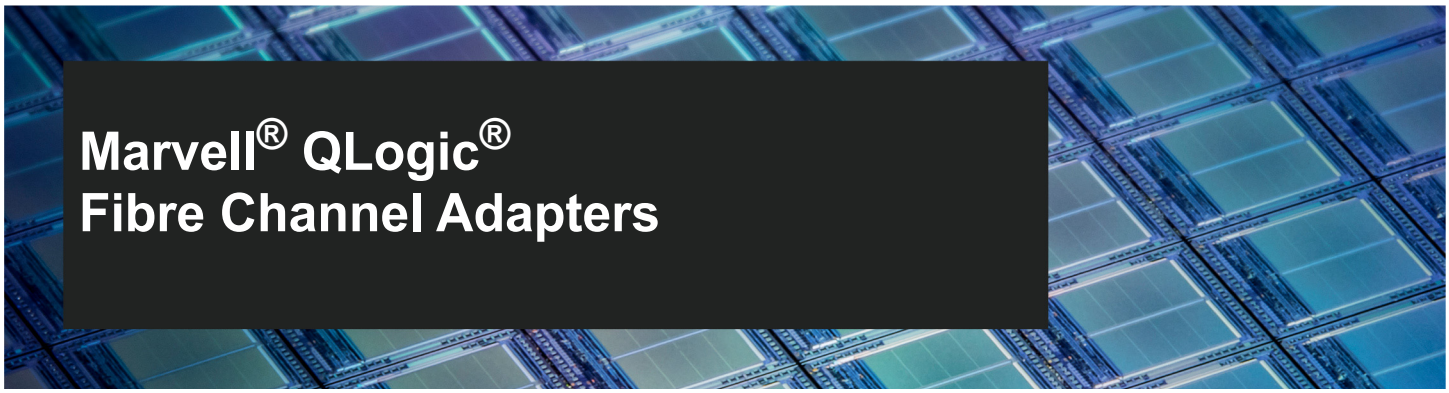

Marvell® QLogic® Fibre Channel Adapters

QLE2660-DEL, QLE2662-DEL, QME2662-DEL,
QLE2690-DEL, QLE2690L-DEL,QLE2690N-DEL, QLE2690NL-DEL,
QLE2692-DEL, QLE2692L-DEL, QLE2692N-DEL, QLE2692NL-DEL, QME2692-DEL

User's Guide

Third party information brought to you courtesy of Dell.

Doc. No. CU0354604-00 Rev. ZD
October 21, 2022

Table of Contents

Introduction

This introductory chapter provides a list of the covered models, describes the intended audience and contents of this guide, lists related documents and the document conventions, describes the product functionality and features, and lists the supported OSs.

This user's guide covers the following products:

- QLogic® QLE2660-DEL single-port, low profile adapter with a full-height bracket
- QLogic QLE2662-DEL dual-port, low profile adapter with a full-height bracket
- QLogic QME2662-DEL dual-port mezzanine adapter
- QLogic QLE2690-DEL, QLE2690N-DEL[1]: single-port, low profile adapters with a full-height bracket
- QLogic QLE2690L-DEL, QLE2690NL-DEL[1]: single-port, low profile adapters with a low profile bracket
- QLogic QLE2692-DEL, QLE2692N-DEL[1]: dual-port, low profile adapters with a full-height bracket
- QLogic QLE2692L-DEL, QLE2692NL-DEL[1]: dual-port, low profile adapters with a low profile bracket
- QLogic QME2692-DEL dual-port mezzanine adapter

NOTE

Throughout this document, the term *adapter* refers to any or all of these products.

This guide provides technical information about the adapters, including how to install and configure the adapter, as well as detailed descriptions of the adapter's various uses and functions.

[1] The QLE2690N-DEL, QLE2690NL-DEL, QLE2692N-DEL, and QLE2692NL-DEL adapters do not have a tFRU component.

Intended Audience

This guide is intended for system administrators and other technical staff members responsible for configuring and managing adapters installed on Dell® PowerEdge® servers in Windows®, Linux®, or VMware® environments.

User's Guide Content

This guide provides information in the following chapters and appendices:

- Chapter 1 Hardware Installation covers the hardware and software requirements, safety precautions, a pre-installation checklist, PCI Express® (PCIe®) slot considerations, and procedures for installing the adapter and connecting it to the network.
- Chapter 2 Driver Installation and Configuration covers the installation of the drivers included with the adapter on Windows, Linux, and VMware OSs.
- Chapter 3 Fibre Channel Configuration provides information about the multiboot image (Fibre Channel Adapter Function Configuration Utility) and instructions for setting Fibre Channel Adapter parameters; persistent binding; configuring the boot device, N_Port ID virtualization (NPIV), and driver parameters; and reassembling out-of-order frames.
- Chapter 4 NVM Express over Fibre Channel provides information on installing and configuring NVM Express on selected Fibre Channel Adapters.
- Appendix A Troubleshooting provides information about Fibre Channel diagnostics and a Fibre Channel troubleshooting diagram.
- Appendix B Specifications defines the physical characteristics and power requirements and lists supported standards, interface, and environmental specifications.
- Appendix C Regulatory Information provides warranty, regulatory, and compliance information.
- Appendix D Adapter LEDs describes the adapter port LEDs.
- Appendix E Revision History describes the changes made in this revision of the guide.

Related Materials

For additional information, refer to the following documents:

- *User's Guide—QConvergeConsole CLI* (part number SN0054667-00) provides specific command line use in both interactive and noninteractive modes.[1]
- *User's Guide—QConvergeConsole Plug-ins for vSphere* (part number TD-000965) provides reference material on using the QConvergeConsole VMware vCenter Server Plug-in and the QConvergeConsole VMware vSphere Web Client Plug-in.[1]
- *User's Guide—Marvell QLogic Fibre Channel QConvergeConsole Extension for Windows® Admin Center* (part number TD-001728) provides information on installing and using the Marvell QLogic Fibre Channel (FC) QConvergeConsole Extension for Windows Admin Center (WAC) to manage the Marvell QLogic 2500/2600/2700 Series Adapters.[1]
- *User's Guide—Marvell® QLogic® PowerKit: Windows® PowerShell® cmdlets for 2500, 2600, 2700 Series Marvell QLogic Fibre Channel Adapters* (part number TD001187) provides information on installing the Marvell QLogic Fibre Channel (MRVLFC) PowerKit (Windows® PowerShell® cmdlets) software to manage Marvell QLogic Fibre Channel (FC) adapters.[1]
- *User's Guide—Marvell ESXCLI Plug-in for VMware vSphere* (part number BC0151101-00) provides information on installing and using the Marvell® ESXCLI VMware® Plug-in. The ESXCLI plug-in extends the capabilities of the ESX® command line interface to manage the following adapters installed in VMware ESX/ESXi hosts.[1]

NOTE

To access Marvell product documents, visit the Marvell Web site, www.marvell.com.

Documentation Conventions

This guide uses the following documentation conventions:

- **NOTE** provides additional information.
- **CAUTION** without an alert symbol indicates the presence of a hazard that could cause damage to equipment or loss of data.

[1] Describes how to configure and use Universal SAN Congestion Mitigation (USCM).

- **WARNING** indicates the presence of a hazard that could cause serious injury or death.
- 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 underlined blue. For example:
  - Table 9-2 lists problems related to the user interface and remote agent.
  - See "Installation Checklist" on page 3-6.
  - For more information, visit www.marvell.com.
- Text in **bold** font indicates user interface elements such as a menu items, buttons, check boxes, or column headings. For example:
  - Click **Start**, 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 command line text. For example:
  - To return to the root directory from anywhere in the file structure, type **`cd /root`**, and then press the ENTER key.
  - Issue the following command: `# sh ./install.bin`
- Key names and key strokes are indicated with UPPERCASE:
  - Press the CTRL+P keys.
  - 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 *QLogic 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 *QConvergeConsole Help* throughout this document.

- QConvergeConsole CLI non-interactive 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:
    - `<hba instance>`

    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`

Functionality and Features

This section provides the following information:

Functional Description

Functional descriptions for the adapters are as follows:

- **QLE2660-DEL:** A low profile, 16GFC, single-port Fibre Channel PCIe adapter.
- **QLE2662-DEL:** A low profile, 16GFC, dual-port Fibre Channel PCIe adapter.
- **QME2662-DEL:** A 16GFC, dual-port Fibre Channel mezzanine adapter for the blade server environment.
- **QLE2690-DEL:** A low profile, Enhanced 16GFC, single-port Fibre Channel-to-PCIe adapter with a full-height bracket installed.
- **QLE2690L-DEL:** A low profile, Enhanced 16GFC, single-port Fibre Channel-to-PCIe adapter with a low profile bracket installed.
- **QLE2690N-DEL**: A low profile, Enhanced 16GFC, single-port Fibre Channel-to-PCIe adapter without a tFRU component and with a full-height bracket installed.
- **QLE2690NL-DEL**: A low profile, Enhanced 16GFC, single-port Fibre Channel-to-PCIe adapter without a tFRU component and with a low profile bracket installed.
- **QLE2692-DEL:** A low profile, Enhanced 16GFC, dual-port Fibre Channel-to-PCIe adapter with a full-height bracket installed.
- **QLE2692L-DEL:** A low profile, Enhanced 16GFC, dual-port Fibre Channel-to-PCIe adapter with a low profile bracket installed.
- **QLE2692N-DEL:** A low profile, Enhanced 16GFC, dual-port Fibre Channel-to-PCIe adapter without a tFru component and with a full-height bracket installed.
- **QLE2692NL-DEL:** A low profile, Enhanced 16GFC, dual-port Fibre Channel-to-PCIe adapter without a tFru component and with a low profile bracket installed.
- **QME2692-DEL:** An Enhanced 16GFC, dual-port, Fibre Channel mezzanine adapter for the blade server environment

Key Features

Key features of the adapters include:

- Centralized device management for SAN
- Universal SAN Congestion Mitigation (USCM)—QLE269x only
  - FPIN Monitoring
  - Profiles

- Connectivity to 16GFC/8GFC/4GFC[1] Fibre Channel networks
- PCIe 3.0 x8 and PCIe 2.0 *x*8
- Full hardware offload for the Fibre Channel protocol
- Message-signaled interrupts (MSI-X) and legacy interrupts (INT-X)
- NPIV
- Boot from SAN
- FC-NVMe (QLE269*x* only)
- Several advanced management features for 2600 Series Adapters:
  - QConvergeConsole CLI is available if you are running Windows or Linux.
  - QConvergeConsole VMware vCenter Server Plug-in is available if you are running VMware ESXi.

Supported Operating Systems

NOTE

Because the *Dell Update Packages Version xx.xx.xxx User's Guide* is not updated in the same cycle as this Fibre Channel adapter user's guide, consider the operating systems listed in this section as the more current.

The adapter supports the following operating systems. To view the most complete and current list, refer to the product release notes.

Microsoft

- Windows Server® 2019
- Windows Server 2022

Red Hat

- Red Hat® Enterprise Linux *(*RHEL®*)* 8.5
- RHEL 8.6
- RHEL 9.0

SuSE

- SuSE® Linux Enterprise Server (SLES®) 15 SP3
- SLES 15 SP4

[1] 4Gb does not apply to the QME2662-DEL and QME2692-DEL; 8Gb does not apply to QME2692-DEL.

VMware

- vSphere: VMware ESXi 7.0 U3
- vSphere: VMware ESXi 8.0

Citrix

- Citrix® Hypervisor 8.2 LTSR

NOTE

For the most current versions of the OS and drivers supported by the adapter, refer to the release notes file (`release.txt`).

1 Hardware Installation

This chapter provides the hardware and software requirements, safety precautions, a pre-installation checklist, PCIe slot considerations, and procedures for installing the adapter and connecting it to the network.

Hardware and Software Requirements

Before you install the adapter, verify that the system meets the following hardware and software requirements.

- **Hardware:**
  - For QLE2660-DEL, QLE2662-DEL, QLE2690-DEL, QLE2690L-DEL, QLE2690N-DEL, QLE2690NL-DEL, QLE2692-DEL, QLE2692L-DEL, QLE2692N-DEL, QLE2692NL-DEL adapter port and slot assignments, refer to the “Expansion Cards” section of the *Hardware Owner’s Manual* for your Dell PowerEdge server.
  - For QME2662-DEL adapter port and slot assignments, refer to the blade and M1000e chassis diagram in the *Dell PowerEdge M1000e Systems Configuration Guide*.
  - For QME2692-DEL adapter port and slot assignments, refer to the blade and MX7000 chassis diagram in the *Dell PowerEdge MX7000 Enclosure Installation and Service Manual*.
- **Software:** For information on the supported operating systems, firmware versions, adapter drivers, and utilities, refer to the product release notes.

Safety Precautions

⚠ WARNING

The adapter is being installed in a system that operates with voltages that can be lethal. Before you open the case of your system, observe the following precautions to protect yourself and to prevent damage to the system components.

For your safety, follow these precautions:

- Remove any metallic objects or jewelry from your hands and wrists.
- Make sure to use only insulated or nonconducting tools.
- Before you touch internal components, verify that the system is powered OFF and is unplugged.
- Install or remove adapters in a static-free environment. The use of a properly grounded wrist strap or other personal anti-static devices and an anti-static mat is strongly recommended.

Pre-Installation Checklist

1. Verify that your system meets the hardware and software requirements listed in “Hardware and Software Requirements” on page 1.
2. Verify that your system is using the latest BIOS.

NOTE

If you acquired the adapter software on a disk or from the Dell support Web site (http://support.dell.com), verify the path to the adapter driver files.

3. Check the adapter for visible signs of damage. Never attempt to install a damaged adapter.

PCIe Bus Slot Considerations

The size of the PCIe bus slot in which you install the adapter will affect the speed throughput. Table 1-1 lists the approximate speed throughput for each bus slot size.

Table 1-1. Approximate Speed Throughput per Bus Slot Size

Slot Size	PCIe Generation (Speed) Throughput		
	PCIe Gen1 (2.5 GT/s)	PCIe 2.0 (5 GT/s)	PCIe 3.0 (8 GT/s)
x8 slot	~20GFC (2.5 × 8)	~40GFC (5 × 8)	~64GFC (8 × 8)
x16 slot	~40GFC (2.5 × 16)	~80GFC (5 × 16)	~128GFC (8 × 16)

Table 1-2 lists the throughput needed to achieve line rates.

Table 1-2. Throughput Requirements to Achieve Line Rate

Number of 16G FC Ports	Model Number	PCIe Throughput Required to Achieve Line Rate	Minimum PCIe Gen and Slot Configurations
1	QLE2660_DEL, QLE2690-DEL, QLE2690L-DEL, QLE2690N-DEL, QLE2690NL-DEL	16GFC (1×16GFC)	PCIe 2.0 x8 and higher PCIe 3.0 x8 and higher
2	QLE2662-DEL, QLE2692-DEL, QLE2692L-DEL, QLE2692N-DEL, QLE2692NL-DEL, QME2692-DEL	32GFC (2×16GFC)	PCIe 2.0 x8 and higher PCIe 3.0 x8 and higher

Installing the Adapter

Follow the instructions for your adapter model number.

QLE2660-DEL, QLE2662-DEL, QLE2690-DEL, QLE2690L-DEL, QLE2690N-DEL, QLE2690NL-DEL, QLE2692-DEL, QLE2692L-DEL, QLE2692N-DEL, QLE2692NL-DEL

To install the QLE2600-DEL Series Adapters:

1. Power off the computer and all attached devices such as monitors, printers, and external components.
2. Disconnect the power cable.
3. Remove the computer cover and find an empty PCIe bus slot. For more information about selecting a bus slot, see “PCIe Bus Slot Considerations” on page 2.
4. Pull out the slot cover (if any).
5. Grasp the adapter by the top edge and seat it firmly into the appropriate slot.
6. Refasten the adapter’s retaining bracket.
7. Close the computer cover.
8. Plug the Fibre Channel cable into the adapter.
9. Plug in the power cable and turn on the computer.

For more detailed information, refer to the *Hardware Owner's Manual* for your Dell PowerEdge server.

QME2662-DEL

For installation instructions, refer to the "I/O Module Mezzanine Cards" and "Guidelines for Installing I/O Modules" sections of the *Dell PowerEdge Modular Systems Hardware Owner's Manual*.

QME2692-DEL

For installation instructions, refer to the *Dell EMC PowerEdge MX7000 Enclosure Installation and Service Manual*.

Connecting to the SAN

Follow the instructions for your adapter model number.

QLE2600-DEL

To connect the QLE2600-DEL Series Adapters to the SAN, refer to the *Hardware Owner's Manual* for your Dell PowerEdge server.

QME2662-DEL

To connect to the SAN, refer to the "Guidelines for Installing I/O Modules" section of the *Dell PowerEdge Modular Systems Hardware Owner's Manual*.

QME2692-DEL

To connect to the SAN, refer to the *Dell EMC PowerEdge MX7000 Enclosure Installation and Service Manual*.

2 Driver Installation and Configuration

NOTE

If you need to update the Flash memory of multiple adapters simultaneously, see “Updating the Dell Firmware” on page 31.

This chapter provides the following information about the drivers included with the adapters:

- “Windows Driver Installation and Configuration” on page 5
- “Linux Driver Installation and Configuration” on page 12
- “VMware Driver Installation and Configuration” on page 17

NOTE

When you disable the firmware (for example, during a firmware dump or during a firmware update) in Windows or Linux with a QConvergeConsole agent, multiple application messages are generated. These messages are generated because the application cannot communicate with the adapter while the firmware is disabled. After the firmware is re-enabled, the errors will go away.

Windows Driver Installation and Configuration

NOTE

If you are using the 2600 Series Adapters to boot from storage in a Windows Server 2019 or Windows Server 2022 environment, your adapters may not recognize the storage targets and LUNs when you load the adapter driver the first time. For the adapter to recognize the storage and LUNs, you must load the adapter drivers a second time.

You can run a software or driver Dell Update Package (DUP) in two ways:

- Running the Dell Update Package in the GUI
- Running the Dell Update Package from the Command Line

Running the Dell Update Package in the GUI

Before you begin, refer to the *Dell Update Packages Version xx.xx.xxx User's Guide*, "Prerequisites and Features for Systems Running Windows" section.

To run the DUP in the GUI:

1. Double-click the icon representing the DUP file. (The actual file name of the DUP varies.)
2. In the Dell Update Package window (Figure 2-1), click **Install**.

