page 29
- "Driver Installation and Verification" on page 29
- "iSCSI Driver qedi" on page 30
- "FCoE Driver qedf" on page 37
- "Configuring Boot from SAN" on page 41

Software Components

The Marvell FastLinQ 41000/45000 Series Converged Network Adapters storage drivers are comprised of a monolithic iSCSI driver (qedi) and a monolithic FCoE driver (qedf). Both of these drivers are PCI drivers, and can support iSCSI and FCoE (respectively), independent of the other protocol components and drivers.

The iSCSI driver (qedi) depends on the VMware infrastructure for specific functionality, for example, iscsid, which is VMware's iSCSI daemon. The iSCSI driver depends on iscsid for initiating and servicing specific IP services (ARP, DHCP), and session and connection management on its behalf.

The iSCSI and FCoE drivers are packaged with a compatible NIC driver for proper interoperability. If multiple protocols must co-exist, the drivers must be compatible to ensure that they do not initialize the hardware differently.

Driver Installation and Verification

To install the driver package, issue the following command:

```
# esxcli software vib install -d /<absolute_path_to_the_driver_bundle.zip>
```

Figure 3-1 shows an example of the driver installation package command.



Figure 3-1. Driver Installation Package Command

A typical package contains driver VIBs for NIC, FCoE, and iSCSI.

Issue the appropriate command to verify that the driver has been loaded:

esxcli software vib list | grep qed

Following is a sample output.

Following installation, you must reboot the system for the changes to take effect.

iSCSI Driver qedi

This section provides the following information for the qedi driver:

- Module Parameters for qedi
- Configuring the qedi Driver
- Verifying iSCSI Driver Presence
- Target Configuration
- Differences from bnx2i
- Differences from bnx2fc

Module Parameters for qedi

Module parameters for the qedi driver include the following:

qedi_debug is the debug message level to print various driver debug information. This is a bit field that enables logging in different subsections.

```
Basic Info = 0x0001,
Additional Debug = 0x0002,
Probe time messages = 0x0004,
Link handling = 0x0008,
Timer related messages = 0x0010,
Connection related messages = 0x0020,
Interface related messages = 0x0040,
Receive Errors = 0x0080,
Transmit Errors = 0x0100,
Command queuing related messages = 0x0200,
Interrupt context messages = 0x0400,
Task Management Logs = 0x8000,
All Debug messages: 0xffff
```

gedi_int_mode forces an interrupt mode other than MSI-X.

```
0 - MSI-X (Default)
```

- 1 INTX
- 2 MSI
- qedi_attemptdumponpanic attempts to save a firmware dump for each function when the system panics (PSOD).

Attempt firmware dump (Default)
 Do not attempt firmware dump

qedi_cpuaffinity enables MQ CPU-Affinity.

```
0 - Disable
1 - Enable (Default)
```

gedi_ooo_enable enables processing and handling of iSCSI out-of-order packets.

```
0 - Disable
1 - Enable (Default)
```

Configuring the qedi Driver

The qedi driver automatically binds to the exposed iSCSI functions of the Converged Network Adapter. Configure the target using ESXCLI (VMware's command line tool) or vCenter/VI Client (VMware's Web GUI) tool. This functionality and operation is similar to that of the bnx2i driver.

Verifying iSCSI Driver Presence

After installing the driver package and rebooting, verify the iSCSI driver presence.

To verify that the iSCSI devices were detected correctly:

1. To verify that the qedi driver module was loaded, issue the following command:

```
~# vmkload_mod -l | grep qedi
qedi 1 1992
```

- 2. To determine if the iSCSI functions were enumerated:
 - a. In the tree pane on the left under **Storage**, click **Storage Adapters** as shown in Figure 3-2.
 - b. In the content pane under the Marvell adapters heading, iSCSI devices are listed. In the following example, two iSCSI CNA devices were detected and listed as **vmhba32** and **vmhba33**.

172.27.16.164	🔄 🔝 🛛 🎯 Actions 🗸								=*
Getting Started Summary Monitor	Configure Permissio	ons VMs Resource Po	ols Datasto	res Networks					
"	Storage Adapters								
✓ Storage	+ 6 1 0	a -						Q Filter	•
Storage Adapters	Adapter	Туре	Status	Identifier	Targets	Devices	Paths		
Storage Devices	HP H220i Host Bus Ad	lapter							
Datastores	🚱 vmhba0	SAS	Unknown	5001438022f56c20	4	4	4		
Host Cache Configuration	Patsburg 6 Port SAT	AHCI Controller							
Protocol Endpoints	🐼 vmhba1	Block SCSI	Unknown		1	1	1		
- Networking	QLogic FastLinQ QL4	1000 Series 10/25/40/50G0	GbE Controller	(ISCSI)					
Virtual switches	🚱 vmhba32	iscsi	Online	qedil-000e1ed662dc(iqn.1998-01.com.vmware:localhost.eit	0	0	0		
VMkernel adapters	Vmhba33	ISCSI	Online	qedil-000e1ed662dd(iqn.1998-01.com.vmware:localhost.ei	0	0	0		
Physical adapters									
TCP//P configuration									
Advanced									
- Virtual Machines	Adapter Details								
Viitual machines	Hupter Details								
VM Startup/Shutdown	Properties Device	s Paths Targets Ne	twork Port Bin	ding Advanced Options					
Agent VM Settings	General								Edit
Swap file location	General								
Default VM Compatibility	Name	vmhba32							
	Model	QLogic FastLinQ QL4100	0 Series 10/2	5/40/50GGbE Controller (iSCSI)					
Licensing	iSCSI Name	iqn.1998-01.com.vmware	localhost.eitl	ab.local:1308594444:32					
Time Configuration	ISCSI Alias	qedil-000e1ed662dc							
Authentication Services	Target Discovery	Send Targets, Static Targ	ets						
Certificate	Authentication								Edit
Power Management									
Advanced System Settings	Method None								
System Resource Reservation									
Security Profile									
System Swap									
Host Profile									
- Uardwara									

c. Under **Adapter Details**, select a vmhba to view more information.

Figure 3-2. Storage Adapters

- 3. To use target functionality, associate the vmnic to the iSCSI adapter:
 - a. Under Adapter Details, click the Network Port Binding tab as shown in Figure 3-3.
 - b. On the Network Port Binding page, associate the vmnic to iSCSI adapter by selecting the respective vmnic for iSCSI vmhba.

😼 172.27.16.164 🛛 🧕 🐉 🕞	👔 💽 🖓 Actions 👻									≡*
Getting Started Summary Monitor	Configure Permissions \	Ms Resource	e Pools Datasto	res Networks						
	Storage Adapters									
✓ Storage									O Filter	•
Storage Adapters	Adapter	Type	Status	Identifier		Tamets	Devices	Paths	(Times	
Storage Devices	HP H220i Host Bus Adapter	1.jpe	Otatos	- Additional - Add		1 m. ge m	Derives	1 Milla		
Datastores	wmbba0	SAS	Unknown	5001438022(56	5c20	4	4	4		
Host Cache Configuration	Patsburg 6 Port SATA AHCI	Controller								
Protocol Endpoints	S vmhba1	Block St	CSI Unknown			1	1	1		
- Networking	QLogic FastLinQ QL41000 S	eries 10/25/40/	50GGbE Controller	(ISC SI)						
Virtual switches	🐼 vmhba32	iSCSI	Online	qedil-000e1ed	662dc(iqn.1998-01.com.vmware:localhost.eit	0	0	0		
VMkernel adapters	🚱 vmhba33	iSCSI	Offline	qedil-000e1ed6	662dd(iqn.1998-01.com.vmware:localhost.ei	0	0	0		
Physical adapters										
TCP/IP configuration										
Advanced										
▼ Virtual Machines	Adapter Details									
VM Startup/Shutdown										
Agent VM Settings	Properties Devices Pal	hs Targets	Network Port Bin	ding Advanced	Options					
Swap file location	+ × 0									
Default VM Compatibility	Port Group	/Mkernel Ad P	Port Group Policy	Path Status	Physical Network Adapter					
- System	Q VMkernel (vSwitch1)	💌 vmk1 (📀 Compliant	🔷 Not used	对 vmnic4 (10 Gbit/s, Full)					
Liconsing										
Time Configuration										
Authentication Services										
Contificato										
Deves Nenagement										
Advanced System Settings										
Advanced System Settings										
System Resource Reservation										
Security Prome										
system swap	44								1 itoms 🕞 Export - 👘	Conv -
Host Profile	ana -								Thents	Cob) +
- Hardwara										

Figure 3-3. Network Port Binding Page

Target Configuration

After successfully completing the network port binding, proceed with either:

- Static Target Configuration
- Dynamic Target Configuration

Static Target Configuration

To configure a static target:

- 1. Complete the Add Static Target Server dialog box as shown in Figure 3-4:
 - a. Type the iSCSI server name in the **iSCSI Server** box.
 - b. Type the port number in the **Port** box.
 - c. Type the iSCSI target name in the **iSCSI Target Name** box.

- d. Select the **Inherit settings from parent** check box (if applicable).
- e. Click OK.

	- digot solitol			
iSCSI Server:	192.168.5.100			
Port:	3260			
iSCSI Target Name:	iqn.1992-04.com.em	ccx.apm0012390442	27.a8	٦
	to * Somerour			

Figure 3-4. Add Static Target Server

The following occurs:

- □ An event is initiated to re-scan the adapter.
- Devices allowed by the target are visible in the result list.

After static configuration is added, it is listed under Adapter Details, Targets page, Static Discovery, as shown in Figure 3-5.

🐻 172.27.16.164 🛛 🕵 🐉 🕞 🛅	🔝 🛛 🎯 Actions 👻							E.
Getting Started Summary Monitor	Configure Permissions VMs	Resource Po	ols Datasto	res Networks				
44	Storage Adapters							
✓ Storage							<u></u>	
Storage Adapters	+ 18 📕 📓 🗗 -						Q Filter	•
Storage Devices	Adapter	Туре	Status	Identifier	Targets	Devices	Paths	
Datastores	HP H220I Host Bus Adapter	040	Linknown	E004420020/E8+20	4		4	
Host Cache Configuration	Patsburg 6 Port SATA AHCI Contr	oller	UNKIOWI	5001436022150020	4	4	4	
Protocol Endpoints	wmhha1	Block SCSI	Unknown		1	1	1	
- Networking	QLogic FastLinQ QL41000 Series	10/25/40/50G0	bE Controller	((SCS))			1	
Virtual switches	Vmhba32	ISCSI	Online	gedil-000e1ed662dc(ign.1998-01.com.vmware:localhost.eit	4	6	6	
VMkernel adapters	😪 vmhba33	ISCSI	Offline	qedil-000e1ed662dd(iqn.1998-01.com.vmware:localhost.ei	0	0	0	
Physical adaptors								
TCP/IP configuration								
Advanced								
- Virtual Machines	Adapter Details			-				
VM Startup/Shutdown								
Agent VM Settings	Properties Devices Paths	Targets Ne	twork Port Bin	ding Advanced Options				
Swan file location	Dynamic Discovery Static Dis	covery						
Default VM Compatibility						bbA	Remove Authentication Advance	d
- Sustem							Trentove [Mathematication] [Mathematication	- dini
◆ System Liconning	ISCSI server		Target Name					
Time Configuration	192.168.5.100:3260		iqn.1992-04	l.com.emc:cx.apm00123904427.a8				
Authentiesties Consistent	192.168.5.101:3260		iqn.1992-04	l.com.emc:cx.apm00123904427.a9				
Authentication Services	192.168.5.102:3260		iqn.1992-04	l.com.emc:cx.apm00123904427.b8				
Certificate	192.168.5.103:3260		iqn.1992-04	l.com.emc:cx.apm00123904427.b9				
Power Management								
Advanced System Settings								
System Resource Reservation								
Security Profile								
System Swap								
Host Profile								
- Uardwara								

Figure 3-5. Static Configuration Added

2. After the rescan, view discovered storage devices per adapter on the Devices or Path pages.

Dynamic Target Configuration

To configure a dynamic target:

- 1. Complete the Add Send Target Server dialog box as shown in Figure 3-6:
 - a. Type the iSCSI server name in the **iSCSI Server** box.
 - b. Type the port number in the **Port** box.
 - c. Select the Inherit settings from parent check box (if applicable).

d. Click **OK**.

vmhba32 - Add	Send Target Server		?
iSCSI Server: Port:	192.168.5.100 3260		
Authentication	Settings		
🗹 Inherit settir	gs from parent		
		OK Cance	

Figure 3-6. Add Send Target Server

The following occurs:

- All targets are discovered under the provided IP.
- □ The iscsid initiates a login to each of these targets.
- An event is initiated to re-scan the adapter.
- Devices allowed by the target are visible in the result list.
- After dynamic configuration is added, it is listed under Adapter
 Details, Targets page, Dynamic Discovery, as shown in Figure 3-7.

172.27.16.164 🐴 🛃 🕞 🛛	🖞 🔝 Actions 🗸								=
Setting Started Summary Monitor	Configure Permissions V	Is Resource Po	ols Datasto	res Networks					
ч	Storage Adapters								
▼ Storage								O Filter	
Storage Adapters	Adapter	Tues	Status	Identifier	Toracto	Devices	Paths	(a rinter	
Storage Devices	HD H220i Host Bus Adapter	() per	Status	Merinine.	1 and a ra	Devides	1 8018		
Datastores	wmhba0	SAS	Unknown	5001438022f56c20	4	4	4		
Host Cache Configuration	Patsburg 6 Port SATA AHCI C	ontroller							
Protocol Endpoints	S vmhba1	Block SCSI	Unknown		1	1	1		
- Networking	QLogic FastLinQ QL41000 Se	ries 10/25/40/50G0	GbE Controller	(ISCSI)					
Virtual switches	🚱 vmhba32	iSCSI	Online	qedil-000e1ed662dc(iqn.1998-01.com.vmware:localhost.eit	4	6	6		
VMkernel adapters	🐼 vmhba33	iscsi	Offline	qedil-000e1ed662dd(iqn.1998-01.com.vmware:localhost.ei	0	0	0		
Physical adapters									
TCP/IP configuration									
Advanced									
Virtual Machines	Adapter Details								
VM Startur/Shutdown									
Anont VM Softings	Properties Devices Path	IS Targets Ne	twork Port Bin	ding Advanced Options					
Swan file location	Dynamic Discovery Stati	Discovery							
Default VM Compatibility						bbA	Remove	Authentication	anced
- System									
Liconsing	ISCSI server								
Time Configuration	192.168.5.100:3260								
Cardienate									
Certificate									
Power Management									
Advanced System Settings									
System Resource Reservation									
Security Profile									
System Swap									

Figure 3-7. Dynamic Target Configuration Added

2. After the rescan, view discovered storage devices per adapter on the Devices or Path pages.

Differences from bnx2i

Significant differences exist between qedi, the driver for the Marvell FastLinQ 41000/45000 Series 10/25/40/50/100GbE Controller (iSCSI) and the previous 8400 (578*xx*) Marvell iSCSI offload drivers, bnx2i (legacy) and qfle3i (native). Differences include:

- qedi directly binds to a PCI function exposed by the CNA device, and hence independently controls a PCI function.
- qedi is not dependent on a network driver.
- qedi can independently initialize and start the hardware.

Differences from bnx2fc

Significant differences exist between qedf, the driver for the Marvell FastLinQ 41000/45000 Series 10/25/40/50/100GbE controller (FCoE) and the previous Marvell FCoE offload driver, bnx2fc (legacy). Differences include:

- qedf binds directly to a PCI function exposed by the CNA device, and therefore independently controls a PCI function.
- qedf is not dependent on a network driver; therefore, no configuration is needed to initiate device discovery (unlike bnx2fc).
- qedf can independently initialize and start the hardware.