Figure 2-1. Dell Update Package Window

3. In the QLogic Super Installer—InstallShield® Wizard's Welcome window (Figure 2-2), click **Next**.

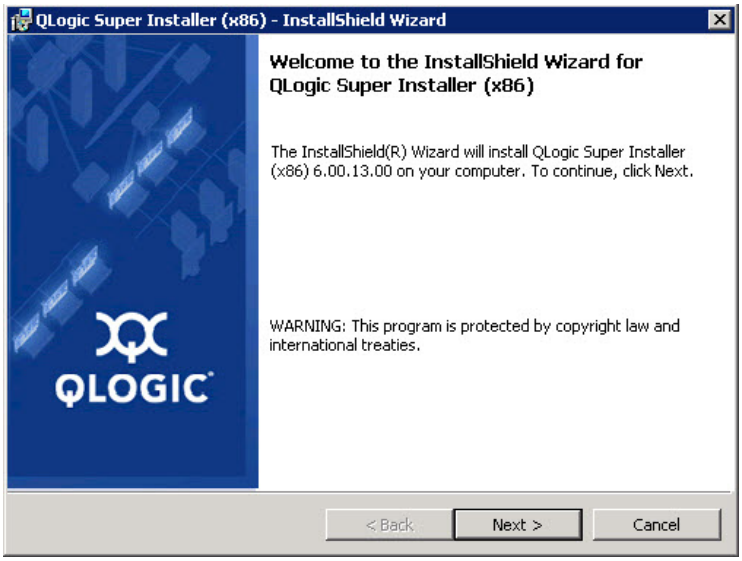

Figure 2-2. QLogic InstallShield Wizard: Welcome Window

4. In the wizard's License Agreement window (Figure 2-3):
   a. Read the QLogic End User Software License Agreement.
   b. To continue, select **I accept the terms in the license agreement**.
   c. Click **Next**.

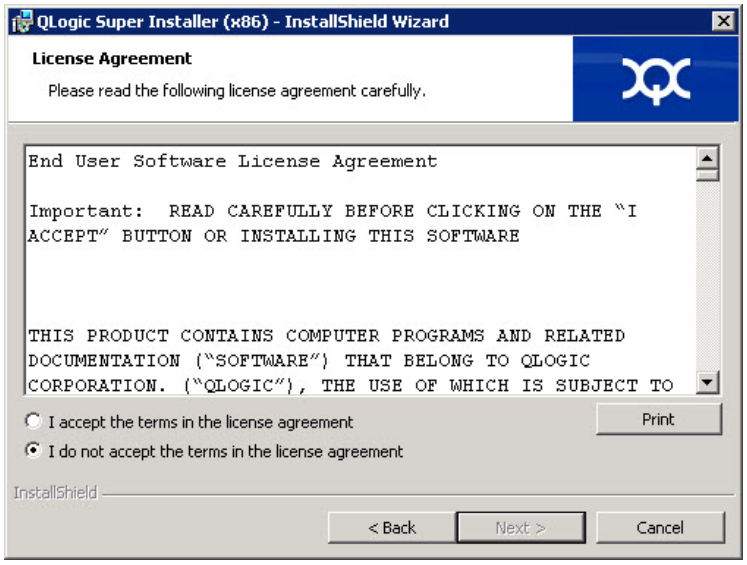

Figure 2-3. QLogic InstallShield Wizard: License Agreement Window

5. Complete the wizard's Setup Type window (Figure 2-4) as follows:
   a. Select one of the following setup types:
      - Click **Complete** to install *all* program features.
      - Click **Custom** to manually select the features to be installed.
   b. To continue, click **Next**.

      If you selected **Complete**, proceed directly to Step 6 b.

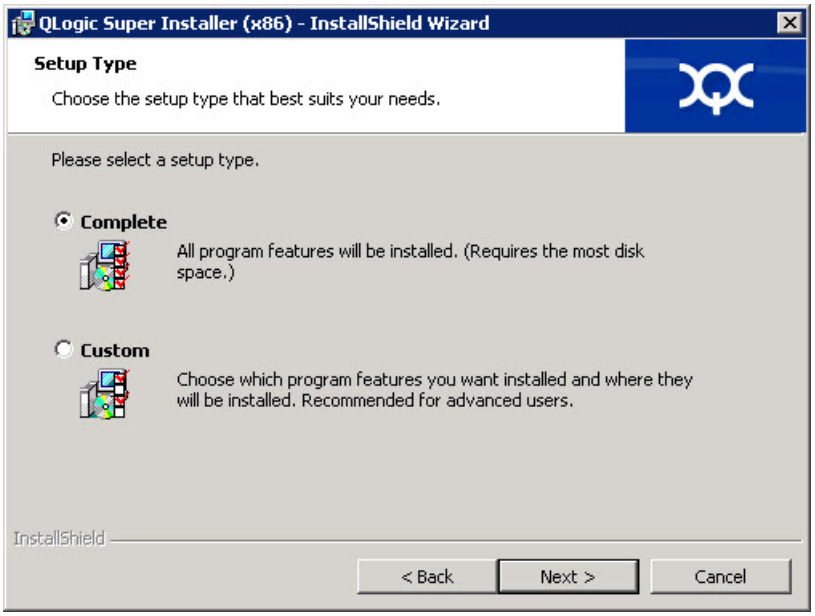

Figure 2-4. InstallShield Wizard: Setup Type Window

6. If you selected **Custom** in Step 5, complete the Custom Setup window (Figure 2-5 on page 9) as follows:
   a. Select the features to install. By default, all features are selected. To change a feature's install setting, click the icon next to it, and then select one of the following options:
      - **This feature will be installed on the local hard drive**—This setting marks the feature for installation without affecting any of its subfeatures.
      - **This feature, and all subfeatures, will be installed on the local hard drive**—This setting marks the feature and all of its subfeatures for installation.
      - **This feature will not be available**—This setting prevents the feature from being installed.

b. Click **Next** to continue.

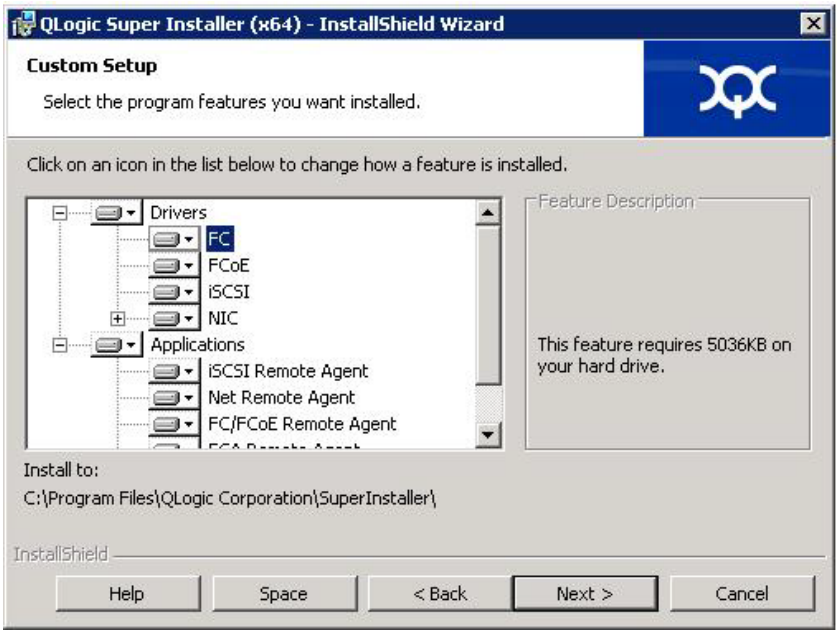

Figure 2-5. InstallShield Wizard: Custom Setup Window

7. In the InstallShield Wizard's Ready To Install window (Figure 2-6), click **Install**.

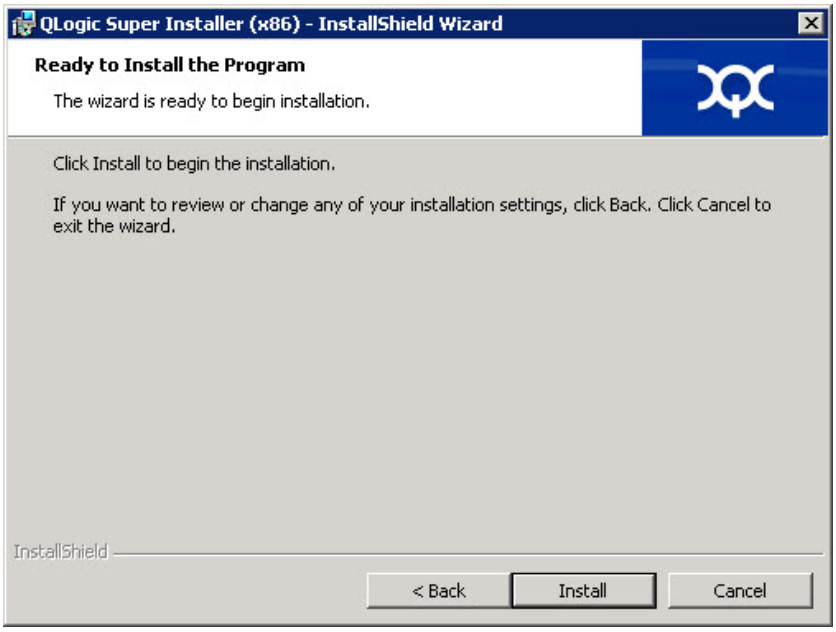

Figure 2-6. InstallShield Wizard: Ready to Install the Program Window

The InstallShield Wizard installs the QLogic Adapter drivers and Management Software Installer.

8. When the installation is complete, the InstallShield Wizard Completed window appears (Figure 2-7). Click **Finish** to dismiss the installer.

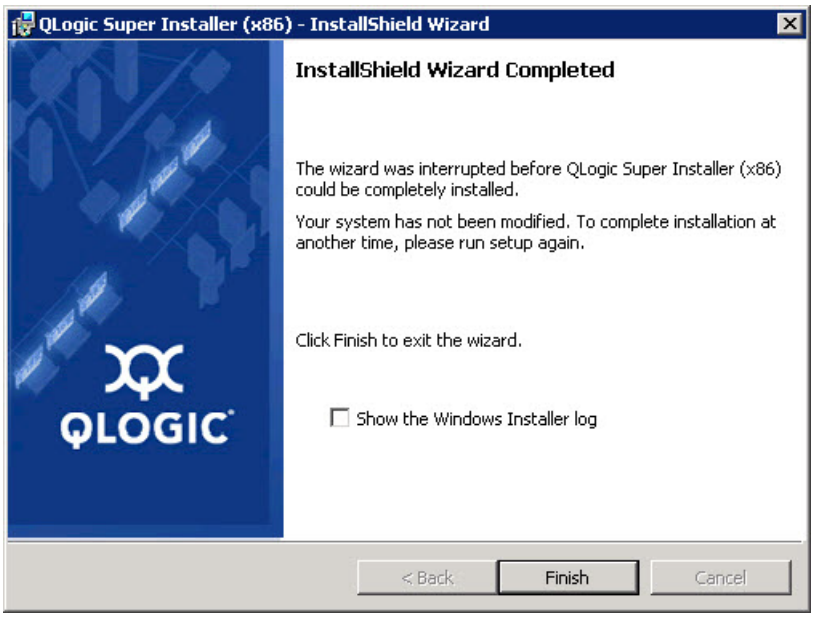

Figure 2-7. InstallShield Wizard: Completed Window

9. In the Dell Update Package (DUP) window (Figure 2-8), “Complete” indicates successful installation.

   - ❑ (Optional) To open the log file, click **View Log**. The log file shows the progress of the DUP installation, any previous installed versions, any error messages, and other information about the installation.
   - ❑ To close the Update Package window, click **OK**.

Figure 2-8. Dell Update Package Window

Running the Dell Update Package from the Command Line

For a list of the CLI options for systems running Windows, a description of each option, and the command syntax, refer to the *Dell Update Packages Version xx.xx.xxx User’s Guide*, “Windows CLI Options” section.

Running the DUP from the command line, with no options specified, results in the same behavior as double-clicking the icon representing the DUP.

NOTE

The actual file name of the DUP varies.

To run the DUP from the command line:

```
C:\><DUP_file_name>.EXE
```

The following shows the syntax for specifying options to customize the Dell update package installation behavior:

```
<DUP_file_name>.exe [/<option1>[=<value1>]] [/<option2>[=<value2>]]...
```

To display the GUI for guided installation, update, or extraction, use no options.

Examples

To update the system silently:

```
<DUP_file_name>.exe /s
```

To extract the update contents to the `C:\mydir\` directory:

```
<DUP_file_name>.exe /s /e=C:\mydir
```

To extract the driver components to the `C:\mydir\` directory:

```
<DUP_file_name>.exe /s /drivers=C:\mydir
```

To install only the driver components:

```
<DUP_file_name>.exe /s /driveronly
```

To change from the default log location to `C:\my path with spaces\log.txt`:

```
<DUP_file_name>.exe /l="C:\my path with spaces\log.txt"
```

Linux Driver Installation and Configuration

This section provides the following procedures for installing drivers on a Linux system:

- Installation Overview
- Installing the Linux Fibre Channel Driver

Installation Overview

To install and configure the adapter drivers on a Linux system, refer to the driver release notes, readme, and installation instructions included in the package.

NOTE

To install the Red Hat Package Manager (RPM), issue the following command as a root user:

```
# rpm -Uvh <rpm name>
```

For example:

```
# rpm -Uvh qla2xxx-kmp-default-
<driver-version_kernel-version>-<release>.x86_64.rpm
```

To uninstall the RPM, issue the following command as a root user:

```
# rpm -e <rpm>
```

For example:

```
# rpm -e qla2xxx-kmp-default-
<driver-version_kernel-version>-<release>
```

Importing a Public Key for Secure Boot

Linux drivers require that you import and enroll the QLogic public key to load the drivers in a Secure Boot environment. Before you begin, ensure that your server supports Secure Boot.

This section is for use with `qla2xxx-kmp-default-<driver-version_kernel-version>-<release>.rpm` package.

NOTE

A public key is required on a Linux virtual machine running on a hypervisor (if secure boot is enabled).

This section provides two methods for importing and enrolling the public key.

To import and enroll the QLogic public key:

1. Download the public key from the following Web page:

   https://ldriver.qlogic.com/Module-public-key

   A warning appears, indicating that the connection is not private. Follow the next steps or equivalent steps to continue.

   a. Click the **Advanced** button.

   b. Click **Proceed to ldriver.qlogic.com (unsafe)**.

2. To install the public key, issue the following command:

   ```
   # mokutil --root-pw --import <cert name>.der
   ```

   Where the `--root-pw` option enables direct use of the root user.

3. Reboot the system.

4. Review the list of certificates that are prepared to be enrolled:

   ```
   # mokutil --list-new
   ```

5. Reboot the system again.

6. When the shim launches MokManager, enter the root password to confirm the certificate importation to the Machine Owner Key (MOK) list.

7. To determine if the newly imported key was enrolled:

   ```
   # mokutil --list-enrolled
   ```

To launch MOK manually and enroll the QLogic public key:

1. Issue the following command:

   ```
   # reboot
   ```

2. In the **GRUB 2** menu, press the C key.
3. Issue the following commands:

   ```
   chainloader $efibootdir/MokManager.efi
   - boot
   ```

4. Select **Enroll key from disk**.
5. Navigate to the `cert.der` file and then press ENTER.
6. Follow the instructions to enroll the key. Generally this includes pressing the 0 (zero) key and then pressing the Y key to confirm.

NOTE

The firmware menu may provide more methods to add a new key to the Signature Database.

For additional information about Secure Boot, refer to the following Web page:

https://www.suse.com/documentation/sled-12/book_sle_admin/data/sec_uefi_secboot.html

Installing the Linux Fibre Channel Driver

This section provides procedures for installing the Linux Fibre Channel driver for the following operating systems:

- Building the Driver for RHEL 8.0 and 9.0
- Building the Driver for SLES 15

Building the Driver for RHEL 8.0 and 9.0

The following instructions apply to the RHEL 8.5 (version 4.18.*x*), RHEL 8.6 (version 4.18.*x*), and RHEL 9.0 (version 5.14.*x*) operating systems.

1. From the directory that contains the source driver file, `qla2xxx-src-vx.xx.xx.xx.xx.x-k.tar.gz`, issue the appropriate commands.

   ```
   # tar -xzvf qla2xxx-src-vx.xx.xx.xx.xx.x-k.tar.gz
   # cd qla2xxx-src-vx.xx.xx.xx.xx.x-k
   ```

2. Build and install the driver modules from the source code by executing the `build.sh` script as follows:

   ```
   # ./extras/build.sh install
   ```

   The `build.sh` script does the following:

   - ❑ Builds the driver `.ko` files.
   - ❑ Copies the `.ko` files to the appropriate `/lib/modules/<kernel version>/extra/qlgc-qla2xxx` directory.
   - ❑ Adds the appropriate directive in the `modprobe.conf` (if applicable).

NOTE

You can optionally complete either or both Step 3 and Step 4 of this procedure.

3. (Optional) To manually load the driver, issue the `insmod` or `modprobe` command:

   - ❑ To directly load the driver from the local build directory, issue the following `insmod` commands in order:

     ```
     # modprobe scsi_tgt (if not already loaded)
     # modprobe scsi_transport_fc
     # insmod ./qla2xxx.ko
     ```

   - ❑ To load the driver using modprobe, issue the following command:

     ```
     # modprobe -v qla2xxx
     ```

   - ❑ To unload the driver using modprobe, issue the following command:

     ```
     # modprobe -r qla2xxx
     ```

4. (Optional) To automatically load the driver each time the system boots, rebuild the RAM disk to include the driver as follows:

   a. Create a backup copy of the RAMDISK image by issuing the following commands:

      ```
      # cd /boot
      # cp initramfs-[kernel version].img initramfs-[kernel version].img.bak
      ```

   b. Create the new RAMDISK by issuing the following command:

      ```
      # dracut -f
      ```

   c. To load the driver, reboot the host.

Building the Driver for SLES 15

The following instructions apply to SLES 15 SP3 (version 5.3.*x*) and SLES 15 SP4 (version 5.3.*x*) operating systems.

1. From the directory that contains the source driver file, `qla2xxx-src-vx.xx.xx.xx.xx.x-k.tar.gz`, issue the following commands:

   ```
   # tar -xzvf qla2xxx-src-vx.xx.xx.xx.xx.x-k.tar.gz
   # cd qla2xxx-src-vx.xx.xx.xx.xx.x-k
   ```

   Where *x.xx.xx.xx.xx* is the applicable version number.

2. Build and install the driver modules from the source code by executing the `build.sh` script as follows:

   ```
   # ./extras/build.sh install
   ```

   The `build.sh` script does the following:

   - ❑ Builds the driver `.ko` files.
   - ❑ Copies the `.ko` files to the appropriate `/lib/modules/<kernel version>/updates` directory.
   - ❑ Adds the appropriate directive in the `modprobe.conf` file (if applicable).

NOTE

You can optionally complete either or both Step 3 and Step 4 of this procedure.

3. (Optional) Manually load the driver for Linux.

   - ❑ Edit the `/etc/modprobe.d/unsupported_modules` file to make the following change:

     `allow_unsupported_modules` **`1`** (replace `0` with `1`)

   - ❑ To load the driver using `modprobe`, issue the following command:

     ```
     # modprobe -v qla2xxx
     ```

   - ❑ To unload the driver using modprobe, issue the following command:

     ```
     # modprobe -r qla2xxx
     ```

4. (Optional) To automatically load the driver each time the system boots, rebuild the RAM disk to include the driver as follows:

   a. Create a backup copy of the RAMDISK image by issuing the following commands:

   ```
   # cd /boot
   # cp initramfs-[kernel version].img initramfs-[kernel version].img.bak
   ```

   b. Create the new RAMDISK by issuing the following command:

   ```
   # dracut -f
   ```

5. To load the driver, reboot the host.

VMware Driver Installation and Configuration

This section provides the following procedures for installing drivers on a VMware system:

- Installation Overview
- Installing the ESXi 7.0 U3 and ESXi 8.0 Fibre Channel Driver
- Installing the QConvergeConsole VMware vCenter Server Plug-in
- Configuring VM-ID

Installation Overview

To install and configure the adapter drivers on a VMware system, refer to the driver release notes and readme files included in the package.

Installing the ESXi 7.0 U3 and ESXi 8.0 Fibre Channel Driver

The operating system manages and controls the driver installation process. To install the ESXi driver, follow the steps in this section.

NOTE

This section provides the most common ways of installing and upgrading the driver. For other installation procedures, consult the operating system manuals and the driver readme file.

This section provides procedures for the following:

- “Updating an Existing Driver or Installing a New Driver for an Existing ESXi 7.0 and 8.0 Installation with esxcli” on page 18
- “Verifying the Version of the Installed Driver” on page 18

Updating an Existing Driver or Installing a New Driver for an Existing ESXi 7.0 and 8.0 Installation with esxcli

To use the driver bundle <offline-bundle>.zip):

1. Copy the driver bundle (`<offline-bundle>.zip`) to this ESXi host.
2. Install the driver bundle (`<offline-bundle>.zip`) as follows:
   a. To create a temporary directory, issue the following commands:
   ```
   mkdir /install
   cd /install
   ```
   b. Unzip the driver bundle in the temporary directory:
   ```
   /install : unzip <offline-bundle>.zip
   ```
   c. Issue the following command:
   ```
   esxcli software vib install -n qlnativefc -d /install
   ```

To use the driver VIB:

1. Copy the driver VIB to this ESXi host by issuing the following command:
   ```
   qlnativefc-<driver-version>-1OEM.<esx-build>.x86_64.vib
   ```
2. Install the driver VIB using the following `esxcli` commands:
   a. To make a temporary directory, issue the following commands:
   ```
   mkdir /install
   cd /install
   ```
   b. Issue the following command:
   ```
   esxcli software vib install -v /install/<driver-vib>
   ```

Verifying the Version of the Installed Driver

To verify the installed package in the system, issue the following command:

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

The driver version is embedded in the VIB version.

For example, the output looks like the following:

```
# esxcli software vib list | grep qlnativefc
qlnativefc                    2.1.50.0-1OEM.600.0.0.2768847
QLogic  VMwareCertified   2017-01-19
```

Installing the QConvergeConsole VMware vCenter Server Plug-in

The following installation sections apply to VI Server and HTML5 Client plug-ins.

It is advised to use HTML5 Plug-in for ESXi 6.5 and later.

To use QConvergeConsole VMware vCenter Server Plug-in, install the following software in this order:

1. **QConvergeConsole VMware vCenter Server Plug-in** on the vCenter Server
2. **QLogic Adapter CIM Provider** on the ESXi server

The following sections explain how to install and uninstall the required software:

- Installation Package Contents
- Installing the QConvergeConsole VMware vCenter Server Plug-in
- Plug-in Unregistration from a Manual Install
- Uninstalling the QConvergeConsole VMware vCenter Server Plug-in
- Installing the QLogic Adapter CIM Provider
- Uninstalling the QLogic Adapter CIM Provider

For information on installing the Plug-in, refer to “Installing the QConvergeConsole VMware vCenter Server Plug-in” on page 20.

Installation Package Contents

The latest version of the QLogic Adapter CIM Provider and QConvergeConsole VMware vCenter Server Plug-in package contains the files needed to install both the Plug-in and the CIM Provider. Required files include the following (where *x_x_x* is the version number):

- `QLogic_Adapter_VI_Plugin_x_x_x.exe`—This file is the QConvergeConsole VMware vCenter Server Plug-in installation package.
- `QLogic_Adapter_vSphere_Client_Plugin_<ver_num>.exe`— This file is the HTML5-based vSphere Client Plug-in installation package (for HTML5-based vSphere Client)
- `qlogic_adapter_provider_vmware_esx67-x.x.x`—This file contains the QLogic Adapter CIM Provider installation package for ESXi 6.7 U3, where *x.x.x* is the version of the CIM Provider.
- `qlogic_adapter_provider_vmware_esx65-x.x.x`—This file contains the QLogic Adapter CIM Provider installation package for ESXi 6.5 U3, where *x.x.x* is the version of the CIM Provider.

- `qlogic_adapter_provider_vmware_esx70-x.x.x`—This file contains the QLogic Adapter CIM Provider installation package for ESXi 7.0, where *`x.x.x`* is the version of the CIM Provider.
- `readme.txt`—This file is the Read Me document that covers hardware and software requirements, operating system support, supported features, installation and removal instructions, known issues and workarounds, and support contact information.
- `release_notes.txt`—This file is the Release Notes document that lists changes, fixes, known issues, and release details.

For detailed information on installing the QConvergeConsole VMware vCenter Server Plug-in, refer to "Installing the QConvergeConsole VMware vCenter Server Plug-in" on page 20.

For detailed information on installing the CIM Provider, refer to "Installing the QLogic Adapter CIM Provider" on page 25.

Installing the QConvergeConsole VMware vCenter Server Plug-in

To install the QConvergeConsole VMware vCenter Server Plug-in:

1. Download the `QLogic_Adapter_VI_Plugin_x_x_x.exe` file (where `x_x_x` is the version number).
2. Start the installation either by double-clicking the `.exe` file, by typing the name of the `.exe` file in a Run window, or by clicking **Browse** and locating the `.exe` file.

   The InstallAnywhere wizard opens, as shown in Figure 2-9.

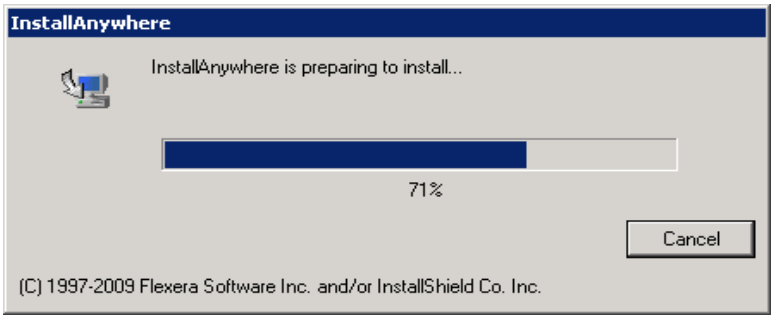

Figure 2-9. InstallAnywhere Initial Window

3. In the QLogic Adapter VI Plug-in wizard's Introduction window (see Figure 2-10), click **Next**.

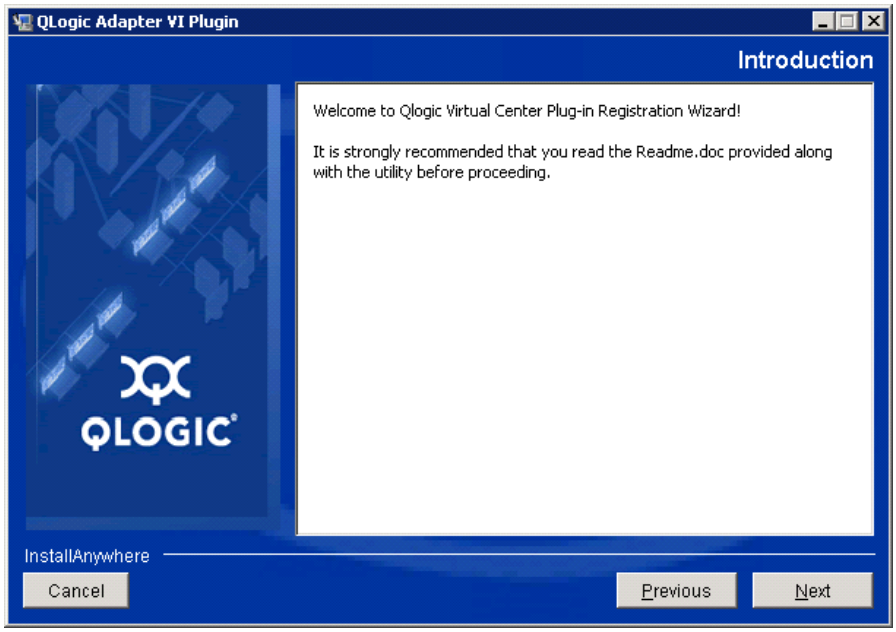

Figure 2-10. QLogic Adapter VI Plug-in Wizard: Introduction

Wait while the wizard configures the plug-in (see Figure 2-11).

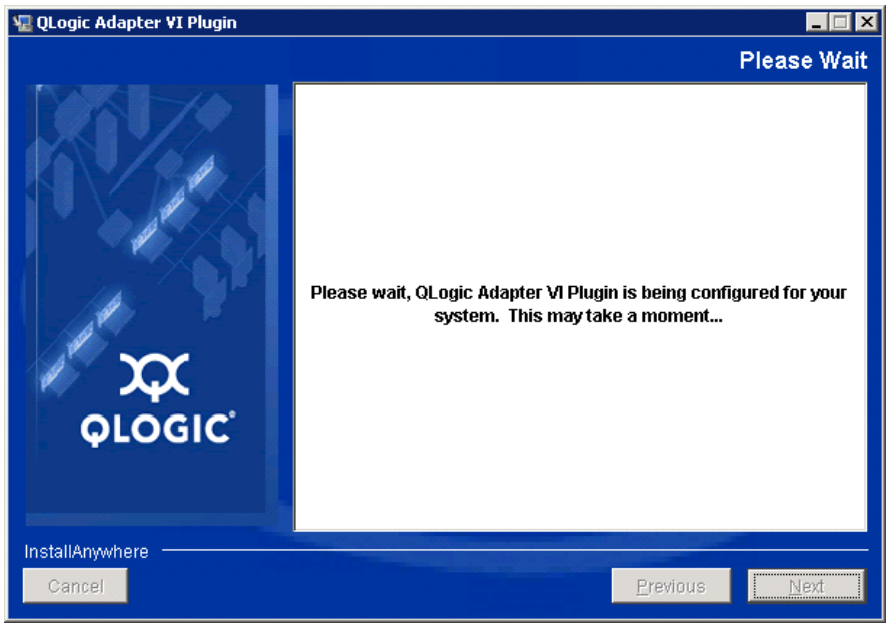

Figure 2-11. QLogic Adapter VI Plug-in Wizard: Configuring the Plug-in

4. Select the installation directory, and then click **Install** (see Figure 2-12).

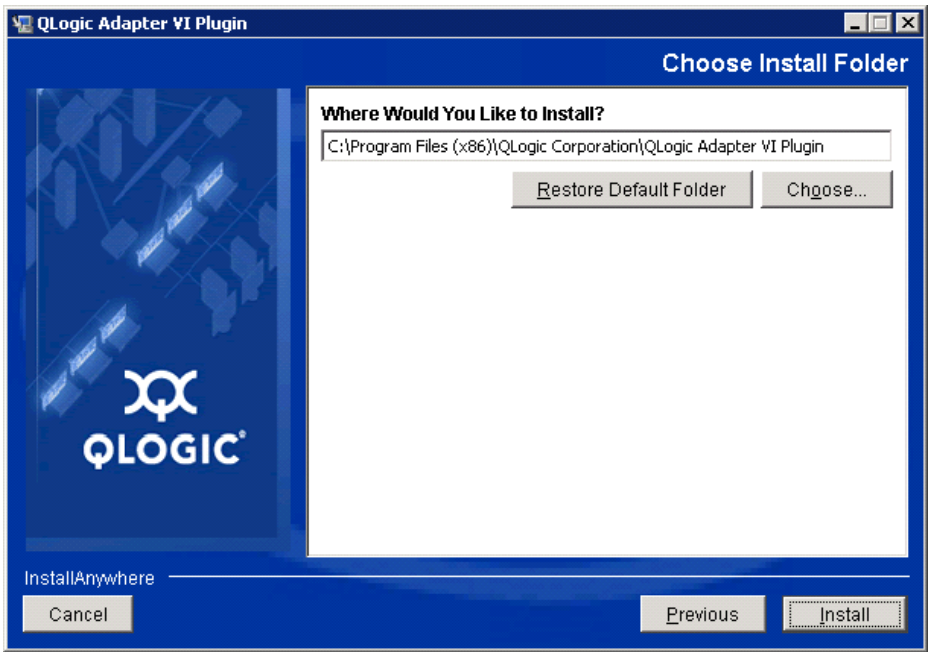

Figure 2-12. QLogic Adapter VI Plug-in Wizard: Choose Install Folder

Wait while the wizard performs the installation (see Figure 2-13).

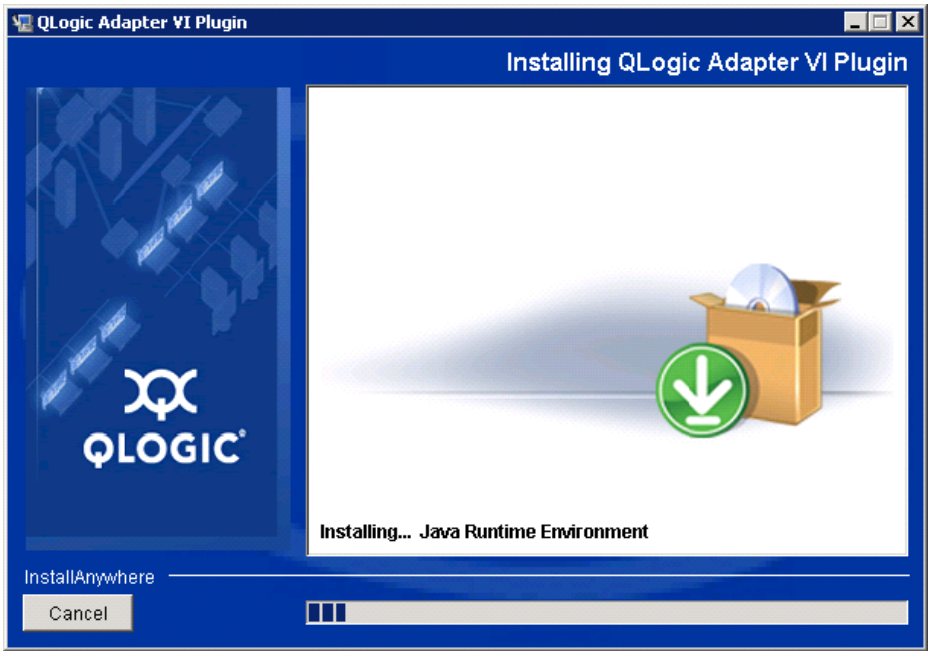

Figure 2-13. QLogic Adapter VI Plug-in Wizard: Installing the Plug-In

5. Type the requested information in the boxes, and then click **Next** (see Figure 2-14).

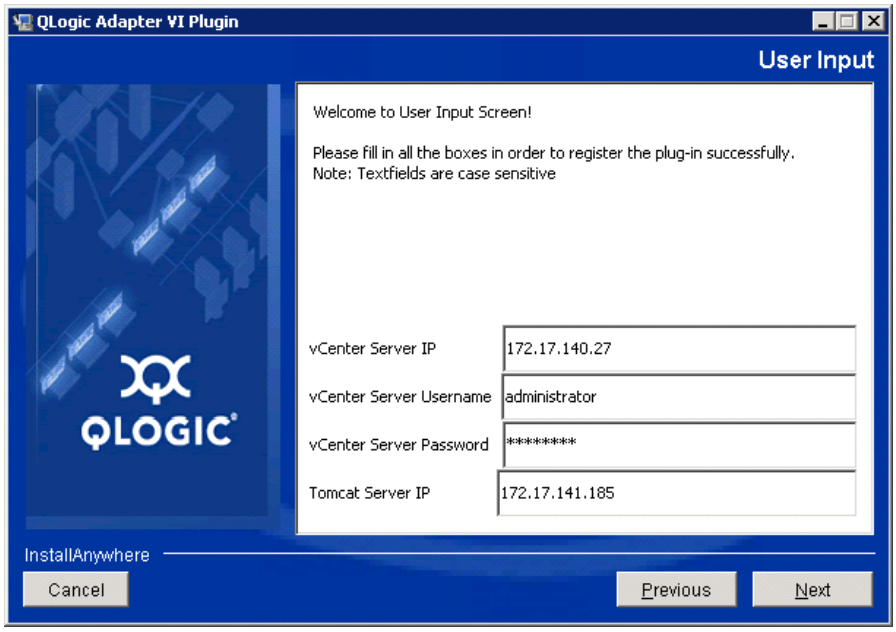

Figure 2-14. QLogic Adapter VI Plug-in Wizard: User Input

Wait while the wizard finishes configuring the plug-in (see Figure 2-15).