FCoE Driver qedf

This section provides the following information for the qedf driver:

- Module Parameters for qedf
- Verifying FCoE Driver Presence
- FCoE Switch and Storage Configuration

Module Parameters for qedf

Module parameters for the qedf driver include the following:

qedf_debug_level is the debug message level to print various driver debug information. This is a bit field that enables logging in different subsections:

```
/* SCSI Task Mgmt */
0x0000080
0x00000100
               /* lport related */
               /* rport related */
0x00000200
0x00000400
               /* ELS logs */
               /* fcoe L2 frame related logs*/
0x00000800
0x00001000
               /* Init logs */
               /* Link discovery events */
0x00002000
               /* Timer events */
0x00004000
0x00008000
               /* Informational logs, e.g. device MFS, * MAC
address, WWPN, WWNN */
               /* Middle Path (MP) related */
0x00010000
               /* log non-fatal errors */
0x00020000
0x00040000
               /* Application based code */
               /* unsolicited event */
0x00080000
               /* Log fcp errors */
0x00100000
               /* Log vlan info */
0x00200000
               /* Log libfc exceptions */
0x01000000
               /* Log libfc min info */
0 \times 02000000
               /* Log libfc medium info */
0x04000000
0x08000000
               /* log libfc extended info */
               /* Session setup, cleanup, etc' */
0x20000000
               /* scatter/gather element information */
0x40000000
               /* Extensively informative messages */
0x80000000
Oxfffffff
               /* LOG all messages */
```

qedf_int_mode forces an interrupt mode other than MSI-X.

```
0 - MSI-X (Default)
1 - INTX
2 - MSI
```

gedf_attemptdumponpanic attempts to save a firmware dump for each function when the system panics (PSOD).

```
    Attempt firmware dump (Default)
    Do not attempt firmware dump
```

■ **qedf_dev_loss_timer** is the time (in seconds) after target device disappears until the NO_CONNECT status is returned to SCSI layer.

```
0-30 - Permitted Range
10 - Default
```

gedf_maxqdepth is the maximum queue depth to report for target devices.

```
1-128 - Permitted Range
64 - Default
```

CAUTION

The <code>qedf_maxqdepth</code> value is sensitive to how internal firmware resources are partitioned. Changing this parameter may negatively impact performance on other targets. Therefore, change this value only after consulting with Marvell Engineers, or at your own risk.

qedf_enable_r_a_tov enables or disables the user-defined R_A_TOV.

```
0 - Use fabric defined R_A_TOV
1 - Use user-defined R A TOV
```

■ **qedf_r_a_tov** sets the user-defined R_A_TOV value. This value is only applicable if qedf_enable_r_a_tov is set to 1.

```
1-20: Permitted Value
10 - Default
```

gedf max luns adjusts the maximum LUNs supported by the driver.

```
1-512 - Permitted Range
256 - Default
```

gedf_devloss_tmo adjusts the device link loss timeout value (in seconds).

1-30: Permitted Range 10: Default

Verifying FCoE Driver Presence

After installing the driver package and rebooting, verify that the FCoE devices are correctly detected.

To verify the FCoE driver:

1. To verify that the qedf driver module was loaded, issue the following command:

~# vmkload_mod -1 | grep qedf
qedf 1 1992

2. Verify that the FCoE functions are listed. In the VC Client, click the **Storage Adapter** tab. Figure 3-8 shows an example where two FCoE CNA devices are detected and listed as **vmhba54** and **vmhba55**.

Getting Started Summary Monitor	Configure Permissions VMs	Resource Poo	ls Datastore	es Networks Update Manager			
44	Storage Adapters						
	+ 🛃 💷 🖾 🕂						
Storage Adapters	Adapter	Туре	Status	Identifier	Targets	Devices	Paths
Storage Devices	🚱 vmhba3	Fibre Cha	Online	20:00:00:24:ff:34:0f:a1 21:00:00:24:ff:34:0f:a1	0	0	0
Datastores	ISP2532-based 8Gb Fibre Channel	to PCI Expres	s HBA				
Host Cache Configuration	🔄 vmhba4	Fibre Cha	Unknown	20:00:00:24:ff:42:4e:14 21:00:00:24:ff:42:4e:14	0	0	0
Protocol Endpoints	🗇 vmhba5	Fibre Cha	Unknown	20:00:00:24:ff:42:4e:15 21:00:00:24:ff:42:4e:15	0	0	0
- Networking	QLogic FastLinQ QL41000 Series	10/25/40/50Gb	E Controller (i S	SC SI)			
Virtual switches	vmhba32	iSCSI	Unbound	qedil-000e1ec4c6c6(iqn.1998-01.com.vmware:localhost:11	0	0	0
VMkernel adapters	🔄 vmhba33	iSCSI	Unbound	qedil-000e1ec4c6c7(iqn.1998-01.com.vmware:localhost:11	0	0	0
Physical adapters	QLogic FastLinQ QL45000 Series	40G/100G Cont	roller (FCoE)				
TCP/IP configuration	🚱 vmhba64	Fibre Cha	Online	10:00:00:0e:1e:c4:c6:c4 20:00:00:0e:1e:c4:c6:c4	3	48	90
Advanced	🔄 vmhba65	Fibre Cha	Online	10:00:00:0e:1e:c4:c6:c5 20:00:00:0e:1e:c4:c6:c5	3	48	90
- Virtual Machinos	USB Storage Controller						
	vmhba35	Block SCSI	Unknown		1	2	2

Figure 3-8. VC Client: Storage Adapter Page

FCoE Switch and Storage Configuration

Complete the following steps after the FCoE driver has been installed from the switch vendor.

To configure the FCoE switch and storage:

- 1. Configure the vfc port on the switch.
- 2. Enable DCBX on the switch in Auto mode (negotiating mode).
- 3. Ensure that Priority Flow Control is enabled.
- 4. Ensure that the initiators and targets are listed on the name server database.
- 5. Add the initiator and the target into the same zone.
- 6. Ensure that the storage is configured to expose the required targets and LUNs to the initiator.
- 7. On each of the initiator ports, issue the Rescan command.
- 8. Ensure that all the exposed targets are visible on the VC Client.

Configuring Boot from SAN

Boot from SAN refers to installing and booting a server's operating system from a storage device attached to the SAN instead of from a locally attached storage device. The OS is installed on one or more LUNs in the SAN array, and the Fibre Channel and Converged Network Adapters are configured to boot from the specified LUN. Boot from SAN provides several benefits including:

- Reducing server costs
- Ease of management
- Improved disaster tolerance
- Improved boot LUN availability

VMware ESXi 7.0 and 8.0 provide boot from SAN capability on Marvell Fibre Channel and Converged Network Adapters. Non-Volatile Memory Express (NVMe) boot from SAN is supported on VMware ESXi 7.0 and later.

For storage array and switch configuration, refer to the appropriate user manual for the device.

For FastLinQ 41000/45000 iSCSI boot from SAN information, see the applicable user's guide.

For VMware ESXi 7.0 and 8.0-specific configuration requirements, refer to the appropriate *VMware vSphere Storage* document:

https://docs.vmware.com/en/VMware-vSphere/index.html

- In switch independent NIC Partitioning (NPAR) mode, the FastLinQ 41000/45000 supports FCoE on the single partition that has FCoE-Offload enabled. But, when the FastLinQ 41000/45000 is in NPAR mode, FCoE boot from SAN is only supported on the first partition of a port and that first partition must have FCoE-Offload enabled.
- On the 41000/45000 in Default and NPAR mode:
 - FCoE boot is supported on the partition that has FCoE-Offload enabled on it.
 - □ Software iSCSI boot is supported on the Ethernet-enabled first partition. iSCSI-Offloaded boot from SAN is not supported on VMware.
- On the QL41000/45000, software iSCSI is supported on any Ethernet enabled partition, while iSCSI-Offload is supported on an iSCSI-Offload enabled partition.
- In Single Function mode, the FastLinQ 41000/45000 supports both FCoE and FCoE boot from SAN, but that port must have FCoE-Offload enabled.