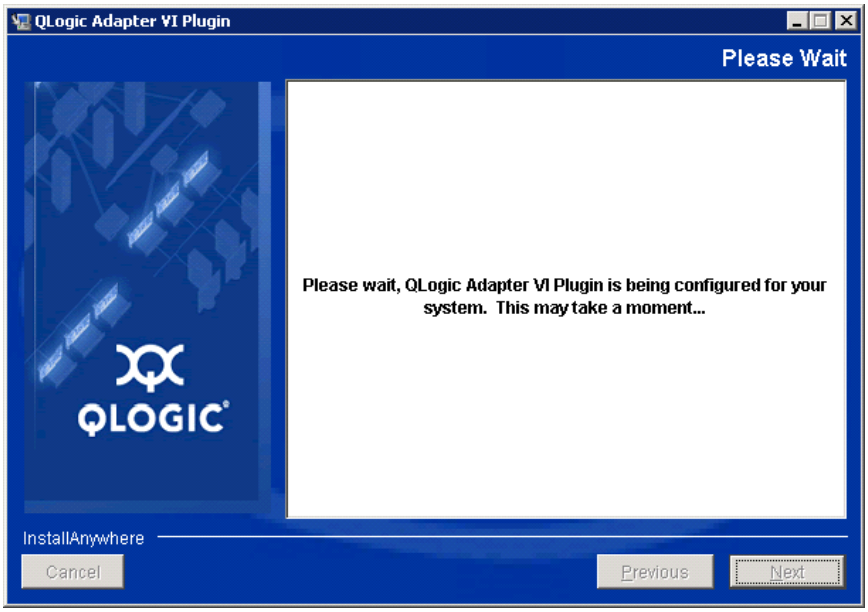

Figure 2-15. QLogic Adapter VI Plug-in Wizard: Configuration in Progress

6. When the Registration Result window (Figure 2-16) appears, click **Finish** to exit.

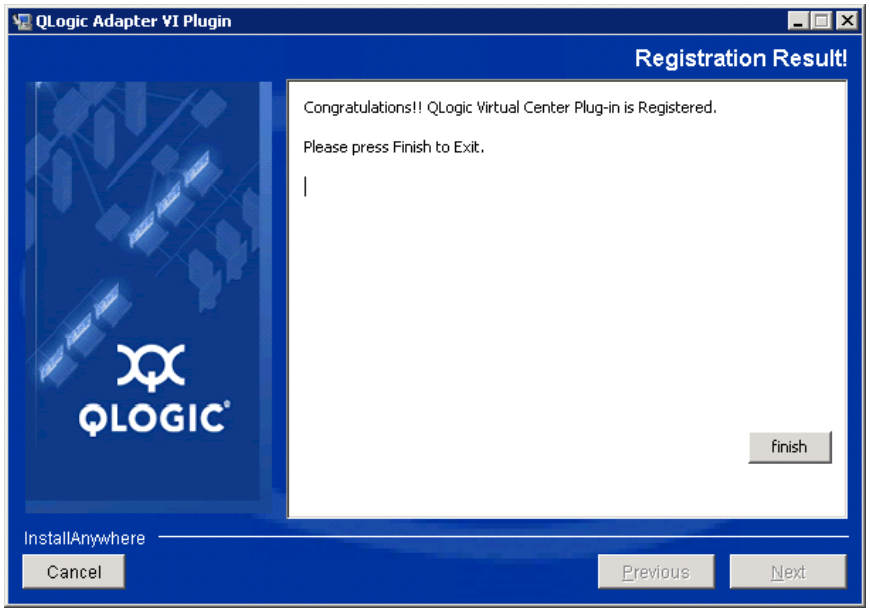

Figure 2-16. QLogic Adapter VI Plug-in Wizard: Registration Result

7. After the installation completes, restart the Tomcat™ service as follows:
   - ❑ If the plug-in is installed on the VMware vCenter Server, restart the VMware Virtual Center Management Web services.
   - ❑ If the plug-in is installed on a server other than the vCenter Server, restart the Apache Tomcat service.

Plug-in Unregistration from a Manual Install

If you have performed a manual install of the QConvergeConsole VMware vCenter Server Plug-in, you must perform a manual uninstall before running the plug-in Installation Wizard.

VMware provides two types of scripts for plug-in registration (and unregistration):

- For PowerShell scripting:

  http://communities.vmware.com/docs/DOC-4521

- For Perl®:

  http://communities.vmware.com/docs/DOC-4530

Before you can use the script, you must download the appropriate VI SDK from VMware:

- For Perl VI SDK, download vSphere SDK for Perl:

  http://www.vmware.com/support/developer/viperltoolkit/

- For PowerShell, download vSphere PowerCLI:

  http://communities.vmware.com/community/vmtn/vsphere/automationtools/powercli

After downloading and installing the SDK and the registration script, follow the VMware instructions to unregister the plug-in. For example, the Perl unregister command is:

```
perl registerPlugin.pl --server="127.0.0.1"
-username="administrator" --password="password"
--key="com.qlogic.QLogicAdapterVIPlugIn" --action="remove"
```

To log into the vCenter Server, replace `server's IP address`, `administrator` and `password` with the correct information.

Uninstalling the QConvergeConsole VMware vCenter Server Plug-in

To remove the QConvergeConsole VMware vCenter Server Plug-in:

1. In the Windows Control Panel, select **Programs and Features**.
2. In the Add or Remove Programs dialog box, select the QConvergeConsole VMware vCenter Server Plug-in, and then click **Change/Remove**.
3. To remove the plug-in, follow the instructions in the QConvergeConsole VMware vCenter Server Plug-in installer.

Installing the QLogic Adapter CIM Provider

This section describes how to install, start, and remove the QLogic Adapter CIM Provider for VMware ESXi. Because there are multiple zip packages, ensure that you pick the zip package that matches your environment: ESXi 7.0 and later.

NOTE

The QLogic Adapter CIM Provider for VMware ESXi was generated as a VIB file. A VIB contains the complete set of files and binaries required to install the provider on VMware ESXi. The file `offline-bundle.zip` contains the VIB and the necessary metadata to install the provider on VMware ESXi.

Initial Installation Methods

Initial installation methods for the CIM Provider include the following:

- **Online**. Refer to "Installing the CIM Provider on an ESXi 6.7 U3 or 7.0 Host" on page 26.
- **Offline**. Refer to "Existing ESXi Installation Using VMware Update Manager" on page 26.
- **Remote**. Refer to "Remote Installation of the CIM Provider on an ESXi Host" on page 27.

Installing the CIM Provider on an ESXi 6.7 U3 or 7.0 Host

1. For an ESXi 6.7/7.0 host, do the following; otherwise proceed to Step 2.
   a. Turn off the firewall on the ESXi 6.7/ESXi 7.0 host by issuing the following command:

      ```
      # esxcli network firewall set --enabled false
      ```

   b. Reboot the ESXi 6.7/7.0 host.
   c. On some platforms, FC adapters are not discovered from ESXi 6.7/7.0 through vCenter. If this happens, restart the CIMON by issuing the following commands:

      ```
      # /etc/init.d/sfcbd-watchdog stop
      # /etc/init.d/sfcbd-watchdog restart
      ```

2. Copy the following file to the root directory (`/`) of the ESXi system (for ESXi 7.0):

   `qlogic_adapter_provider_vmware_esx70-1.7.28`

   This CIM provider is applicable for ESXi versions 6.*x* and later.

3. Issue the `esxcli` commands as follows:

   ```
   # cd /
   # esxcli software acceptance set --level=CommunitySupported
   # esxcli software vib install -d
   file://<offline bundle>.zip --maintenance-mode
   ```

4. Reboot the system as required.

Existing ESXi Installation Using VMware Update Manager

An existing ESXi host has asynchronous drivers installed using VMware Update Manager (VUM). For more information, see "Using vSphere ESXi Image Builder CLI" in the *vSphere Installation and Setup Guide* at:

http://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-pubs.html

To install the asynchronous drivers:

1. Extract the contents of the asynchronous driver zip file.
2. Identify the `offline-bundle.zip` file or files.
3. From vCenter Server, go to **Home**, and then select **Update Manager**.
4. Click the **Patch Repository** tab.
5. Click the **Import Patches** link at the top right of the window.
6. Click **Finish**.

   The asynchronous driver is now added to the patch repository.
7. Create a baseline and remediate the ESXi host. For more information, refer to *Installing and Administering VMware vSphere Update Manager:*

   http://www.vmware.com/support/pubs/vum_pubs.html.

Remote Installation of the CIM Provider on an ESXi Host

NOTE

Before performing this procedure, ensure that the remote ESXi system is in Maintenance Mode. To do so using vSphere Client, select **Inventory**, select **Host**, and then select **Enter Maintenance Mode**.

1. Copy the `offline-bundle.zip` file to any location on the host where either the vSphere CLI package is installed or vSphere Management Assistant (vMA) is hosted.
2. Navigate to the location of the `offline-bundle.zip` file.
3. To install the offline bundle, issue the `vihostupdate` command as follows:

   ```
   # vihostupdate.pl <conn_options> --install --bundle
   offline-bundle.zip --nosigcheck
   ```
4. Follow the on-screen instructions to complete the installation. You might need to reboot the ESXi system.

NOTE

For more details on the `vihostupdate` command, see the *vSphere Command-Line Interface Installation and Reference Guide* at:

http://www.vmware.com/pdf/vsphere4/r40/vsp_40_vcli.pdf

Subsequent Update Installation

To update the QLogic Adapter CIM Provider after a prior VIB installation, follow the instructions in “Uninstalling the QLogic Adapter CIM Provider” on page 28 to remove the existing VIB. After completing the VIB removal, use the same steps in “Initial Installation Methods” on page 26 to install the new VIB.

Starting the QLogic Adapter CIM Provider

After a system startup, the SFCB (Small-Footprint CIM Broker) CIMOM (CIM object manager) in the ESXi system should start automatically and load the QLogic Adapter CIM Provider, when necessary.

If the CIM Provider does not start automatically, you can manually stop, start, or restart the SFCB CIMOM by issuing the following commands.

- To stop the SFCB CIMOM and the QLogic Adapter CIM Provider:

  ```
  # /etc/init.d/sfcbd-watchdog stop
  ```

- To start the SFCB CIMOM and the QLogic Adapter CIM Provider:

  ```
  # /etc/init.d/sfcbd-watchdog start
  ```

- To restart the SFCB CIMOM and the QLogic Adapter CIM Provider:

  ```
  # /etc/init.d/sfcbd-watchdog restart
  ```

After starting the SFCB CIMOM, use a CIM client utility to query the QLogic Adapter CIM Provider for information.

Uninstalling the QLogic Adapter CIM Provider

You can uninstall the QLogic Adapter CIM Provider for your version of VMware. For information about removing the QLogic Adapter CIM Provider through a remote host, see the *QLogic Adapter CIM Provider and vCenter Plug-in for VMware ESX/ESXi Readme* file.

To uninstall the CIM Provider from an ESXi 6.*x* or 7.0 host:

1. To view the VIB list, issue the following command:

   ```
   # esxcli software vib list
   ```

2. To remove the QLogic Adapter CIM Provider, issue the following command:

   ```
   # esxcli software vib remove --vibname viplugin-cimprovider
   --maintenance-mode -f
   ```

To uninstall the CIM Provider from a remote host:

NOTE

Before performing this procedure, make sure that the ESXi system is in Maintenance Mode. To do so using the vSphere Client, select **Inventory**, select **Host**, and then select **Enter Maintenance Mode**.

1. From a console on the host where the vSphere CLI package is installed or vMA is hosted, query and find the Bulletin ID of the existing provider:

   ```
   # vihostupdate.pl <conn_options> --query
   ```

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

   ```
   # vihostupdate.pl <conn_options> --remove --bulletin
   <bulletinID>
   ```

NOTE

For more details on `vihostupdate`, see the *vSphere Command-Line Interface Installation and Reference Guide*, located here:

http://www.vmware.com/pdf/vsphere4/r40/vsp_40_vcli.pdf

Configuring VM-ID

Virtual machine (VM) storage I/O activity automatically generates VM-ID tagging in FC frames. VM-ID tags I/O frames with an ID that identifies the virtual machine that sent or received them. This ID allows management applications to monitor traffic flows down to the virtual machine level. VM-ID is disabled by default and requires a Brocade switch running a supported FOS version. For detailed VM-ID information, see the documentation for your switch.

- To enable VM-ID (and disable QoS), issue the following command and then reboot the system:

  ```
  # esxcfg-module -s "ql2xvmidsupport=1,
  ql2xfabricpriorityqos=0" qlnativefc
  ```

- To disable VM-ID, issue the following command, and then reboot the system:

  ```
  # esxcfg-module -s "ql2xvmidsupport=0" qlnativefc
  ```

For details about configuring VM-ID, see the *User's Guide: Marvell Fibre Channel and Converged Network Adapters for VMware ESXi 6.x, 7.x, and 8.x* (SN0154529-00).

3 Fibre Channel Configuration

This chapter provides the following information about configuring Fibre Channel:

NOTE

For information on configuring the Fibre Channel Adapter using VMware vCenter Server, refer to the *User's Guide—QConvergeConsole Plug-ins for vSphere* (part number TD-000965), available at:

www.marvell.com.

Updating the Dell Firmware

The firmware Dell Update Package (DUP) is a Flash update utility only; it is not used for adapter configuration. To run the firmware DUP, double-click the executable file. Alternatively, run the firmware DUP from the command line, which supports several command line options.

NOTE

If you update the firmware using the DUP with the Integrated Dell Remote Access Controller (iDRAC) or Lifecycle Controller (LC), it shows FFV as the QLogic Fibre Channel adapter MBI version. This issue occurs only the first time that you run the firmware DUP; following a reboot, FFV will be shown correctly on subsequent firmware updates with the DUP.

Running the Firmware Update by Double-Clicking

For additional information, refer to the *Dell Update Packages Version xx.xx.xxx User's Guide*, "Using Dell Update Packages" section.

To run the firmware DUP by double-clicking the executable file:

1. Double-click the icon representing the firmware DUP file.

2. When the DUP splash screen appears (Figure 3-1), click **Install** to continue.

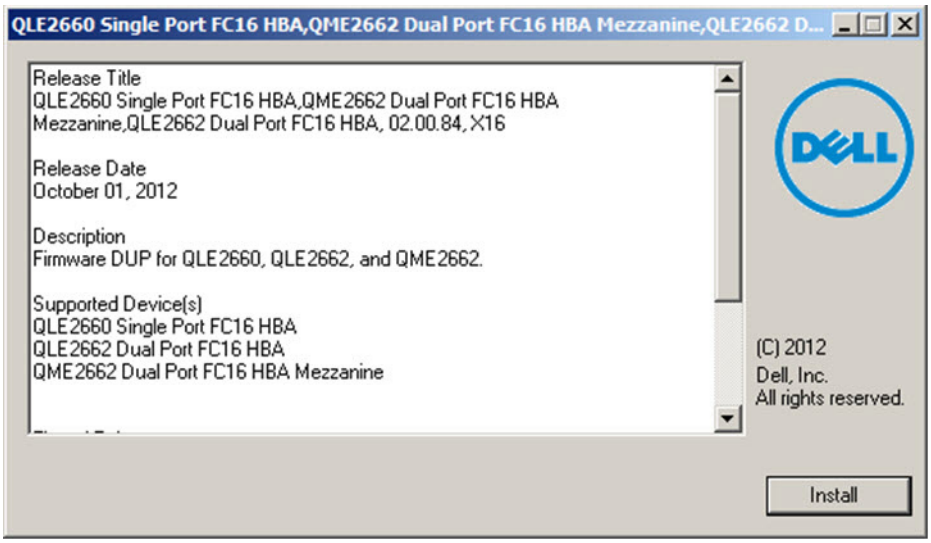

Figure 3-1. Dell Update Package Splash Screen

3. Follow the installer instructions. If a warning message appears (Figure 3-2), click **Yes** to continue the installation.

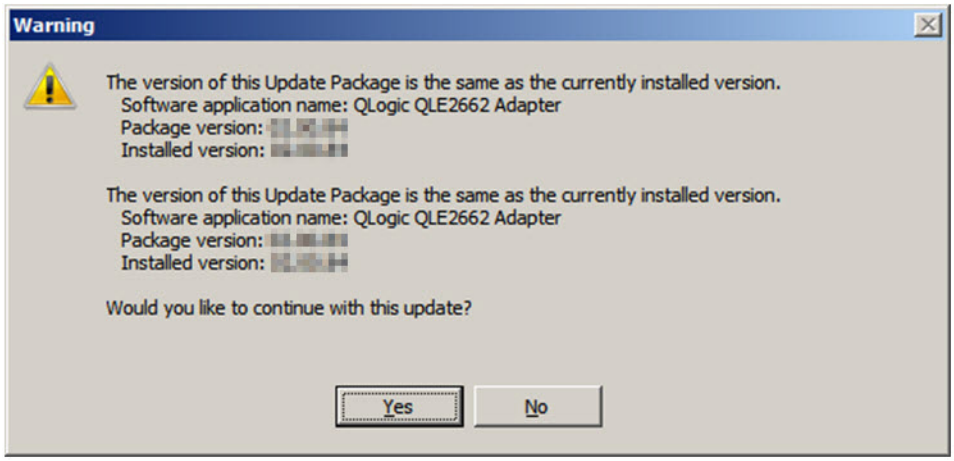

Figure 3-2. Continue Dell Update Package Installation

The installer indicates that it is loading the new firmware (Figure 3-3).

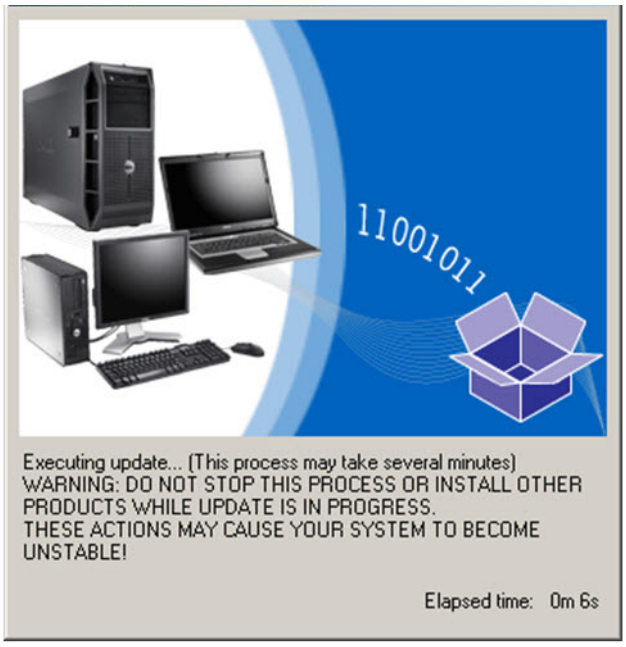

Figure 3-3. Loading New Firmware

4. When complete, the installer indicates the result of the installation (Figure 3-4). Click **Yes** to elect to reboot.

NOTE

You are prompted to reboot the system only if the adapter firmware includes components that require a power-on reset to be activated.

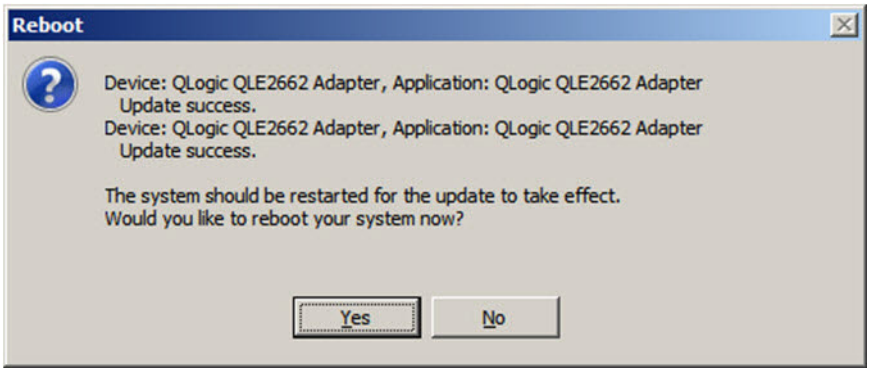

Figure 3-4. Result of Installation

5. To complete the installation, click **Finish** (Figure 3-5).

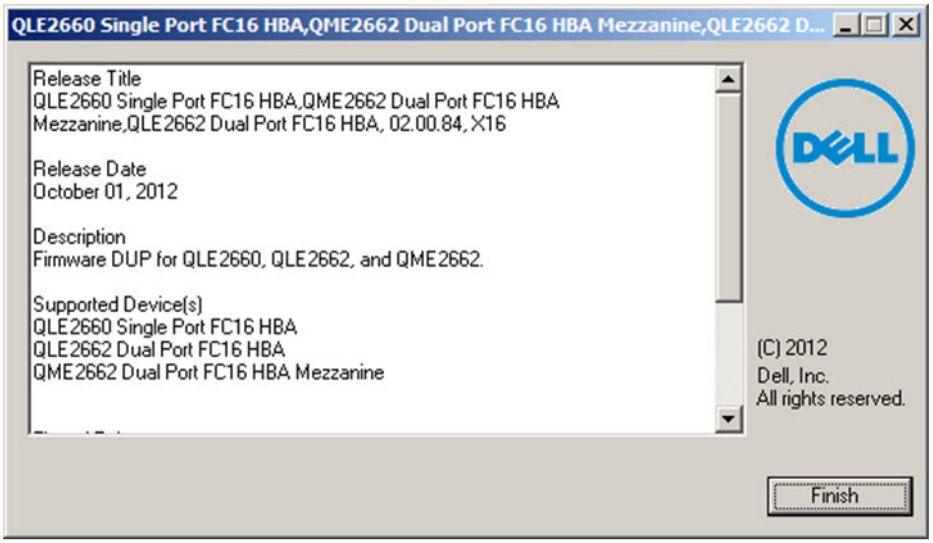

Figure 3-5. Finish Installation

Running the Firmware Update from the Command Line

Running the firmware DUP from the command line, with no options specified, results in the same behavior as double-clicking the icon representing the DUP. Note that the actual file name of the DUP will vary.

```
<adapter>_Firmware_<release>_<OS>_<version>.EXE
```

Figure 3-6 shows the options you can use to customize the DUP installation behavior.

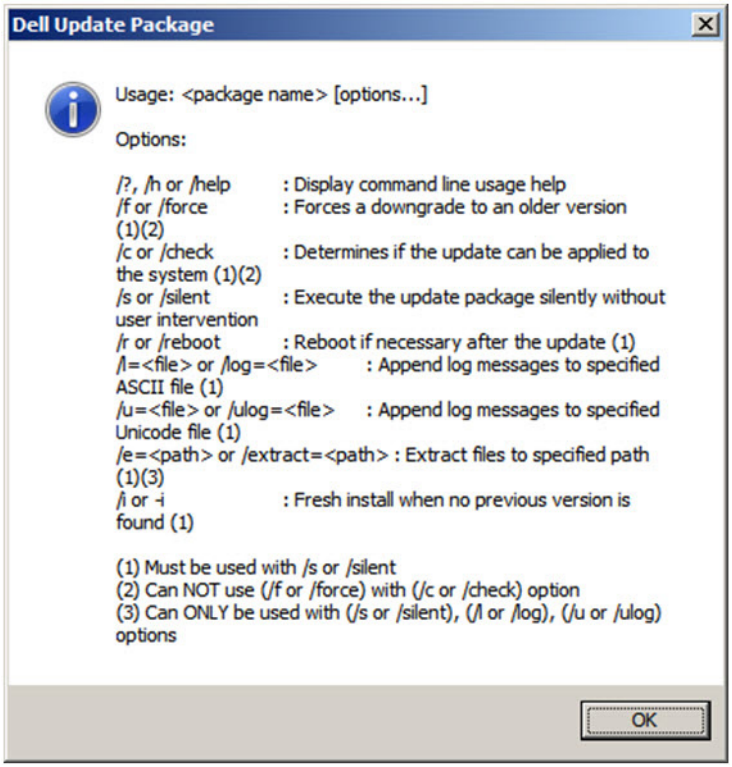

Figure 3-6. Command Line Options

Using *Fast!*UTIL for Custom Configuration

This section provides detailed configuration information for advanced users who want to customize the configuration of the adapters and connected devices using *Fast!*UTIL (the QLogic Fibre Channel Adapter BIOS utility) in a pre-OS environment.

To configure adapters using *Fast!*UTIL:

1. During the Host Bus Adapter BIOS initialization, press ALT+Q or CTRL+Q. It may take a few seconds for the *Fast!*UTIL menu to appear.
2. When the Select Host Adapter window appears, select the adapter type.
3. In the Adapter Settings window (see “Adapter Settings” on page 37), configure the adapter settings.

4. A message indicates that the configuration settings have been modified. Select **Save changes**.
5. To load the new parameters, reboot the system.

CAUTION

If the configuration settings are incorrect, your 2600 Series Adapter may not function properly.

The ***Fast!***UTIL **Options** menu provides the following options:

- Configuration Settings
- Scan Fibre Devices
- Fibre Disk Utility
- Loopback Data Test
- Select Adapter
- Exit *Fast!UTIL*

Configuration Settings

The following sections show and describe the 2600 Series Adapters **Configuration Settings** menu options in *Fast!*UTIL:

- Adapter Settings
- Selectable Boot Settings
- Restore Default Settings
- Raw NVRAM Data
- Advanced Adapter Settings

Adapter Settings

Table 3-1 describes the adapter settings.

Table 3-1. Adapter Settings

Setting	Values	Default	Description
Host Bus Adapter BIOS	Enabled, Disabled	Disabled	Disabling this setting disables the ROM BIOS on the 2600 Series Adapter, which frees space in upper memory. If you are booting from a Fibre Channel disk drive attached to the 2600 Series Adapter, you must enable this setting. See “Selectable Boot Settings” on page 38 for details on specifying the boot device.
Frame Size	512, 1024, 2048	2048	This setting specifies the maximum frame length supported by the 2600 Series Adapter.
Loop Reset Delay	0–60 seconds	5 seconds	After resetting the loop, the firmware refrains from initiating any loop activity for the number of seconds specified in this setting.
Adapter Hard Loop ID	Enabled, Disabled	Disabled	If the Adapter Hard Loop ID setting is enabled, the adapter attempts to use the ID specified in the Hard Loop ID setting.
Hard Loop ID	0–125	0	If the Adapter Hard Loop ID setting is enabled, the adapter attempts to use the ID specified in this setting.
Connection Options	0, 1, 2	2	This setting defines the type of connection: ■ 0—loop only (for the 4Gb and 8Gb data rate) ■ 1—point-to-point only (for the 4Gb, 8Gb, and 16Gb data rate) ■ 2—loop preferred, then point-to-point (for the 4Gb, 8Gb, and 16Gb data rate)
Fibre Channel Tape Support	Enabled, Disabled	Enabled	This setting enables FCP-2 recovery.
Data Rate	2, 3, 4, 5	2	This setting determines the data rate: ■ 2—The adapter autonegotiates and determines the data rate. ■ 3—The adapter runs at 4GFC[a]. ■ 4—The adapter runs at 8GFC[b]. ■ 5—The adapter runs at 16GFC. Loop mode is not supported.

[a] QME2662-DEL and QME2692-DEL do not support the 4GFC data rate.

[b] QME2692-DEL does not support the 8GFC data rate.

Selectable Boot Settings

The **Configuration Settings** menu provides access to the **Selectable Boot Settings** options. If you enable the Host Bus Adapter BIOS in the adapter settings, you can select the boot device, as shown in Table 3-2.

Table 3-2. Selectable Boot Settings

Enable Selectable Boot	Boot WWPN/ LUN	Device Boot
No	—	The BIOS configures the first disk drive it finds as the boot device.
Yes	None specified	The BIOS configures the first disk drive it finds that is also LUN 0 as the boot device.
Yes	Specified	The BIOS scans through the specified boot worldwide port name (WWPN) and LUN list until it finds a disk drive. This drive is configured as the boot device.

The **Selectable Boot Settings** option allows you to specify up to four WWPNs or LUNs as boot devices.

To enable the Selectable Boot option and insert devices into boot device locations:

1. On the **Configuration Settings** menu, select **Host Adapter Settings**.
2. Enable the **Host Bus Adapter BIOS** and then press the ESC key to return to the **Configuration Settings** menu.
3. On the **Configuration Settings** menu, select **Selectable Boot Settings**.
4. To enable or disable (toggle) the **Selectable Boot** option, press the ENTER key. To move to the Primary location of the **Selectable Boot List** menu, press the DOWN ARROW key.
5. To see a list of accessible devices (in the Select Fibre Channel Device menu), press ENTER.
6. Select a drive by pressing the UP ARROW and DOWN ARROW keys or the PAGE UP and PAGE DOWN keys to scroll to the device that you want to put into your **Selectable Boot** menu list. Then press ENTER to select the requested device and load it into the **Selectable Boot** menu list.
7. To specify an alternate boot device on the 2600 Series Adapter, use the arrow keys to move to the next available alternate entry, press ENTER, select the disk drive you want, and press ENTER again. You can specify up to three alternate boot devices in this way.

NOTE

System BIOS compatibility: To boot from a QLogic host 2600 Series Adapter in a computer system with a multiboot system BIOS, the 2600 Series Adapter must be included in the list of bootable devices in the system's **Boot** menu.

Restore Default Settings

The **Restore Defaults** option restores the NVRAM settings back to default.

Raw NVRAM Data

The **Raw NVRAM Data** option displays the 2600 Series Adapter's NVRAM contents in hexadecimal format. This option is a QLogic troubleshooting tool; you cannot modify the raw NVRAM data.

Advanced Adapter Settings

On the **Configuration Settings** menu in *Fast!*UTIL, select **Advanced Adapter Settings**. Table 3-3 describes the default settings for the 2600 Series Adapters.

Table 3-3. Advanced Adapter Settings

Setting	Values	Default	Description
LUNs per Target	0, 8, 16, 32, 64, 128, 256	128	This setting specifies the number of LUNs supported per target if the target does not support the Report LUN command. Multiple LUN support is typically for redundant array of independent disks (RAID) boxes that use LUNs to map drives.
Enable LIP Reset	Yes, No	No	This setting determines the type of loop initialization process (LIP) reset used when the OS initiates a bus reset routine as follows: ■ Yes = The driver initiates a global LIP reset to reset the target devices. ■ No = The driver initiates a global LIP reset with full login.
Enable LIP Full Login	Yes, No	Yes	When this setting is Yes, the 2600 Series Adapter's reduced instruction set computer (RISC) controller logs in to all ports after any LIP.
Enable Target Reset	Yes, No	Yes	When this setting is Yes, the drivers issue a Target Reset command to all devices on the loop when a SCSI Bus Reset command is issued.
Login Retry Count	0–255	8	This setting specifies the number of times the software tries to log in to a device.

Table 3-3. Advanced Adapter Settings (Continued)

Setting	Values	Default	Description
Port Down Retry Count	0–255 seconds	30 seconds	This setting specifies the number of seconds the software waits before resending a command to a port whose status indicates that the port is down.
Link Down Timeout	0–255 seconds	30 seconds	This setting specifies the number of seconds the software waits for a link that is down to come up.
Operation Mode	0, 5, 6	6	This setting specifies the reduced interrupt operation (RIO) modes if they are supported by the software driver. The RIO modes allow posting multiple command completions in a single interrupt. This setting supports the following modes: ■ 0—Interrupt for every I/O completion ■ 5—Interrupt when the interrupt delay timer expires ■ 6—Interrupt when the interrupt delay timer expires or no active I/Os
Interrupt Delay Timer	0–255 (100µs units)	2 (200µs)	This setting contains the value used by a timer to set the wait time before generating an interrupt.
Enable Interrupt	Yes, No	No	This setting enables or disables the interrupt: ■ Yes—This setting enables the BIOS to use the interrupt request (IRQ) assigned to the 2600 Series Adapter. ■ No—The BIOS polls for the 2600 Series Adapter's RISC controller mailbox command completion status.

Scan Fibre Devices

The **Scan Fibre Devices** option scans the Fibre Channel loop and lists all the connected devices by loop ID. It lists information about each device; for example, vendor name, product name, and revision. This information helps when configuring your 2600 Series Adapter and attached devices.

Fibre Disk Utility

The **Fibre Disk Utility** option scans the Fibre Channel loop and lists all the connected devices by loop ID. You can select a Fibre Channel hard disk and do one of the following tasks:

- Perform a low-level format
- Verify the disk media
- Verify the disk data (Some targets do not support this feature.)
- Select a disk device

CAUTION

Performing a low-level format destroys all data on the disk.

Loopback Data Test

The **Loopback Data Test** option allows you to perform the following loopback tests:

- External loopback data test
- Internal loopback data test

NOTE

When running an external loopback data test, make sure that the Fibre Channel link is up or a loopback plug is attached to the 2600 Series Adapter before starting the test.

Select Adapter

If you have a multiport or multiple 2600 Series Adapters in the system, use the **Select Adapter** option to select and then configure or view the settings of a specific 2600 Series Adapter or one of its ports.

Exit *Fast!*UTIL

The **Exit *Fast!*UTIL** option allows you to exit the utility and reboot the system. After making changes to the 2600 Series Adapter in *Fast!*UTIL, make sure you save the changes before you reboot the system.

Setting Fibre Channel Adapter Parameters

Use QConvergeConsole to set Fibre Channel adapter parameters:

- Setting Fibre Channel Adapter Parameters with Interactive QConvergeConsole CLI
- Setting Fibre Channel Adapter Parameters with Noninteractive QConvergeConsole CLI

Setting Fibre Channel Adapter Parameters with Interactive QConvergeConsole CLI

To set Fibre Channel adapter parameters with interactive QConvergeConsole CLI:

1. On the **FC Configuration** menu, select **HBA Parameters**.

2. On the port menu, select a port to open the **HBA Parameters** menu.
3. On the **HBA Parameters** menu, select the **Change Settings** option to open the **Configure Parameters** menu.

For more information, refer to the section about the Fibre Channel interactive commands in the *User's Guide—QConvergeConsole CLI* (part number SN0054667-00).

Setting Fibre Channel Adapter Parameters with Noninteractive QConvergeConsole CLI

For information about how to set Fibre Channel adapter parameters using the noninteractive mode of QConvergeConsole CLI, refer to the section about the Fibre Channel noninteractive commands in the *User's Guide—QConvergeConsole CLI* (part number SN0054667-00).

Configuring Target Persistent Binding

NOTE

Target persistent binding is supported only on Windows.

Use QConvergeConsole to configure target persistent binding:

- Configuring Persistent Binding with Interactive QConvergeConsole CLI
- Configuring Persistent Binding with Noninteractive QConvergeConsole CLI

Configuring Persistent Binding with Interactive QConvergeConsole CLI

To bind a target using interactive QConvergeConsole CLI:

1. On the **FC Configuration** menu, select **Target Persistent Bindings**.
2. On the port menu, select a port and then select **Bind Target(s)** to bind a target to a port.
3. Select a target and specify a target ID to open the **Target Persistent Binding – FC Port Configuration** menu with options to select more targets, save changes, or cancel the binding operation.

For more information, refer to the section about Fibre Channel interactive commands in the *User's Guide—QConvergeConsole CLI* (part number SN0054667-00).

Configuring Persistent Binding with Noninteractive QConvergeConsole CLI

To bind a selected target to a specific adapter using the noninteractive mode of QConvergeConsole CLI, issue the following command:

```
# qaucli -pr fc -p (<hba instance> | <hba wwpn>) (<target wwnn>
<target wwpn> <target port id> <target id>)
```

Where:

`hba instance` = Adapter number (use the `-g` command to find)
`hba wwnn` = Worldwide node name of the adapter
`hba wwpn` = Worldwide port name of the adapter
`target wwnn` = Worldwide node name of the target
`target wwpn` = Worldwide port name of the target
`target port id` = Port ID of the target
`target id` = ID to which the target is bound

For more information, refer to the section about Fibre Channel noninteractive commands in the *User's Guide—QConvergeConsole CLI* (part number SN0054667-00).

Configuring Boot Devices

To configure boot devices, follow the instructions in one of the following sections:

- Configuring Boot Devices with Interactive QConvergeConsole CLI
- Configuring Boot Devices with Noninteractive QConvergeConsole CLI
- Configuring Boot Devices with the BIOS
- Configuring Boot Devices with QConvergeConsole Plug-in for VMware vCenter Server
- Configuring Boot Devices with QConvergeConsole Plug-in for VMware vSphere HTML5 Client

Configuring Boot Devices with Interactive QConvergeConsole CLI

In the interactive mode of QConvergeConsole CLI, open the **Adapter Configuration** menu, and then select **Configure Boot Devices**. Depending on the selected adapter, you can configure the boot device with BIOS or UEFI, or only UEFI.

For information, refer to the section about the Fibre Channel interactive commands in the *User's Guide—QConvergeConsole CLI* (part number SN0054667-00).

Configuring Boot Devices with Noninteractive QConvergeConsole CLI

In the noninteractive mode of QConvergeConsole CLI, set a specific target as the boot device (BIOS boot mode) for a specific adapter by issuing the following command:

```
# qaucli -pr fc -e (<hba instance> | <hba wwpn>) <target wwnn>
<target wwpn> <target id> <lun id>
```

Where:

`hba instance`	=	Adapter number (issue the `-g` command to find)
`hba wwpn`	=	Worldwide port name of the adapter
`target wwnn`	=	Worldwide node name of the target
`target wwpn`	=	Worldwide port name of the target
`target id`	=	ID to which the target is bound
`lun id`	=	ID of the LUN

In the noninteractive mode of QConvergeConsole CLI, set a specific target as the boot device (UEFI boot mode) for a specific adapter by issuing the following command:

```
qaucli -pr fc -e <hba instance>|<hba wwpn>  <storage WWNN> <storage
WWPN> <storage nsid> <storage flag> <boot entry> <boot mode>
[<storage NQN>] [<storage ctrlid>]
```

For more information, refer to the "Fibre Channel Noninteractive Commands" section of the *User's Guide—QConvergeConsole CLI* (part number SN0054667-00).