The FastLinQ 41000/45000 supports iSCSI-Offload on the port that has iSCSI-Offload enabled. The FastLinQ 41000/45000 supports software iSCSI boot from SAN only on VMware in both Single Function and NPAR modes. The FastLinQ 41000/45000 also supports software iSCSI on the Ethernet-enabled port.

The FastLinQ 41000/45000 NPAR mode per partition Quality of Service Minimum Bandwidth (Relative Bandwidth Weight) settings are superseded by their Data Center Bridging (DCB) per traffic class Enhanced Transmission Selection (ETS) settings.

The FastLinQ 41000/45000 NIC extended partitioning (NPAReP) mode per partition QoS Minimum Bandwidth (Relative Bandwidth Weight) settings are used with their DCB per traffic class ETS settings.

NPAR/NPAReP mode's Minimum Bandwidth (Relative Bandwidth Weight) settings provide a minimum guaranteed send bandwidth to each individual partition when there is more bandwidth to send (over two or more partitions) than available on their shared physical port. NPAR/NPAReP mode settings are only for send traffic.

DCB mode's ETS settings provide a minimum guaranteed bandwidth to each individual designated traffic class type (or priority group) when there is more bandwidth (over two or more traffic class types) than available on their shared physical port. DCB ETS settings are normally bidirectional.

In all cases, the NPAR/NPAReP per partition Maximum Bandwidth settings are always used.

Additionally, DCB is required for lossless iSCSI-Offload-TLV over DCB and FastLinQ 41000/45000 RoCE and RoCEv2 traffic types.

To configure the Marvell Converged Network Adapter for boot from SAN using Fast!UTIL:

- 1. After installing the adapter, boot the server into BIOS.
- 2. To start the *Fast*!UTIL configuration utility, at the Marvell BIOS banner screen (Figure 3-9), press the CTRL+Q keys.

```
QLE8362 PCI FCoE ROM BIOS VERSION 3.07 - Beta 6
Copyright (C) QLogic Corporation 1993-2012. All rights reserved.
www.qlogic.com
Press <CTRL-Q> or <ALT-Q> for Fast!UTIL
BIOS for Adapter 0 is disabled
BIOS for Adapter 1 is disabled
ROM BIOS NOT INSTALLED
<CTRL-Q> Detected, Initialization in progress, Please wait...
```

Figure 3-9. Starting Fast!UTIL from the BIOS Screen

- 3. In the *Fast*!UTIL utility, choose the appropriate adapter port to be configured.
- 4. On the *Fast*!UTIL Options page, select **Configuration Settings**, and then press ENTER.
- 5. On the Configuration Settings page, select **Host Adapter Settings**, and then press ENTER.
- 6. To set the BIOS to search for SCSI devices:
 - a. On the Host Adapter Settings page, select **Host Adapter BIOS**.
 - b. To toggle the value to **Enabled**, press ENTER.
 - c. To exit, press ESC.

7. After the **Host Adapter BIOS** is enabled, select the boot device as shown in the following table.

Enable Selectable Boot	WWPN or LUN Boot List	Device Boot
No	х	BIOS configures the first disk drive it finds as boot device.
Yes	None specified	BIOS configures the first disk drive it finds that is also a LUN 0 as boot device.
Yes	Specified	BIOS scans through the specified boot WWPN or LUN list until it finds a disk drive, and then configures it as the boot device.

- a. Choose Selectable Boot Settings, and then press ENTER.
- b. On the Selectable Boot Settings page, choose **Selectable Boot**.
- c. Toggle the value to **Enabled**, and then press ENTER.
- 8. After the Selectable Boot is enabled, select the boot device. Marvell Converged Network Adapters allow the specification of up to four WWPNs or LUNs as boot devices. To select up to four boot devices, follow these steps:
 - a. To move to the **Primary** location of the **Selectable Boot List** menu, press the DOWN ARROW key.
 - b. To see a list of accessible devices in the **Select Fibre Channel Device** menu, press ENTER.
 - c. Press the DOWN ARROW key to scroll down to the device you want to put into your **Selectable Boot** menu list.
 - d. To select the requested device and load it into the **Selectable Boot** menu list, press ENTER.
 - e. (Optional) Repeat this step to specify up to three alternate boot devices.

Now the ESX system is ready to boot from SAN from the specified boot devices.

NOTE

Marvell recommends that when both the adapter BIOS and the **Selectable Boot** option are enabled, you always select a device and put it in the **Primary** boot device location of the **Selectable Boot** menu list.

To configure the Marvell Converged Network Adapter for boot from SAN using the QConvergeConsole plug-in for VMware vCenter Server:

For detailed information on how to use the vCenter Server plug-in to configure boot from SAN, refer to the *User's Guide: Marvell QLogic QConvergeConsole Plug-ins for vSphere*, in the section Boot Configuration—FCoE Function.

For instructions on downloading documentation from the Marvell Web site, refer to "Downloading Updates and Documentation" on page xiii.

A Troubleshooting

This appendix provides the following methods of determining the status of the adapters and finding solutions to common problems:

- Troubleshooting with LEDs
- "Dynamic Extended Logging" on page 46
- "Troubleshooting with Driver Logs" on page 48
- "Debugging Using vmkmgmt_keyval Information" on page 51
- Collecting System Logs for Troubleshooting" on page 52
- "Frequently Seen Issues" on page 52

Troubleshooting with LEDs

Marvell Fibre Channel and Converged Network Adapters have LEDs that can help you diagnose problems. Typically, administrators observe the LED patterns to perform first-level troubleshooting that can indicate some obvious firmware or hardware issues.

For detailed information on the LED patterns for a specific Marvell Fibre Channel and Converged Network Adapter, refer to the appropriate user's guide for that adapter (see "Related Materials" on page ix). For instructions on downloading documentation from the Marvell Web site, see "Downloading Updates and Documentation" on page xiii.

Dynamic Extended Logging

This section applies to 2600/2700/2800 Series Fibre Channel Adapters (does *not* apply to 41000/45000 Series Adapters).

The Marvell qlnativefc driver provides a mechanism to dynamically enable extended error logging about driver operation. These logs provide debugging and error information to help you understand driver behavior regarding issues related to the adapter. These logs are saved in a system log file named vmkernel, in the /var/log/ directory.

To enable extended error logging, issue the following command:

```
# /usr/lib/vmware/vmkmgmt_keyval/vmkmgmt_keyval -i MOD_PARM/qlogic
-s scsi-qlaenable-log -k DRIVERINFO
```

To disable extended error logging, issue the following command:

```
# /usr/lib/vmware/vmkmgmt_keyval/vmkmgmt_keyval -i MOD_PARM/qlogic
-s scsi-qladisable-log -k DRIVERINFO
```

When logging is enabled, the driver version in the $vmkmgmt_keyval$ output is appended with a -debug flag as shown in the following examples.

ESXi debug example:

/usr/lib/vmware/vmkmgmt_keyval/vmkmgmt_keyval -i vmhba<x>/qlogic -k ADAPTER -g

Where vmhba < x > is one of the adapter's keyval numbers.

QLogic 16Gb 2-port FC to PCIe Gen3 x8 Adapter for QLE2692: FC Firmware Version: 8.03.03 (d0d5), Driver version 3.1.29.0-debug Host Device Name vmhba70 BIOS version 3.33 FCODE version 4.11 EFI version 6.13 Flash FW version 8.03.03 ISP: ISP261, Serial# RFD1723T38823 MSI-X enabled Request Queue = 0x430b994d4000, Response Queue = 0x430b994f5000 Request Queue count = 2048, Response Queue count = 512 Number of response queues for CPU affinity operation: 4 CPU Affinity mode enabled

NOTE

Although the driver logs provide debugging and error information, they also increase I/O latency. Therefore, Marvell does not recommend that you enable extended logging during normal operations.