Configuring Boot Devices with the BIOS

For information about boot device configuration with the *Fast!*UTIL BIOS utility, refer to "Using *Fast!UTIL for Custom Configuration" on page 35* for information.

Configuring Boot Devices with QConvergeConsole Plug-in for VMware vCenter Server

Follow these steps to configure boot devices with QConvergeConsole VMware vCenter Server Plug-in.

To configure boot devices:

1. In the QConvergeConsole left pane, expand the tree and then select a port.
2. In the content pane on the right, click the **Boot** tab.
3. Enable boot from the port by setting the appropriate settings:
   a. Select the appropriate check box options.
   b. For each applicable **Boot Name**, enter the information in the **Target-WWN** and **LUN ID** boxes.
4. Click **Save Configuration**.

   Figure 3-7 shows an example.

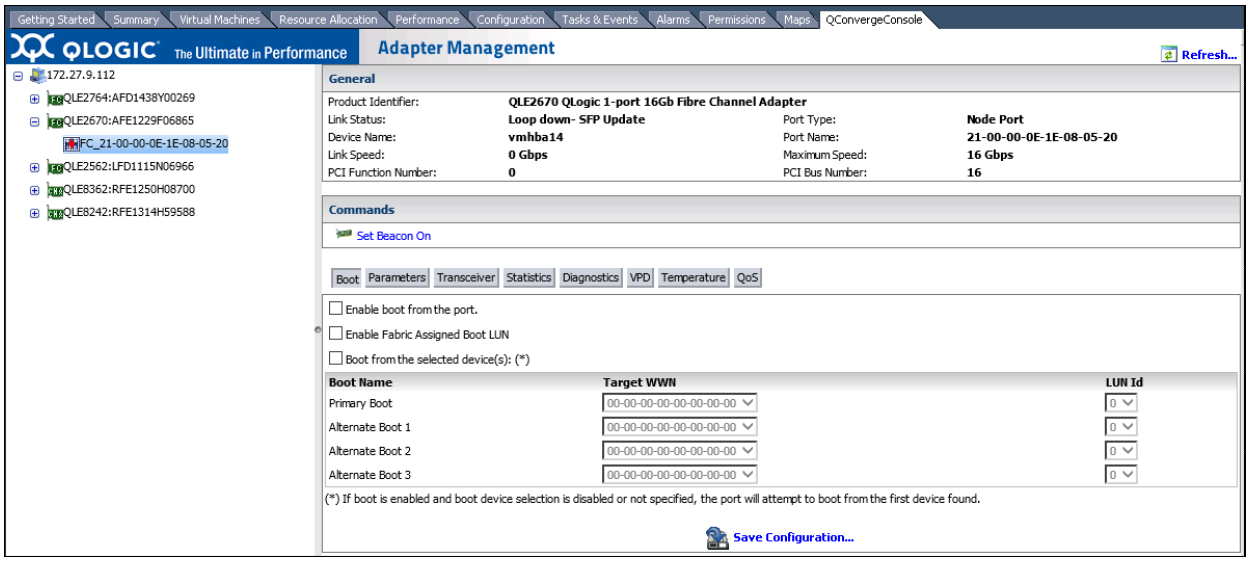

Figure 3-7. Configuring Boot Devices with QConvergeConsole Plug-in for VMware vCenter Server

Configuring Boot Devices with QConvergeConsole Plug-in for VMware vSphere HTML5 Client

NOTE

This procedure is not supported in VMware ESXi 8.0.

To configure boot devices with QConvergeConsole Plug-in for VMware vSphere HTML Client, follow the instructions in the “Boot Configuration—Fibre Channel Port” section of the *User’s Guide—QConvergeConsole Plug-ins for vSphere* (part number TD-000965).

Configuring Virtual Ports (NPIV)

Use QConvergeConsole to configure virtual ports for N_Port ID virtualization and quality of service:

- Configuring NPIV with Interactive QConvergeConsole CLI
- Configuring NPIV with Noninteractive QConvergeConsole CLI
- Configuring NPIV Quality of Service

Configuring NPIV with Interactive QConvergeConsole CLI

For information about using virtualization (NPIV) in the interactive mode of QConvergeConsole CLI, refer to the *User’s Guide—QConvergeConsole CLI* (part number SN0054667-00). To download the guide, go to www.marvell.com.

Configuring NPIV with Noninteractive QConvergeConsole CLI

You can configure a virtual port in the noninteractive mode of QConvergeConsole CLI with either an automatic or a specific WWPN.

To create a virtual port with an automatic WWPN:

- Issue the following command:

```
# qaucli -vp <hba instance>|<hba wwpn> create auto
[<num_of_vport>]
```

Where:

`hba instance`	=	Adapter number (issue the `-g` command to find)
`hba wwpn`	=	Worldwide port name of the adapter
`num_of_vport`	=	Number of virtual ports to be created.

If the key word 'max' is specified, then maximum number of virtual ports will be automatically created.

To create a virtual port with a specific WWPN:

1. Issue the following command:

   `# qaucli -pr fc -vp (<hba instance> | <hba wwpn>) create <vport hex>`

   Where:

   `hba instance` = Adapter number (issue the `-g` command to find)

   `hba wwpn` = Worldwide port name of the adapter

   `vport hex` = World wide port name of the virtual port with the two hexadecimal digits in byte three supplied by the user

2. When prompted, type two hexadecimal digits. The system checks these digits to ensure that they are unique and, if they are, puts them into byte 1 of the WWPN.

For more information, refer to the section about Fibre Channel noninteractive commands in the *User's Guide—QConvergeConsole CLI* (part number SN0054667-00).

Configuring NPIV Quality of Service

The QLogic 2600 Series Adapters solution provides for standards-based quality of service (QoS), ensuring high-quality performance for applications that require preferential delivery. The QLogic QoS solution is based on assigning QoS levels to virtual ports (NPIV ports), in contrast to end-to-end QoS, which assigns QoS levels from initiator to target across the fabric. You can configure the QoS using the **priority method** or the **bandwidth method**, as shown in Figure 3-8.

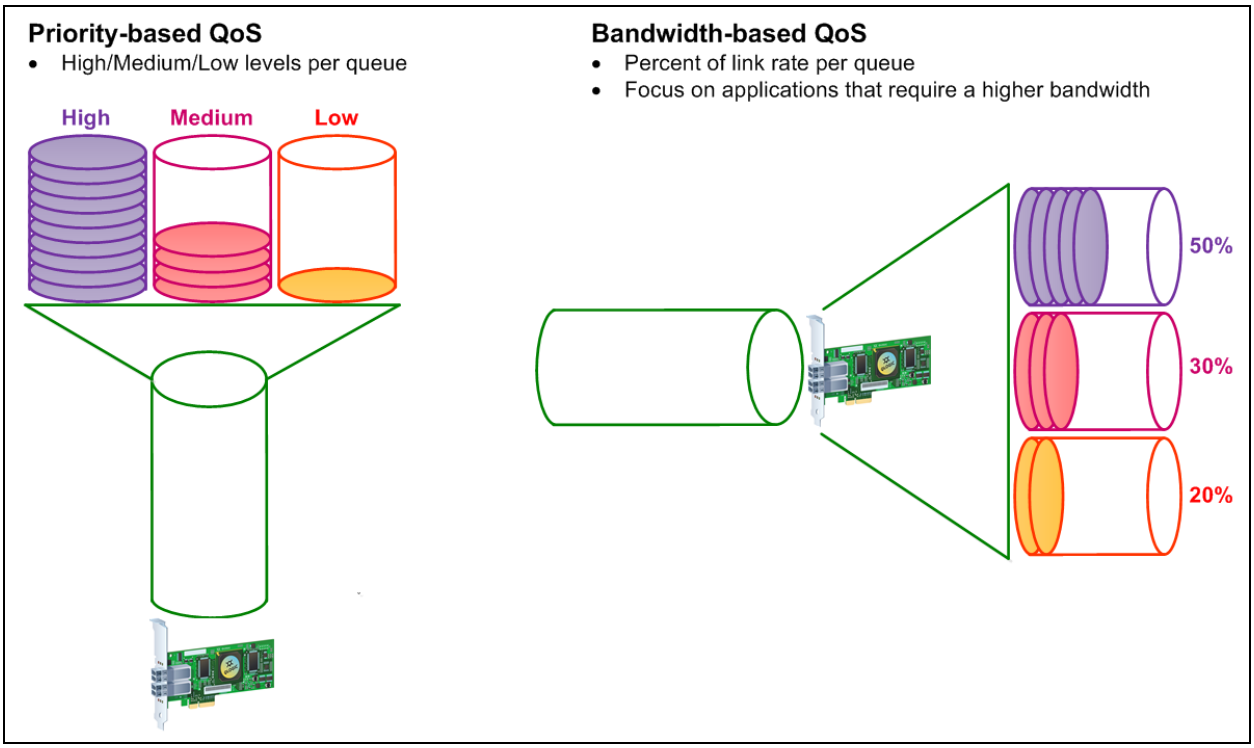

Figure 3-8. Priority and Bandwidth-Based QoS

On a single physical port, you can configure QoS using the priority method or the bandwidth method, but not both.

- The **priority method** assigns priority levels (low, medium, or high) to virtual ports based on the quality of service required for the applications or VMs using the virtual port. This method provides for various levels of latency requirements across virtual ports.

- The **bandwidth method** assigns a bandwidth percentage (0–100) or a bandwidth rate (for example, 6GFC) to applications or VMs that are using a virtual port. This method provides a minimum bandwidth guarantee. The bandwidth-based QoS configuration applies only when the total bandwidth requirement of the applications exceeds the available bandwidth. The bandwidth is not reserved for a specific virtual port, and unused bandwidth is shared among the other virtual ports.

The physical port or WWN always has high priority, independent of the bandwidth or priority scheme chosen for created virtual ports. You need not assign a QoS level to every virtual port that you create. If a QoS level is not assigned to a virtual port, the default is low priority (best effort), independent of the bandwidth or priority scheme. This flexibility allows you to focus on critical applications to ensure that the highest service level is provided. You can modify QoS levels for virtual ports using QConvergeConsole CLI.

To modify QoS levels, the configuration must meet the following hardware and software requirements:

- Server operating system: Microsoft® Windows Server 2019 or later
- Server support for MSI-X
- Marvell QLogic 2600 Series Adapter
- Latest version of the QLogic STOR miniport driver, available from Dell:

  http://support.dell.com
- Fibre Channel switch that supports NPIV
- Physical port attached in a point-to-point connection

Setting QoS by Priority

To set the QoS by priority with interactive QConvergeConsole CLI:

1. From the **Adapter Configuration** menu, select the **NPIV QoS** option.
2. From the port menu, select a port to open the **QoS Type Menu**. Select **1: QoS Priority** to set the priority.

Setting QoS by Bandwidth

Setting the QoS by bandwidth allocates up to 80 percent of the physical port's bandwidth to its virtual ports. The remaining 20 percent is reserved to handle non-QoS applications and virtual ports.

Setting QoS by bandwidth guarantees bandwidth for each virtual port and to the application or VM associated with that virtual port. The proper QoS setting can prevent bottlenecks that occur when VMs contend for port bandwidth.

Consider how much bandwidth is required to maximize application or VM performance and set the QoS level to 105 percent of that value. For example, if a mission-critical application on a VM requires 1Gb of bandwidth to transmit its data over the fabric, Marvell recommends a 1.05Gb bandwidth for the virtual port.

Alternatively, you can set the virtual port QoS values as a percentage of the total available bandwidth.

NOTE

Setting QoS by bandwidth guarantees a minimum bandwidth to a virtual port, but not a maximum limit. If the application or virtual port requires more bandwidth than the QoS bandwidth value allows and the physical port has available bandwidth, the application receives the additional bandwidth. The port allocates additional bandwidth on demand up to the port capacity or until there is contention for bandwidth among the virtual ports.

When the physical port is partitioned into four virtual ports, the port bandwidth is divided between the virtual ports according to traffic demands.

QConvergeConsole lets you configure QoS for each virtual port by setting minimum and maximum percentages of the physical port's bandwidth for each virtual port. This feature lets you guarantee a transmission rate for each port that requires a specific bandwidth to run mission-critical applications for business continuity using virtual ports. The setting for a specific QoS can resolve bottlenecks that exist when VMs contend for port bandwidth.

To set the QoS by bandwidth speed with interactive QConvergeConsole CLI:

1. On the **Adapter Configuration** menu, select the **NPIV QoS** option.
2. From the port menu, select a port to open the **NPIV QoS** menu. Select **2: QoS Bandwidth** to set the bandwidth speed.

Configuring Fibre Channel Driver Parameters

Use QConvergeConsole GUI, CLI, or VMware plug-in to configure Fibre Channel driver parameters:

- Configuring Fibre Channel Driver Parameters with Interactive QConvergeConsole CLI
- Configuring Fibre Channel Driver Parameters with Noninteractive QConvergeConsole CLI
- Configuring Fibre Channel Driver Parameters with QConvergeConsole VMware vCenter Server Plug-in

Configuring Fibre Channel Driver Parameters with Interactive QConvergeConsole CLI

To configure Fibre Channel driver parameters using interactive QConvergeConsole CLI:

1. On the **Adapter Configuration** menu, select **HBA Parameters**.
2. On the port menu, select a port to open the **HBA Parameters** menu.
3. Select options on the **Change Settings** menu to present persistently bound targets plus new targets with driver default, present persistently bound targets only, bind devices by WWPN, and bind devices by port ID.

For more information, refer to the section about the Fibre Channel interactive commands in the *User's Guide—QConvergeConsole CLI* (part number SN0054667-00).

Configuring Fibre Channel Driver Parameters with Noninteractive QConvergeConsole CLI

To configure driver settings, issue the following command:

```
# qaucli -pr fc -fs (<hba instance> | <hba wwpn>)
{(<param name> | <param alias>) <param value>}
```

Where:

`hba instance` = Adapter number (use the `-g` command to find)

`hba wwpn` = Worldwide port name of the adapter

`param name` = Name of the parameter

`param alias` = Alias of the parameter, either:

- `PO` Persistently bound targets only
- `PN` Present persistently bound targets plus any new targets with driver default
- `NO` Present targets with driver default
- `BW` Bind devices by WWPN
- `BP` Bind devices by port ID

`param value` = New value of the parameter

For more information, refer to the section about the Fibre Channel noninteractive commands in the *User's Guide—QConvergeConsole CLI* (part number SN0054667-00).

Configuring Fibre Channel Driver Parameters with QConvergeConsole VMware vCenter Server Plug-in

For details on how to configure the Fibre Channel driver parameters using the QConvergeConsole plug-in, see the *User's Guide—QConvergeConsole Plug-ins for vSphere* (part number TD-000965).

Configuring Selective LUNs

To configure selective LUNs with the interactive QConvergeConsole CLI, refer to the section about the Fibre Channel interactive commands in the *User's Guide—QConvergeConsole CLI* (part number SN0054667-00).

To configure selective LUNs with the noninteractive QConvergeConsole CLI, refer to the section about the Fibre Channel noninteractive commands in the *User's Guide—QConvergeConsole CLI* (part number SN0054667-00).

Configuring the UEFI Driver

This section provides information on configuring the unified extensible firmware interface (UEFI) driver contained in the multiboot image. UEFI 2.*x* systems use a human interface infrastructure (HII) to configure motherboard devices and plug-in adapters. QLogic Fibre Channel Adapters use an HII to configure adapter parameters and FCP and FC-NVMe boot-from-SAN settings.

To configure the UEFI driver for Dell:

1. During system boot, press the key corresponding to your platform.
2. On the Dell System Setup window, select **Device Settings**, and then press the ENTER key.

The Device Settings window opens and lists the devices installed in the system. Each listed device supports the HII. QLogic devices have one device settings entry per port. Each entry contains the adapter name and WWPN of the port. Figure 3-9 shows an example of the System Setup window for Device Settings.

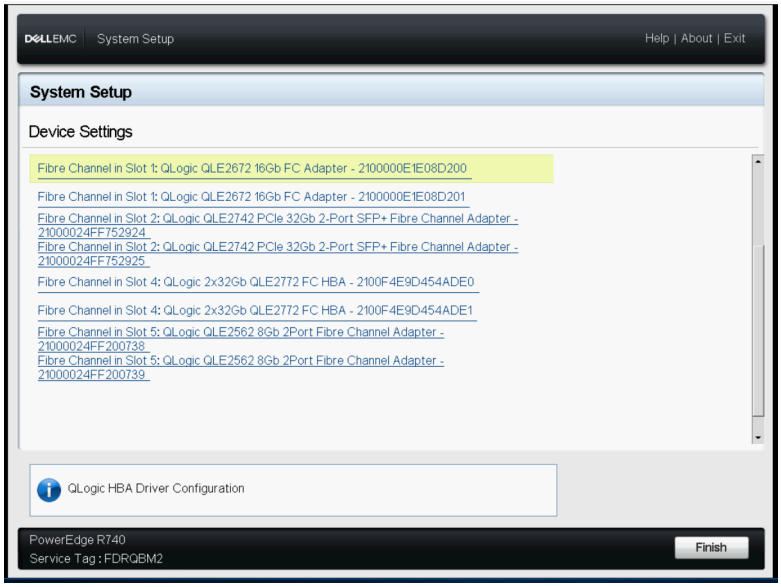

Figure 3-9. Dell System Setup: Device Settings

3. In the Device Settings window, select your device and then press ENTER.

   The Main Configuration Page for the selected device appears. The Main Configuration Page is the top-level menu of the adapter configuration pages. Figure 3-10 shows an example of the Main Configuration Page.

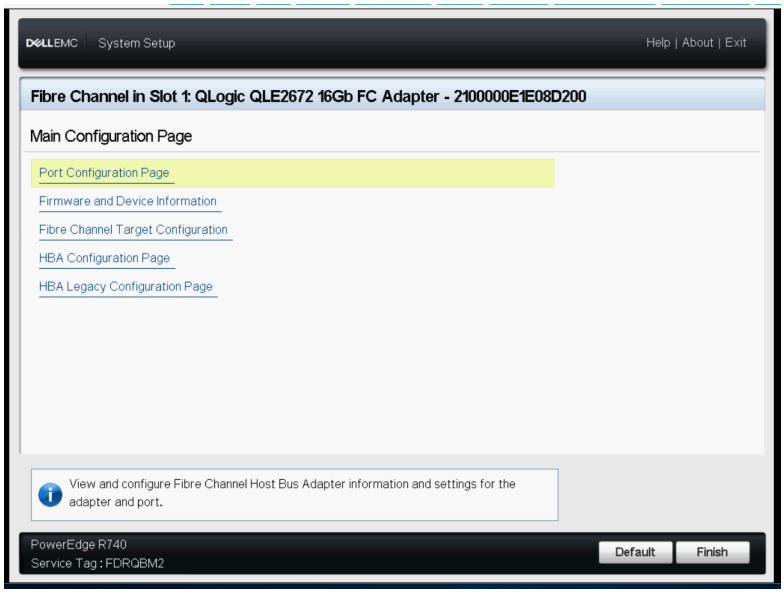

Figure 3-10. Dell System Setup: Main Configuration

4. On the Main Configuration Page, select **Port Configuration Page** and press ENTER.

The Port Configuration Page appears; Figure 3-11 shows an example.

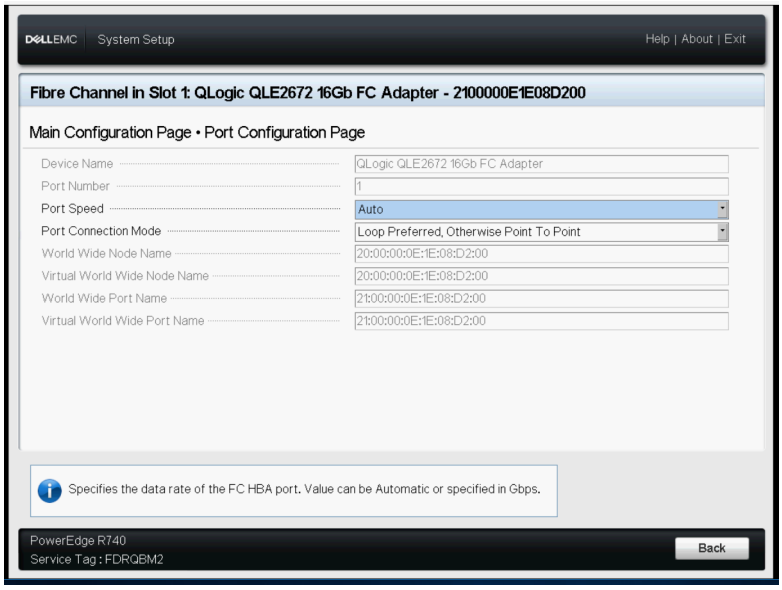

Figure 3-11. Dell System Setup: Port Configuration

5. Use the Port Configuration Page to change adapter operational parameters, such as port speed. In most cases, use the default values. Table 3-4 describes the configurable options on the Port Configuration Page.

Table 3-4. Port Configuration Page Settings

Setting	Description
Device Name	Indicates the name of the adapter.
Port Number	Indicates the index of the selected port (1 to N). Figure 3-11 shows port 1 in a two-port adapter.
Port Speed	Indicates the data rate of the Fibre Channel adapter port. This value can be automatic or specified in GFC. The adapters support 4Gbps, 8Gbps, and 16Gbps.
Port Connection Mode	Indicates the connection mode of the Fibre Channel adapter port. The adapters support the following: ■ Loop Only ■ Point to Point ■ Loop Preferred, Otherwise Point To Point
World Wide Node Name	Indicates the adapter port's unique worldwide node name (WWNN).

Table 3-4. Port Configuration Page Settings (Continued)

Setting	Description
Virtual World Wide Node Name	Indicates the adapter port's unique virtual WWNN.
World Wide Port Name	Indicates the adapter port's unique WWPN.
Virtual World Wide Port Name	Indicates the adapter port's unique virtual WWPN.

6. To return to the Main Configuration Page, click **Back**.
7. On the Main Configuration Page, select **Firmware and Device Information**, and then press ENTER.

   The Firmware and Device Information page provides the adapter and firmware version information, and port address information. Figure 3-12 shows an example.

n

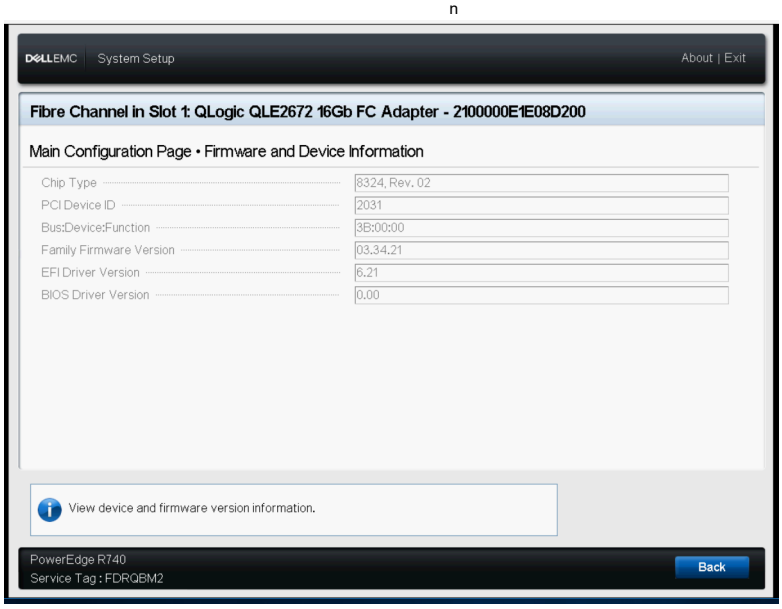

Figure 3-12. Dell System Setup: Firmware and Device Information

Table 3-5 describes the Firmware and Device Information settings.

Table 3-5. Firmware and Device Information Settings

Setting	Description
Chip Type	Type of intelligent storage peripheral (ISP) controller used in the adapter.
PCI Device ID	Unique PCI device ID of the adapter.
Bus:Device:Function	PCI address of the adapter in the Bus:Device:Function (BDF) notation.
Family Firmware Version	Adapter firmware version information.
EFI Driver Version	Adapter extensible firmware interface (EFI) driver version information.

8. To return to the Main Configuration Page, click **Back**.
9. On the Main Configuration Page, select **Fibre Channel Target Configuration,** and then press ENTER.

   The Fibre Channel Target Configuration page appears. Figure 3-13 shows an example.

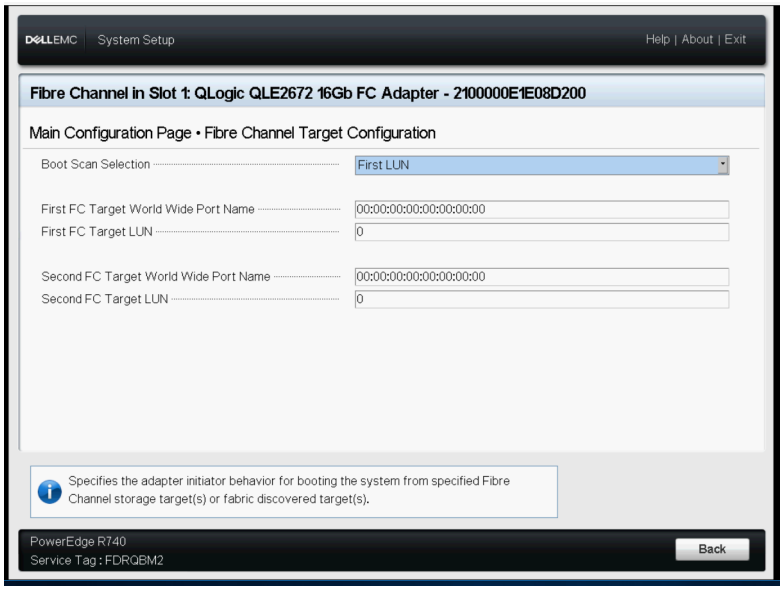

Figure 3-13. Fibre Channel Target Configuration

10. On the Fibre Channel Target Configuration page, select SAN storage devices for boot from SAN. With selective boot scan enabled, the target configuration determines which target LUNs to map. All other devices are ignored. Table 3-6 describes the Fibre Channel Target Configuration settings.

Table 3-6. Fibre Channel Target Configuration Settings

Setting	Description
Boot Scan Selection	Specifies the adapter initiator behavior for booting the system from the specified Fibre Channel storage target(s) or fabric discovered target(s).
First FC Target World Wide Port Name	Specifies the WWPN of the first Fibre Channel storage target.
First FC Target LUN	Specifies the LUN of the first Fibre Channel storage target.
Second FC Target World Wide Port Name	Specifies the WWPN of the second Fibre Channel storage target.
Second FC Target LUN	Specifies the LUN of the second Fibre Channel storage target.

NOTE

System UEFI compatibility: To boot from a QLogic host 2600 Series Adapter in a computer system with a multiboot system UEFI, the 2600 Series Adapter must be included as the first device in the list of bootable devices in the system's **Boot** menu.

Restoring defaults in HII does not zero out target WWPN or target LUN information. This is expected behavior.

11. To return to the Main Configuration Page, click **Back**.

12. On the Main Configuration Page, select **HBA Configuration Page,** and then press ENTER.

    The HBA Configuration Page appears; Figure 3-14 shows an example.

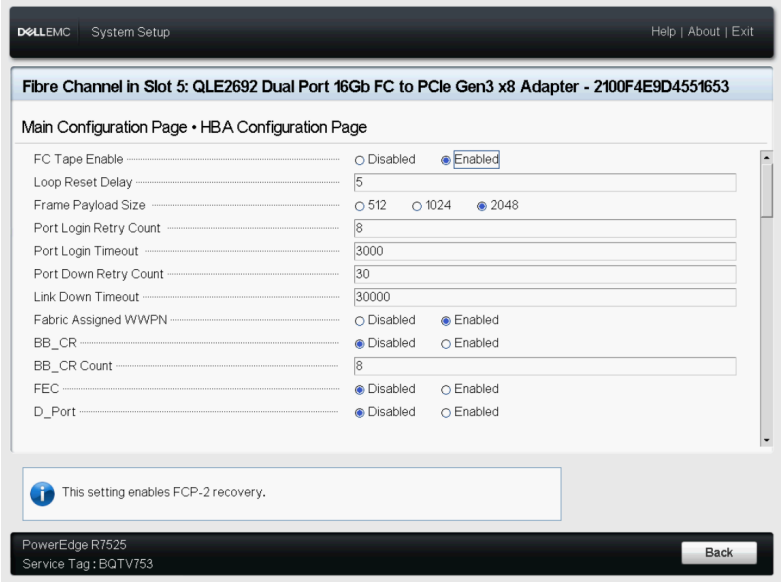

Figure 3-14. Dell System Setup: HBA Configuration

NOTE

FC-NVMe requires one or more of the QLE2690-DEL, QLE2690L-DEL, QLE2690N-DEL, QLE2690NL-DEL, QLE2692-DEL, QLE2692L-DEL, QLE2692N-DEL, QLE2692NL-DEL, or QME2692-DEL adapters.

13. Use the HBA Configuration Page to configure adapter operational parameters. In most cases, use the default values.

Table 3-7 describes the fields on the HBA Configuration Page.

Table 3-7. HBA Configuration Page Settings

Setting	Description
FC Tape Enable	Enables or disables Fibre Channel protocol (FCP-2) recovery.
Loop Reset Delay	Specifies the Fibre Channel arbitrated loop reset delay for the port. After resetting the loop, the adapter firmware refrains from initiating any loop activity for the number of seconds specified. The range is 0 to 60.
Frame Payload Size	Specifies the maximum Fibre Channel frame payload size.
Port Login Retry Count	Specifies the number of times the adapter firmware initiator tries to log in to a target device port. The range is 0 to 255.
Port Login Timeout	Specifies the time-out in milliseconds that the initiator uses when attempting to log in to the target device port. The range is 0 to 255000.
Port Down Retry Count	Specifies the number of seconds the software retries a command to a port returning port-down status. The range is 0 to 255.
Link Down Timeout	Specifies the number of milliseconds the Fibre Channel uplink port can be offline before the system is notified that the SAN fabric connectivity is lost and the uplink port is marked as being down. The range is 1 to 255000.
Fabric Assigned WWPN	Enabled Fabric-assigned worldwide port name (FA-WWPN) allows you to enable an adapter port to use a switch-assigned WWPN rather than the physical adapter port WWPN for communication. (Also known as FA-PWWN)
D_Port	D_Port (diagnostic port) mode allows the adapter to identify and isolate link failures resulting from faulty modules (link, cable, or SFP). D_Port mode requires a Brocade® Fibre Channel switch with the ClearLink® D_Port feature installed. For additional D_Port details, see the Brocade switch documentation. This feature applies only to QLE2690 models.

Table 3-7. HBA Configuration Page Settings (Continued)

Setting	Description
BB_CR	Buffer-to-buffer credit recovery (BB-CR) enables two FC peer ports (N_Port, F_Port, or E_Port) to periodically send and receive the quantity of receiver ready (R_RDY) signals transmitted. Enable the BB-CR feature to allow the peer port to recover from possible R_RDY signals lost over a lossy link. This feature applies only to QLE2690 models.
BB_CR Count	Specifies the quantity of frame RX/TX counters maintained by the port to track R_RDYs and frames received. This feature applies only to QLE2690 models.
FEC	Forward error correction (FEC) improves performance and link integrity to support higher end-to-end data rates by automatically recovering from transmission errors. This setting applies only to 16Gbps speeds. This feature applies only to QLE2690 models.
USCM	Enables Universal SAN Congestion Mitigation. The USCM feature detects and prevents both potential or actual congestion occurrences in the FC SAN environment.
Enable NVMe Boot from SAN	Enables the NVME storage feature. When enabled, the driver will search for NVME storage, based on the settings below.

Table 3-7. HBA Configuration Page Settings (Continued)

Setting	Description
Host NQN	This field identifies the FC adapter. NVME storage devices typically use this value in access lists. The default value of this field is based on the system UUID. All FC adapters in a system will have the same default Host NQN. This field can be changed, but many storage arrays expect a specific Host NQN format. See the NVME Base spec for more details about the Host NQN format. The OS driver also uses a Host NQN. The UEFI driver and OS driver must use the same Host NQN. To view the OS driver Host NQN value, use the QCC CLI tool. Use “qaucli -i” to display the OS driver Host NQN value. If the Host NQN starts with `nqn.2014-08.org.nvmexpress:uuid:`, the UUID value is automatically populated by the driver and cannot be changed. User defined Host NQNs must not start with `nqn.2014-08.org.nvmexpress:uuid:`. An example of a user defined Host NQN is: `nqn.2014-08.com.example:nvme.host.sys.xyz`. **Configuring custom NQN in HII for NVME boot-from-SAN (BFS) is not supported in all operating systems. For more information, see “NVM Express over Fibre Channel” on page 101.**
Host ID	This field overrides the default Host ID and is used when the NVME storage requires a specific Host ID. The Host ID is a sequence of 32 hex digits (non-hex characters cannot be used). Example: 00112233445566778899AABBCCDDEEFF. If an invalid Host ID is entered, the driver saves the value as 32 zero-digits and there is no error message displayed. If the Host NQN field starts with `nqn.2014-08.org.nvmexpress:uuid:` and the UUID changes, the Host ID will be updated to match the new UUID.
Connect to First NVME Target	Enables a specific NVME storage device. When enabled, the device is mapped by the UEFI FC driver.
First NVME Target World Wide Port Name	The World Wide Port Name of an NVME storage device. The WWPN can be obtained from the NVME storage device configuration tool.

Table 3-7. HBA Configuration Page Settings (Continued)

Setting	Description
First NVME Target World Wide Node Name	The World Wide Node Name of an NVME storage device. This field is optional. A value of 00:00:00:00:00:00:00:00 means ignore this field. The WWNN can be obtained from the NVME storage device configuration tool.
First NVME Target NVME Qualified Name	Identifies an NVME storage Subsystem. Multiple Subsystems can be attached to an NVME storage device WWPN. The Storage NQN can be obtained from the NVME storage device configuration tool.
First NVME Target Controller ID	Identifies a Controller attached to an NVME Subsystem. Multiple Controllers can be attached to an NVME Subsystem. In most cases, FFFFh is the proper value for this field. A value of FFFFh means any available Controller.
First NVME Target Namespace ID	Identifies a Namespace attached to an NVME Controller. Multiple Namespaces can be attached to an NVME Controller. The Namespace ID can be obtained from the NVME storage device configuration tool.

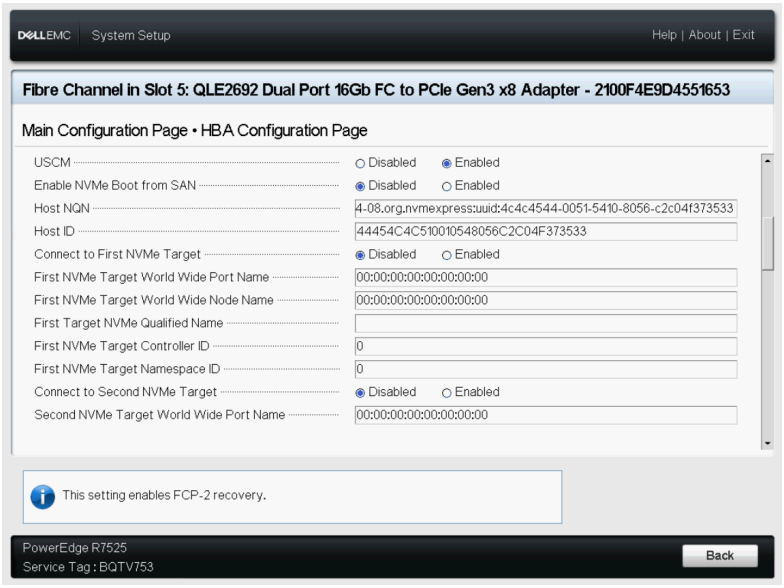

Figure 3-15. NVME options

Each mapped NVME storage device requires setting the following fields: Connect to First NVME Target, NVME Target WWPN, NVME Target NQN, NVME Target Controller ID, NVME Target Namespace ID. Up to eight storage devices can be mapped.

NVME storage device configuration values can be obtained from the NVME storage configuration tool. Target WWPN, WWNN, and NQN information are provided by the configuration tool.