Troubleshooting with Driver Logs

All messages from the Marvell Fibre Channel Adapter qlnativefc driver are logged in /var/log/vmkernel.log. These messages provide additional information about the driver operation, both normal and driver warnings. This section applies to 2600/2700/2800 Series Fibre Channel Adapters (does *not* apply to 41000/45000 Series Adapters).

- Messages from the driver are usually preceded by qlnativefc: vmhbaX(CX:TX.LX), where:
 - □ x is the adapter instance number.
 - C is the channel number.
 - \Box T is the target number.
 - L is the LUN number.

Figure A-1 shows a snapshot from a vmkernel log file

2013-02-2011/:44:12.2/02	CD00:130181) ALEKI: <f< th=""><th>2> Customize Syste</th><th>em/Installer_Screen.</th></f<>	2> Customize Syste	em/Installer_Screen.
2013-02-20T17:44:14.797Z	cpu1:32779) qlnativefc:	vmhba10(41:0.6):	DCBX Completed 0000 0000 0000
2013-02-20T17:44:14.997Z	cpu1:32779) qlnativefc:	vmhba10(41:0.6):	<pre>qla24xx_report_id_acquisition:format 1: scsi(3): VP[0] enabled - state</pre>
261			
2013-02-20T17:44:15.016Z	cpu1:32779) glnativefc:	vmhba10(41:0.6):	Asynchronous LOOP UP (10 Gbps).
2013-02-20T17:44:15.016Z	cpu1:32779) qlnativefc:	vmhba10(41:0.6):	LOOP UP detected (10 Gbps).
2013-02-20T17:44:15.016Z	cpu1:32780) glnativefc:	vmhba10(41:0.6):	Asynchronous PORT UPDATE.
2013-02-20T17:44:15.016Z	cpu1:32780) glnativefc:	vmhba10(41:0.6):	Port database changed ffff 0006 0000.
2013-02-20T17:44:15.021Z	cpu1:32929) qlnativefc:	vmhba10(41:0.6):	Asynchronous PORT UPDATE ignored 0000/0007/0b00.
2013-02-20T17:44:15.127Z	cpu1:33302)qlnativefc:	vmhba10(41:0.6):	<pre>qla2x00_reset_marker()</pre>
2013-02-20T17:44:15.127Z	cpu1:33302) glnativefc:	vmhba10(41:0.6):	<pre>qla2x00_loop_resync() - start</pre>
2013-02-20T17:44:15.140Z	cpu1:33302) qlnativefc:	vmhba10(41:0.6):	F/W Ready - OK
2013-02-20T17:44:15.155Z	cpu1:33302) glnativefc:	vmhba10(41:0.6):	fw_state=3 curr time=144801797978070.
2013-02-20T17:44:15.170Z	cpu1:33302) qlnativefc:	vmhba10(41:0.6):	vlan_id: 0x3ea
2013-02-20T17:44:15.170Z	cpu1:33302)qlnativefc:	vmhba10(41:0.6):	vn_port_mac_address: 0e:fc:00:e6:02:61
2013-02-20T17:44:15.170Z	cpu1:33302)qlnativefc:	vmhba10(41:0.6):	Configure loop dpc flags = 0x10860
2013-02-20T17:44:15.201Z	cpu1:33302) qlnativefc:	vmhba10(41:0.6):	RSCN queue entry[0] = [00/000000].
2013-02-20T17:44:15.201Z	cpu1:33302) qlnativefc:	vmhba10(41:0.6):	device_resync: rscn overflow.
2013-02-20T17:44:15.235Z	cpu1:33302)qlnativefc:	vmhba10(41:0.6):	RHBA exiting normally.
2013-02-20T17:44:15.251Z	cpu1:33302) glnativefc:	vmhba10(41:0.6):	RPA exiting normally.
2013-02-20T17:44:15.281Z	cpu1:33302)qlnativefc:	vmhba10(41:0.6):	RFT_ID exiting normally.
2013-02-20T17:44:15.297Z	cpu1:33302)qlnativefc:	vmhba10(41:0.6):	RFF_ID exiting normally.
2013-02-20T17:44:15.313Z	cpu1:33302)qlnativefc:	vmhba10(41:0.6):	RNN_ID exiting normally.
2013-02-20T17:44:15.328Z	cpu1:33302) qlnativefc:	vmhba10(41:0.6):	RSNN_NN exiting normally.
2013-02-20T17:44:15.422Z	cpu1:33302)qlnativefc:	vmhba10(41:0.6):	GID_PT entry - nn 2000000e1e048d43 pn 2100000e1e048d43 portid=e60261.
2013-02-20T17:44:15.438Z	cpu1:33302)qlnativefc:	vmhba10(41:0.6):	GID_PT entry - nn 2000000e1e048d47 pn 2100000e1e048d47 portid=e6028a.
2013-02-20T17:44:15.454Z	cpu1:33302) glnativefc:	vmhba10(41:0.6):	GID_PT entry - nn 2000000e1e05387b pn 2100000e1e05387b portid=e60293.
2013-02-20T17:44:15.469Z	cpu1:33302) glnativefc:	vmhba10(41:0.6):	GID_PT entry - nn 50060160c1e0f364 pn 5006016a41e0f364 portid=e602ef.
2013-02-20T17:44:15.469Z	cpu1:33302) glnativefc:	vmhba10(41:0.6):	device wrap (e602ef)
2013-02-20T17:44:15.469Z	cpu1:33302) glnativefc:	vmhba10(41:0.6):	Trying Fabric Login w/loop id 0x0000 for port e6028a.
2013-02-20T17:44:15.485Z	cpu1:33302)qlnativefc:	vmhba10(41:0.6):	<pre>qla24xx_login_fabric(3): failed to complete IOCB completion status</pre>
6.			
2013-02-20T17:44:15.485Z	cpu1:33302) glnativefc:	vmhba10(41:0.6):	Trying Fabric Login w/loop id 0x0001 for port e6028a.
2013-02-20T17:44:15.519Z	cpu1:32779) qlnativefc:	vmhba11(41:0.7):	Asynchronous PORT UPDATE ignored 0003/0004/0600.
2013-02-20T17:44:15.530Z	cpu1:33302) qlnativefc:	vmhba10(41:0.6):	Fabric Login successful w/loop id 0x0001 for port e6028a.
2013-02-20T17:44:15.530Z	cpu1:33302) glnativefc:	vmhba10(41:0.6):	Trying Fabric Login w/loop id 0x0002 for port e60293.
2013-02-20T17:44:15.576Z	cpu1:33302)qlnativefc:	vmhba10(41:0.6):	Fabric Login successful w/loop id 0x0002 for port e60293.
2013-02-20T17:44:15.576Z	cpu1:33302)qlnativefc:	vmhba10(41:0.6):	Trying Fabric Login w/loop id 0x0003 for port e602ef.

Figure A-1. Snapshot of Driver Log: ESXi 7.0 and 8.0

 Table A-1 lists the most common driver log messages and provides an explanation of each.

Table A-1. Common Driver Log Messages: ESXi 7.0 and 8.0

qlnativefc:<dev driver string>: Found an ISPXXXX, iobase zzz

Driver is reporting which adapter it has found during initialization.

qlnativefc: <dev driver string>: LIP reset occured (f8f7)

Driver received a LIP asynchronous event from the firmware.

qlnativefc: <dev driver string>:: LOOP UP detected (x Gbps).

Driver received a loop up asynchronous event from the firmware.

qlnativefc: <dev driver string>:: LOOP DOWN detected mbx1=xxxh mbx2=yyyh
mbx3=zzzh mbx4=iiih

Driver received a loop down asynchronous event from the firmware.

qlnativefc: <dev driver string>: Asynchronous P2P MODE received.

Driver received a point-to-point asynchronous event from the firmware.

QLogic Fibre Channel HBA Driver: 1.1.1.0

Driver is reporting information discovered during its initialization. This information includes the driver version.

qlnativefc: <dev driver string>: Loop down - aborting ISP.

Indicates that the driver is attempting to restart the loop by resetting the adapter. Usually done by the driver when sync is not detected by the firmware for a long time (4+ minutes), and usually means that the adapter port is not connected to the switch or loop.

qlnativefc: <dev driver string>: <dev driver string>: ISP System Error mbx1=x mbx2=x mbx3=x mbx7=x

Driver received an asynchronous ISP system error event from the firmware. Additional information follows the message (that is, mailbox values from the firmware).

qlnativefc: <dev driver string>: Configuration change detected: value=x.

Driver received a change in connection async event from the firmware. Additional information follows the message (that is, mailbox 1 value from the firmware).

qlnativefc: <dev driver string>: Port database changed xx yy zz.

Driver received a port database asynchronous event from the firmware. Additional information follows the message (that is, mailbox 1 value from the firmware).

Table A-1. Common Driver Log Messages: ESXi 7.0 and 8.0 (Continued)

qlnativefc: <dev driver string>: RSCN database changed -- xx yy zz. Driver received a registered state change notification (RSCN) asynchronous event from the firmware. Additional information follows the message (that is, mailbox values from the firmware). qlnativefc: <dev driver string>: Cannot get topology - retrying Firmware return status indicating it is busy. qlnativefc: <dev driver string>: Cxx:Tyy:Lzz: DEVICE RESET ISSUED. Indicates a device reset is being issued to (channel:target:lun). qlnativefc: <dev driver string>: Cxx:Tyy:Lzz: BUS RESET ISSUED. Indicates a loop reset is being issued to (channel:target:lun). qlnativefc: <dev driver string>: Cxx:Tyy:Lzz: ADAPTER RESET ISSUED. Indicates an adapter reset is being issued to (channel:target:lun). qlnativefc: <dev driver string>: Unknown Device State: x Indicates that the status returned from the firmware is not supported. %x-%x is the completion-scsi statuses. qlnativefc: <dev driver string>: Cable is unplugged... Indicates that the firmware state is in LOSS OF SYNC; therefore, the cable must be missing. qlnativefc: <dev driver string>: Performing ISP error recovery - ha=p. Indicates that the driver has started performing an adapter reset. qlnativefc: <dev driver string>: qla2x00 abort isp: **** FAILED **** Indicates that the driver failed to perform an adapter reset. qlnativefc: <dev driver string>: (Txx:Lyy): Mid-layer underflow detected (X of Y bytes) Indicates that an underflow was detected. qlnativefc: <dev driver string>: ERROR -- Unable to get host loop ID. Firmware failed to return the adapter loop ID. WARNING qlnativefc: <dev driver string>: [ERROR] Failed to allocate memory for adapter\n Indicates that the driver could not allocate all the kernel memory it needed. WARNING qlnativefc: <dev driver string>: Failed to initialize adapter-Adapter flags x.

Table A-1. Common Driver Log Messages: ESXi 7.0 and 8.0 (Continued)

Indicates that a previously occurring error is preventing the adapter instance from initializing normally.

```
WARNING qlnativefc: <dev driver string>: already in use.
```

Indicates that the driver could not register for the interrupt IRQ because it is in use by another driver.

WARNING qlnativefc: <dev driver string>: ISP Request Transfer Error (x)

Driver received a Request Transfer Error asynchronous event from the firmware.

WARNING qlnativefc: <dev driver string>: ISP Response Transfer Error

Driver received a Response Transfer Error asynchronous event from the firmware.

WARNING Error entry invalid handle

Driver detected an invalid entry in the ISP response queue from the firmware. This error will cause an ISP reset.

qlnativefc: <dev driver string>: MS entry - invalid handle

Driver detected a management server command time-out.

Debugging Using vmkmgmt_keyval Information

The vmkmgmt interface uses key-value pairs. The qlnativefc driver, through the $vmkmgmt_keyval$ command interface, provides first-hand information about the current state of the driver, which is helpful when troubleshooting problems. This section applies to 2600/2700/2800 Series Fibre Channel Adapters (does *not* apply to 41000/45000 Series Adapters).

Table A-2 lists the vmkmgmt_keyval output fields and provides a description of each.

Output Field	Description
Boot Code Version	Provides the version of the boot code present on the adapter. A version of 0.00 indicates that the boot code is corrupted or not present, and must be updated to perform a <i>boot from SAN</i> (see "Configuring Boot from SAN" on page 24).
Number of ISP aborts	Provides the quantity of times the system chip has been reset. Typically, whenever the system NVRAM or Flash is updated, the chip is reset to use the updated information. Otherwise, the chip is reset only if an error occurs. This number thus provides a close approximation of whether the chip has been operating properly.

Table A-2. Output Fields for vmkmgmt_keyval

Output Field	Description
Host adapter FC link state	Provides the connection state of the specified adapter. Possible valuesinclude the following:READY:Ready to perform I/Os.UPDATE:The topology is being rediscovered.DEAD:DOWN:Intermittent state, link temporarily down.
NPIV Supported	Indicates whether the given adapter node supports NPIV. If $_{\rm Yes}$, the subsection provides the maximum quantity of virtual ports supported.
SCSI Device Informa- tion	Provides the WWN of the given adapter.
FC Port Information	Lists all Fibre Channel ports to which the specified adapter is connected. The last entry in this field indicates the login status of the port and whether the port is operating in <i>target</i> or <i>initiator</i> mode.
SCSI LUN Information	Lists all the LUNs visible to the specified adapter.

Table A-2. Output Fields for vmkmgmt_keyval (Continued)

Collecting System Logs for Troubleshooting

To collect all relevant data on the running system into one tar file, issue the following command:

vm-support

Frequently Seen Issues

Table A-3 lists some common issues and questions and provides tips for their resolution.

Table A-3. Frequently Seen Issues

Issue	Troubleshooting Tip
One specific adapter in the system has perfor- mance issues, while all other adapters in the sys- tem are okay.	Ensure that the adapter is plugged into a slot that allows the adapter to use the full bandwidth avail- able. Also check that the bus the adapter is on is not shared by other high-bandwidth PCI cards.

Issue	Troubleshooting Tip
The following error occurs when the system is booted up. In addition, the ALT+Q QLogic BIOS banner is not shown; or <i>Fast</i> /UTIL is hanging when trying to flash the adapters. Device Resource allocation failure	There are too many devices taking up extended BIOS data area (EBDA) memory resources. To free up enough resources to allow the adapters to work properly, you can disable unused on-board services such as wake on LAN (WoL), PXE boot, and so on.
Different target numbers can be assigned to the various storage devices for the adapters on the SAN. What is the correct method to assign per- sistent bindings?	Marvell recommends that you assign the target numbers of the devices in the same order on multi- ple adapters. For example, if controller A of an array is assigned to target 0 on HBA0, you should assign controller B of an array to target 0 on HBA1, and then continue this same methodology for all targets and adapters. Note that the driver and OS must be reloaded for the changes to take effect.
Rescanning an adapter port for target or LUN dis-	<pre># esxcfg-rescan vmhba<x></x></pre>
	You can find the vmhba value corresponding to an adapter port by examining the /var/log/vmkernel.log, where the adapter is assigned a vmhba <x> vector.</x>

B Revision History

Document Revision History				
Revision A, May 30, 2014				
Revision B, March 27, 2015				
Revision C, February 22, 2016				
Revision D, January 17, 2017				
Revision E, March 7, 2017				
Revision F, September 22, 2017				
Revision G, May 11, 2018				
Revision H, January 18, 2019				
Revision J, September 20, 2019				
Revision K, October 24, 2019				
Revision L, October 31, 2019				
Revision M, November 22, 2019				
Revision N, May 8, 2020				
Revision P, February 22, 2021				
Revision R, April 9, 2021				
Revision T, August 30, 2021				
Revision W, October 12, 2022				
Revision X, March 25, 2023				
Revision Y, November 17, 2023				
Revision Z, June 10, 2024				
Changes	Sections Affected			
Removed support for 8300 and 8400 Series Mar- vell Converged Network Adapters.	All sections.			
Removed support for Marvell QLogic 8Gb Fibre Channel Adapters.	All sections.			
Changed references of VM-ID to VMID.	All sections.			
Added ESXi 8.0U3 support.	"Supported Operating Systems" on page viii.			

Changed install command to apply command: esxcli software component apply -n / <path>/ driver-component.zip.</path>	"Installing the Asynchronous Fibre Channel Driver" on page 10.
Added a NOTE about ESXi 8.0 U3 and how it works with VC UUIDs.	"Deploying VMID for VMware ESXi 7.0 and 8.0" on page 24
Added section.	"Enabling USCM FPIN-LI/MPIO" on page 24
Removed text describing Ethernet adapters.	"Configuring Boot from SAN" on page 24
Added section. (Ethernet content removed from "Configuring Boot from SAN" on page 24 moved here)	"Configuring Boot from SAN" on page 41

Glossary

adapter

The board that interfaces between the host system and the target devices. Adapter is synonymous with *host bus adapter (HBA), host channel adapter (HCA), host adapter, and adapter board.*

arbitrated loop

A circular (ring) topology (versus point-to-point) where two or more ports can be interconnected, but only two ports can communicate at a time. All communication passes through all ports connected to the loop.

adapter port

A port on the adapter board.

address resolution protocol

See ARP.

API

Application programming interface. A set of routines, protocols, and tools for building software applications. API simplifies development by providing the building blocks.

ARP

Address resolution protocol. A TCP/IP function to associate an IP address with a link-level address.

basic input output system

See BIOS.

BIOS

Basic input output system (typically in Flash PROM). The program (or utility) that serves as an interface between the hardware and the operating system and allows booting from the adapter at startup.

boot code

The program that initializes a system or an adapter. Boot code is the first program to run when a system or a device within a system, such as an adapter, is powered on. FCode, BIOS, and EFI (enhanced firmware interface) are all forms of boot code for specific hardware and operating system environments.

Boot code for Marvell Fibre Channel Adapters is required if the computer system is booting from a storage device (disk drive) attached to the adapter. The primary function of the boot code is communication with the external boot device before the operating system is up and running. Boot code can also perform secondary functions, including managing the setup for the adapter and initializing and testing the adapter's ISP.

boot device

The device, usually a the hard disk, that contains the operating system the BIOS uses to boot from when the computer is started.

boot from SAN

The ability for each server on a network to boot their operating system from a Fibre Channel RAID unit located on the SAN, rather than from a local disk or direct-attached storage (DAS). This enables easier SAN management because you can replace a server and boot it from the Fibre Channel RAID unit.

CIM

Common information model. Provides a common definition of management information for systems, networks, applications, and services, and allows for vendor extensions. CIM's common definitions enable vendors to exchange semantically rich management information between systems throughout the network.

CIM is composed of a specification and a schema. The schema provides the actual model descriptions, while the specification defines the details for integration with other management models. SMI-S 1.5.0 is based on CIM Schema 2.23.

Converged Network Adapter

Marvell Converged Network Adapters support both data networking (TCP/IP) and storage networking (Fibre Channel) traffic on a single I/O adapter using two new technologies: Enhanced Ethernet and Fibre Channel over Ethernet (FCoE).

data center bridging

See DCB.

data center bridging exchange

See DCBX.

DCB

Data center bridging. Provides enhancements to existing 802.1 bridge specifications to satisfy the requirements of protocols and applications in the data center. Because existing high-performance data centers typically comprise multiple application-specific networks that run on different link layer technologies (Fibre Channel for storage and Ethernet for network management and LAN connectivity), DCB enables 802.1 bridges to be used for the deployment of a converged network where all applications can be run over a single physical infrastructure.

DCBX

Data center bridging exchange. A protocol used by DCB devices to exchange configuration information with directly connected peers. The protocol may also be used for misconfiguration detection and for configuration of the peer.

device

A target, typically a disk drive. Hardware such as a disk drive, tape drive, printer, or keyboard that is installed in or connected to a system. In Fibre Channel, a *target* device.

DHCP

Dynamic host configuration protocol. Enables computers on an IP network to extract their configuration from servers that have information about the computer only after it is requested.

driver

The software that interfaces between the file system and a physical data storage device or network media.

dynamic host configuration protocol

See DHCP.

E_Port

Expansion port. A port in a Fibre Channel switch that connects to another Fibre Channel switch or bridge device by an inter-switch link (ISL)). E_Ports are used to link Fibre Channel switches to form a multi-switch fabric.

Enhanced Ethernet

Also called *data center Ethernet* or *converged enhanced Ethernet*. Refers to new enhancements to the existing Ethernet standard that eliminate Ethernet's inherent lossy nature and make 10Gb Ethernet a viable storage networking transport.

Ethernet

The most widely used LAN technology that transmits information between computers, typically at speeds of 10 and 100 million bits per second (Mbps).

expansion port

See E_Port.

F_Port

Fabric port. The "fabric" port in a Fibre Channel fabric switch provides a point-to-point link attachment to a single N_Port. F_Ports are intermediate ports in virtual point-to-point links between end ports, for example N_Port to F_Port to F_Port to N_Port using a single Fibre Channel fabric switch.

fabric

A fabric consists of cross-connected Fibre Channel devices and switches.

fabric loop port

See FL_Port.

fabric port

See F_Port.

fabric switch

Also, switched fabric. Connects multiple devices from independent Fibre Channel-arbitrated loops (FC-ALs) and point-to-point topologies into a fabric using Fibre Channel switches.

Fast!UTIL

Marvell *Fast*!UTIL[™] Fibre Channel Adapter BIOS utility.

FC-NVMe

NVM Express over Fibre Channel (FC) is a feature of Marvell QLogic FC adapters that provides low latency and high performance data transfer between a host and a peripheral target storage device or system.

Fibre Channel

A high-speed serial interface technology that supports other higher layer protocols such as SCSI and IP.

FCode

Forth code. A type of boot code for use on Sun[®]'s SPARC[®] or Macintosh[®] hardware platforms. See also boot code and Flash.

FCoE

Fibre Channel over Ethernet. A new technology defined by the T11 standards body that allows traditional Fibre Channel storage networking traffic to travel over an Ethernet link by encapsulating Fibre Channel frames inside Layer 2 Ethernet frames. For more information, visit www.fcoe.com.

Fibre Channel over Ethernet

See FCoE.
firmware

Low-level software typically loaded into read-only memory and used to boot and operate an intelligent device.

FL_Port

Fabric loop port. In a Fibre Channel, the fabric switch is capable of Fibre Channel Arbitrated Loop operations and is connected to one or more NL_Ports by a Fibre Channel Arbitrated Loop. An FL_Port becomes a shared entry point for public NL_Port devices to a Fibre Channel fabric. FL_Ports are intermediate ports in virtual point-to-point links between end ports that do not reside on the same loop, for example NL_Port to FL_Port to F_Port to N_Port through a single Fibre Channel fabric switch.

Flash

Non-volatile memory where the boot code is saved. At times, the terms *Flash* and *boot code* are used interchangeably.

generic port

See G_Port.

G_Port

Generic port. A port that can operate as either an E_Port or an F_Port. A G_Port can determine operating mode at switch port initialization, F_Port when an N_Port attachment is determined, E_Port when an E_Port attachment is determined. See E_Port, F_Port, FL_Port, L_Port, N_Port, NL_Port

HBA

Host bus adapter. See adapter.

HII

Human interface infrastructure is a specification (part of UEFI 2.1) for managing user input, localized strings, fonts, and forms, that allows OEMs to develop graphical interfaces for preboot configuration.

human interface infrastructure

See HII.

iiDMA

Intelligent interleaved direct memory access. A Marvell feature that ensures maximum link efficiency.

input/output control

See ioctl.

intelligent interleaved direct memory access

See iiDMA.

intelligent storage peripheral

See ISP.

ioctl

Input/output control. A system call in UNIX and Linux systems. Allows an application to control or communicate with a device driver outside usual read/write operations.

ISP

Intelligent storage peripheral. Marvell trademark and family of Fibre Channel and SCSI controller chips that replace network interface chips in network adapters, servers, and storage.

L_Port

Loop port. Does arbitrated loop functions and protocols. NL_Ports and FL_Ports are examples of loop-capable ports. See E_Port, F_Port, FL_Port, G_Port, N_Port, NL_Port.

LED

Light-emitting diode. Status indicator on a switch, router, adapter, or other device.

light-emitting diode

See LED.

LIP

Loop initialization process. The initialization process in an arbitrated loop that occurs when the loop is powered up or a new device is added. One function of a LIP is to assign addresses. All data transmission on the loop is suspended during a LIP.

loop initialization process

See LIP.

logical unit number

See LUN.

loopback

Diagnostic tool that routes transmit data through a loopback connector back to the same adapter.

LUN

Logical unit number, a subdivision of a SCSI target. It is the small integer handle that differentiates an individual disk drive or partition (volume) within a common SCSI target device such as a disk array.

Technically, a LUN can be a single physical disk drive, multiple physical disk drives, or a portion (volume) of a single physical disk drive. However, LUNs are typically not entire disk drives but rather virtual partitions (volumes) of a RAID set.

Using LUNs, the Fibre Channel host can address multiple peripheral devices that may share a common controller.

media

Physical-layer information carriers. Fibre Channel supports several different physical media: copper, multimode optical, and single-mode optical. All Fibre Channel protocols are supported on all media.

message signaled interrupts

See MSI, MSI-X.

MSI, MSI-X

Message signaled interrupts. An alternate way of generating an interrupts with special messages to allow PCI to emulate a pin assertion or deassertion. Message signaled interrupts allow the device to write a small amount of data to a special address in memory space. The chipset will deliver the corresponding interrupt to a CPU. MSI-X (defined in PCI 3.0) allows a larger number of interrupts (up to 2048), and gives each one a separate target address and data word.

N_Port

Node port. A port that connects by a point-to-point link to either a single N_Port or a single F_Port. N_Ports handle creation, detection, and flow of message units to and from the connected systems. N_Ports are end ports in virtual point-to-point links through a fabric, for example N_Port to F_Port to F_Port to N_Port using a single Fibre Channel fabric switch. See also FL_Port.

N_Port ID virtualization

See NPIV.

NIC

Network interface card or network interface controller. A computer circuit board or card that is installed in a computer so that it can be connected to a network.

NL_Port

Node loop port. A port capable of arbitrated loop functions and protocols. An NL Port connects through an arbitrated loop to other NL Port and at most a single FL Port. NL Ports handle creation, detection, and flow of message units to and from the connected systems. NL Ports are end ports in virtual point-to-point links through a fabric, for example NL Port to F Port to F Port to N Port using a single Fibre Channel fabric switch. In the absence of a fabric switch FL Port. NL Ports can communicate with other NL Ports in virtual point-to-point links through a FC AL open loop circuit often through FC AL (arbitrated Loop) hub or loop switch devices. See: E Port, F Port, FL Port, G Port, N Port.

node loop port

See NL_Port.

node port

See N_Port.

NPIV

The ability for a single physical Fibre Channel end point (N_Port) to support multiple, uniquely addressable, logical end points. With NPIV, a host Fibre Channel Adapter is shared in such a way that each virtual adapter is assigned to a virtual server and is separately identifiable within the fabric. Connectivity and access privileges within the fabric are controlled by identification of each virtual adapter and, hence, the virtual server using each virtual adapter.

NVRAM

Non-volatile random access memory. A type of memory that retains data (configuration settings) even when power is removed. You can manually configure NVRAM settings or restore them from a file.

path

A path to a device is a combination of a adapter port instance and a target port as distinct from internal paths in the fabric network. A fabric network appears to the operating system as an opaque network between the adapter (initiator) and the target.

Because a path is a combination of an adapter and a target port, it is distinct from another path if it is accessed through a different adapter or it is accessing a different target port. Consequently, when switching from one path to another, the driver might be selecting a different adapter (initiator), a different target port, or both.

This is important to the driver when selecting the proper method of failover notification. It can make a difference to the target device, which might have to take different actions when receiving retries of the request from another initiator or on a different port.

PLOGI

Port login. A port login occurs in a Fibre Channel SAN when two node ports establish a connection between each other (typically a device such as a Fibre Channel Adapter connecting to a switch).

point-to-point

Also FC-P2P. Two Fibre Channel nodes directly connected (not in a loop).

port

Access points in a device where a link attaches. The most common port types are:

- N_Port—a Fibre Channel port that supports point-to-point topology.
- NL_Port—a Fibre Channel port that supports loop topology.
- F_Port—a port in a fabric where an N_Port can attach.
- FL_Port—a port in a fabric where an NL_Port can attach.

port instance

The number of the port in the system. Each adapter may have one or multiple ports, identified with regard to the adapter as port 0, port 1 and so forth. to avoid confusion when dealing with a system containing numerous ports, each port is assigned a port instance number when the system boots up. So Port 0 on an adapter might have a port instance number of, for example, 8 if it is the eighth port discovered by the system.

port login

See PLOGI.

QoS

Quality of service. Refers to the methods used to prevent bottlenecks and ensure business continuity when transmitting data over virtual ports by setting priorities and allocating bandwidth.

quality of service

See QoS.

registered state change notification

See RSCN.

RSCN

Registered state change notification (RSCN) is a Fibre Channel fabric notification sent to all specified nodes when any major fabric changes occur. This notification allows nodes to immediately gain knowledge about the fabric and react accordingly.

SAN

Storage area network. Multiple storage units (disk drives) and servers connected by networking topology.

SCSI

Small computer system interface. A high-speed interface used to connect devices—such as hard drives, CD drives, printers, and scanners— to a computer. The SCSI can connect many devices using a single controller. Each device is accessed by an individual identification number on the SCSI controller bus.

small computer system interface

See SCSI.

target

The storage-device endpoint of a SCSI session. Initiators request data from targets (usually disk-drives, tape-drives, or other media devices). Typically, a SCSI peripheral device is the target but an adapter may, in some cases, be a target. A target can contain many LUNs.

A target is a device that responds to a request by an initiator (the host system). Peripherals are targets, but for some commands (for example, a SCSI COPY command), the peripheral may act as an initiator.

UEFI

Unified extensible firmware interface. A specification detailing an interface that helps hand off control of the system for the pre-boot environment (that is, after the system is powered on, but before the operating system starts) to an operating system, such as Windows or Linux. UEFI provides a clean interface between operating systems and platform firmware at boot time, and supports an architecture-independent mechanism for initializing add-in cards.

unified extensible firmware interface

See UEFI.

USCM

SAN congestion management (SCM) is a common noun, and describes a standards-based Fibre Channel technology. Universal SAN Congestion Mitigation (USCM) is Marvell's IP, and describes Marvell's SCM feature set.

world wide port name

See WWPN.

WWPN

World wide port name. Unique 64-bit address assigned to each port on a device. One WWNN may contain multiple WWPN addresses.



Marvell first revolutionized the digital storage industry by moving information at speeds never thought possible. Today, that same breakthrough innovation remains at the heart of the company's storage, networking and connectivity solutions. With leading intellectual property and deep system-level knowledge, Marvell semiconductor solutions continue to transform the enterprise, cloud, automotive, industrial, and consumer markets. For more information, visit <u>www.marvell.com</u>.

© 2024 Marvell. All rights reserved. The MARVELL mark and M logo are registered and/or common law trademarks of Marvell and/or its Affiliates in the US and/or other countries. This document may also contain other registered or common law trademarks of Marvell and/or its Affiliates.

Doc No. SN0154529-00 Rev. Z Revised: June 10, 2024