NOTE

Marvell recommends performing Windows Boot from SAN using a single path, regardless of FC or FC-NVMe protocol. Refer to the following link for the Dell *Boot from SAN Implementation and Best Practices Guide*.

https://www.dell.com/community/Host-Systems/Boot-from-SAN-Implementation-and-Best-Practices-Guide/td-p/6984863

14. Click **Finish** to save your changes and exit.

Setting Queue Depth

For a detailed explanation of queue depth, including “Setting the Host Bus Adapter Queue Depth,” refer to the *Best Practices Guide—Fibre Channel Host Bus Adapters on Microsoft Windows 2012 and VMware ESXi 5.x, 2500 Series and 2600 Series* (part number SN0454502-00), available for download from the Marvell Web site.

Queue depth refers to the maximum LUN queue depth to report for target devices. This section provides the default values and instructions for modifying the queue depth value configured by the Windows, Linux, and ESXi drivers.

Changing the Windows Queue Depth

The default queue depth value on Windows is 20 and the range is 0–254.

To change the default value:

1. Click **Start**, select **Run**, and open the REGEDIT/REGEDT32 program.
2. Select `HKEY_LOCAL_MACHINE` and then expand the tree structure down to the QLogic driver as follows:

```
HKEY_LOCAL_MACHINE
     SYSTEM
          CurrentControlSet
               Services
                    Ql2300
```

```
Parameters
    Device
```

3. Double-click **`DriverParameter:REG_SZ:qd=32`**.
4. If the string `qd=` does not appear, append it to end of the string as **`;qd=32`**.
5. Click **OK**.
6. Exit the Windows Registry Editor, and then shutdown and reboot the system.

Changing the Linux Queue Depth

The default queue depth value on Linux is 64 and the range is 1–65535.

To change the default value:

1. Set the module parameter by adding the following line to the `/etc/modprobe.d/qla2xxx.conf` file:

   `options qla2xxx ql2xmaxqdepth=X`

   where `X` is the new queue depth value.
2. Save and close the `modprobe.conf` file.

Changing the ESXi Queue Depth

The default queue depth value on ESXi is 64 and the range is 0–65535.

To change the default value:

1. Issue the following command:

   `$ esxcfg-module -s "ql2xmaxqdepth=X" qlnativefc`
   where `X` is the new queue depth value.
2. Restart the system.

Setting an FA-WWPN

This section provides the following information for fabric-assigned world wide port name (FA-WWPN):

- Setting the FA-WWPN from the Adapter
- Setting the FA-WWPN from the Brocade Switch

When using the FA-WWPN feature:

- FA-WWPN must be enabled on both the initiator (2600 Series Adapter) and the Brocade switch for the FA-WWPN to be assigned to the 2600 Series Adapter.

- If FA-WWPN is disabled on the 2600 Series Adapter and enabled on the Brocade switch, the FA-WWPN cannot be assigned to the 2600 Series Adapter.
- If FA-WWPN is enabled on the 2600 Series Adapter and disabled on the Brocade switch, the FA-WWPN cannot be assigned to the 2600 Series Adapter.

Setting the FA-WWPN from the Adapter

From the adapter side, you can set the FA-WWPN using the following utilities:

- Setting FA-WWPN with Interactive QConvergeConsole CLI
- Setting FA-WWPN with QConvergeConsole VMware Plug-ins

Setting FA-WWPN with Interactive QConvergeConsole CLI

To set the FA-WWPN from the adapter using interactive QConvergeConsole CLI:

1. From the main menu, select the adapter type (**Fibre Channel Adapter**).
2. On the **Adapter Configuration** menu, select **HBA Parameters**.
3. On the port menu, select a port to open the **HBA Parameters Menu**.
4. Select **Change Settings**.
5. Select **Fabric Assigned WWN**.
6. Under **Enable Fabric Assigned WWN Menu**:
   - ❑ Type **1** to enable FA-WWPN.
   - ❑ Type **2** to disable FA-WWPN.

The following shows an example of setting the FA-WWPN:

```
HBA Desc.      : QLE2672 QLogic 2-port 16Gb Fibre Channel Adapter
FW Version     : 7.04.00
WWPN           : 21:00:00:oe:1e:11:4a:38
WWNN           : 21:00:00:oe:1e:11:4a:38
Host NQN       : nqn.2014-08.org.nvmexpress:uuid:
                 4c4c4544-0047-5110-8034-c7c04f514432
Host ID        : 44454c4c470010518034c7c04f514432
Link           : Online
================================================================

    1:  Connection Options
    2:  Data Rate
    3:  Frame Size
```

```
 4:  HBA Hard Loop ID
 5:  Hard Loop ID
 6:  Loop Reset Delay (seconds)
 7:  Host Adapter BIOS
 8:  Fibre Channel Tape Support
 9:  Operation Mode
10:  Interrupt Delay Timer (100 microseconds)
11:  Execution Throttle[1]
12:  Login Retry Count
13:  Port Down Retry Count
14:  LIP Full Login
15:  Link Down Timeout (seconds)
16:  Target Reset
17:  LUNs per Target
18:  LR Extended Credits
19:  Fabric Assign WWN
20:  Prefer FCP Support
21:  USCM Support
22:  Virtual Lane
23:  Commit Changes
24:  Abort Changes
     (p or 0: Previous Menu; m or 98: Main Menu; ex or 99: Quit)
     Please enter selection: 19

QConvergeConsoleCLI

CLI - Version 2.6.x (Build xx)


 Fabric Assign WWN Menu

     1: Enable
     2: Disable (current)

     (p or 0: Previous Menu; m or 98: Main Menu; ex or 99: Quit)
     Please enter selection: 1
```

[1] The `Execution Throttle` parameter is no longer supported for Fibre Channel 16G and higher.

After refreshing the Brocade switch, the **Adapter Information** shows the port name as shown in the following example.

```
QConvergeConsoleCLI

CLI – Version 2.6.x (Build xx)

 FC Adapter Information

     1: FC Adapter Information
     2: FC Port Information
     3: FC VPD Information
     4: FC Storage Device Information
     5: FC Hyper-V VFC Information
     6: FC Adapter FMB Information

        (p or 0: Previous Menu; m or 98: Main Menu; ex or 99: Quit)
        Please enter selection: 2

       QConvergeConsoleCLI
       CLI - Version 2.6.x (Build xx)

  Adapter Information
  HBA Model QLE2672 SN: RFE12340H08416
    1: Port   1: WWPN: 21:00:00:0f:1f:11:4a:30 Link Down
    2: Port   2: WWPN: 50:00:53:37:63:fa:00:05 Online
```

Setting FA-WWPN with QConvergeConsole VMware Plug-ins

For information about setting a FA-WWPN using the QConvergeConsole VMware vCenter Server Plug-in or QConvergeConsole VMware vSphere Web Client Plug-in, see the section “Managing a Fibre Channel Adapter Port” in the *User’s Guide—QConvergeConsole Plug-ins for vSphere* (part number TD-000965).

Setting the FA-WWPN from the Brocade Switch

NOTE

When setting up this feature from the switch, Brocade identifies it as FA-PWWN.

From the Brocade switch, you can set a FA-WWPN and a static FA-WWPN. For the most current information, refer to the Brocade switch documentation.

Setting the FA-WWPN from the Switch

Figure 3-16 shows the FA-WWPN setting from the Brocade switch.

```
portdisable 17
fapwwn --enable -port 17
portenable 17
fapwwn --show all
-----------------------------------------------------------------------------------------------------------
       AG Port              port     Device Port WWN            Virtual Port WWN        PID   Enable   MapType
-----------------------------------------------------------------------------------------------------------
--:--:--:--:--:--:--:--/---   17    --:--:--:--:--:--:--:--   50:00:53:37:63:fa:00:05    --     Yes    Port/User
```

Figure 3-16. Setting the FA-WWPN from the Brocade Switch

After refreshing the Brocade switch, the **Device Port WWN** appears in Web Tools as shown in Figure 3-17.

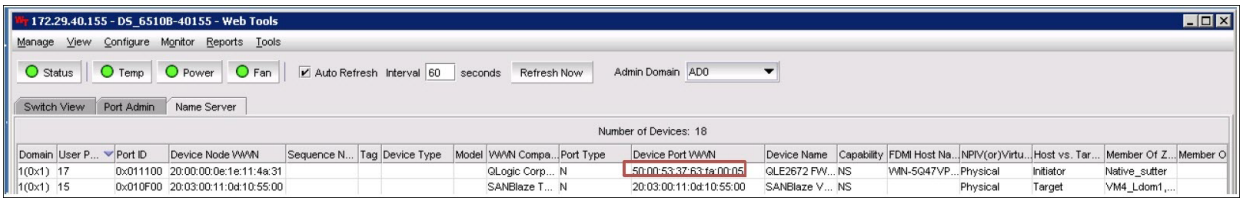

Figure 3-17. Device Port WWN for Brocade Switch

Setting the Static FA-WWPN from the Switch

Figure 3-18 shows the static FA-WWPN setting from the Brocade switch.

```
portdisable 17
fapwwn --assign -port 17 -v 51:00:53:37:63:fa:02:11
portenable 17
fapwwn --show all
-----------------------------------------------------------------------------------------------------------
       AG Port              port     Device Port WWN            Virtual Port WWN        PID   Enable   MapType
-----------------------------------------------------------------------------------------------------------
--:--:--:--:--:--:--:--/---   17    --:--:--:--:--:--:--:--   51:00:53:37:63:fa:02:11    --     Yes    Port/User
                                                            50:00:53:37:63:fa:00:05
```

Figure 3-18. Setting the Static FA-WWPN from the Brocade Switch

After refreshing the Brocade switch, the **Device Port WWN** appears in Web Tools as shown in Figure 3-19.

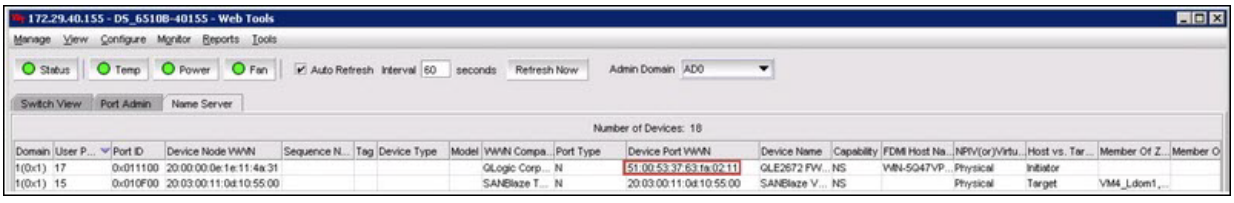

Figure 3-19. Device Port WWN (Static) for Brocade Switch

Configuring and Verifying FA-BLD

This section provides the following information on fabric-assigned boot LUN discovery (FA-BLD):

- Configuring FA-BLD from the Adapter
- Verifying that FA-BLD is Operational
- Adapter Side Restrictions

Configuring FA-BLD from the Adapter

On the adapter side, follow these procedures to configure fabric-assigned boot LUN discovery (FA-BLD) using QConvergeConsole CLI interactive mode.

Configuring the Adapter and Boot Devices with Interactive QConvergeConsole CLI

To configure adapter and boot devices:

1. In QConvergeConsole CLI, navigate to Adapter Configuration and then select **Configure Boot Devices**.
2. Select the adapter port.
3. In the Configure Boot Devices menu, type one of the following, and then press ENTER:
   - ❑ **1** (Legacy BIOS Boot Mode) (see For Legacy BIOS Boot Mode:)
   - ❑ **2** (UEFI Boot Mode) (see For UEFI Boot Mode:)

For Legacy BIOS Boot Mode:

1. In the Boot from SAN BIOS Configuration menu, type **1** (Boot-from-SAN and Drive Mapping Parameters) and then press ENTER.
2. In the Boot from SAN BIOS Parameters menu:
   a. Type **2** (Configure BIOS Boot Parameters) and then press ENTER.
   b. Type **3** (Fabric Assign Boot LUN), and then press ENTER.
   c. Type **1** (Enable) and then press ENTER.

For UEFI Boot Mode:

1. In the Boot from SAN BIOS Configuration menu, type **2** (UEFI Boot Mode) and then press ENTER.
2. In the Boot from SAN UEFI Parameter menu, type **5** (Fabric Assign Boot LUN), and then press ENTER.
3. In the Fabric Assign Boot LUN menu, type **1** (Enable), and then press ENTER.

On the adapter side, issue the following command to configure fabric-assigned FA-BLD using QConvergeConsole CLI non-interactive mode.

```
#qaucli.exe -e 0 enable EnableFabricAssignBootLun  <boot mode>
```

Where `<boot mode>` is one of the following:

`--bios` = BIOS or legacy BIOS

`--uefi` = UEFI (FCP boot target)

`--nvme` = UEFI (NMVe target)

Configuring a Zone in a Brocade Switch

To configure a zone in a Brocade switch:

1. Create a boot LUN configuration as follows:

   ```
   bootluncfg --add 50:00:53:37:63:FA:00:05
   20:02:00:11:0d:51:5c:01 0000000000000000
   ```

2. Create a zone named with the following:
   - ❑ Device's WWPN
   - ❑ Specific information about the target WWN
   - ❑ LUN to be booted from

3. Add a zone to the zoneset configuration as follows:

   ```
   cfgadd config, "BFA_5000533763fa0005_BLUN"
   ```

4. To save the zoneset, issue the following command:

   ```
   cfgsave
   ```

5. Enable zoneset configuration to have it actively running:

   ```
   cfgenable config
   ```

Figure 3-20 shows an example of a configured zone in a Brocade switch.

```
Name of zone contains the PWWN: "BFA_5000533763fa0005_BLUN", and the 4 members of the zone are described as shown below:
1st member       00:00:00:00:                    20:02:00:11 [Target's PWWN, first 4 octets]
2nd member       00:00:00:01:                    0d:51:5c:01 [Target's PWWN, last 4 octets]
3rd member       00:00:00:02:                    00:00:00:00 [Target's LUN, first 4 octets]
4th member       00:00:00:03:                    00:00:00:00 [Target's LUN, last 4 octets]

                 zonecreate "BFA_5000533763fa0005_BLUN", "00:00:00:00:20:02:00:11; 00:00:00:01:0d:51:5c:01;
                 00:00:00:02:00:00:00:00; 00:00:00:03:00:00:00:00"
```

Figure 3-20. Configuring a Zone in a Brocade Switch

Verifying that FA-BLD is Operational

To verify if the FA-BLD is operational, use either *Fast!*UTIL or the system.

Figure 3-21 shows *Fast!*UTIL displaying the boot LUN.

Figure 3-21. Verifying FA-BLD with Fast!UTIL

Figure 3-22 shows using the system with a LUN ready to install or OS booted from LUN.

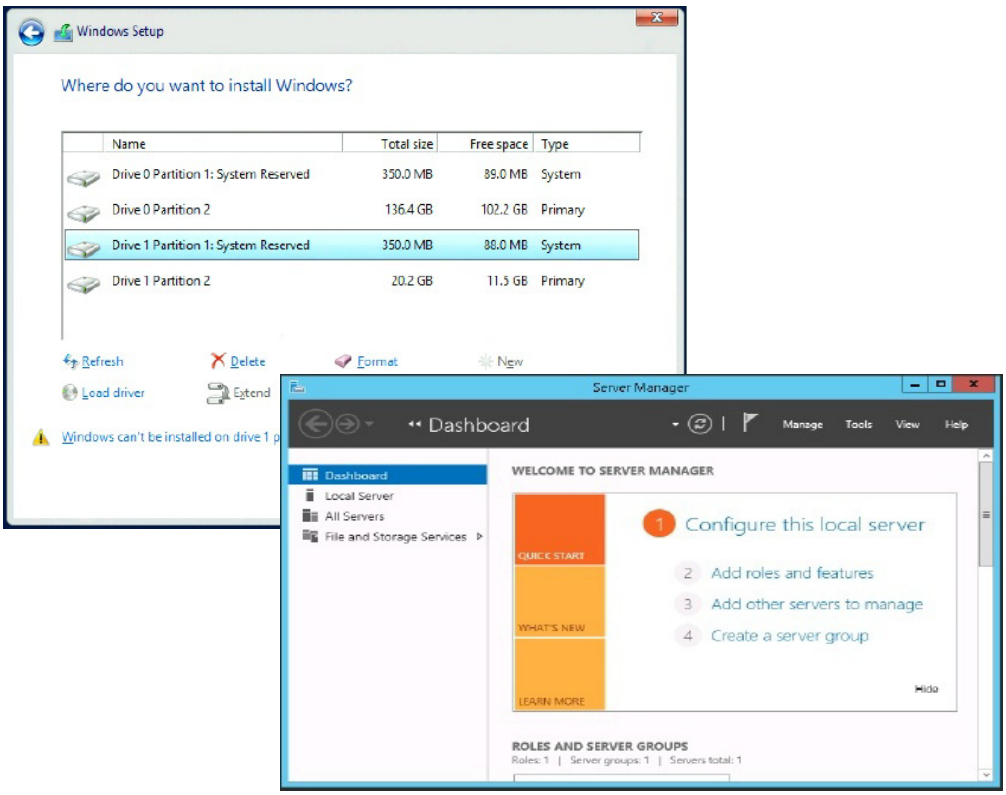

Figure 3-22. Verifying FA-BLD from the System

Adapter Side Restrictions

The known fabric-assigned port world wide name (FA-PWNN) and FA-BLD restrictions include the following:

- Cannot be fully enabled from *Fast!*UTIL
- Do not use **Restore Defaults** from *Fast!*UTIL, QConvergeConsole GUI, or QConvergeConsole CLI. Doing so causes these capabilities to become disabled.

Using a Fabric-Assigned Boot LUN

This section describes how to use a fabric-assigned boot LUN in the following utilities:

- Using a Fabric-Assigned Boot LUN in Interactive QConvergeConsole CLI
- Using a Fabric-Assigned Boot LUN with Noninteractive QConvergeConsole CLI
- Using a Fabric-Assigned Boot LUN with QConvergeConsole Plug-ins

Using a Fabric-Assigned Boot LUN in Interactive QConvergeConsole CLI

To use a fabric-assigned boot LUN with interactive QConvergeConsole CLI:

1. From the main menu, select the adapter type (**Fibre Channel Adapter**).
2. Select an adapter port.
3. On the **Adapter Configuration** menu, select **Configure Boot Devices**.
4. On the **Boot from SAN BIOS Configuration** menu, select **Boot-from-SAN Drive Mapping Parameters**, and then press ENTER.
5. Select **Configure BIOS Boot Parameters**, and then press ENTER.
6. On the **Boot from SAN BIOS Parameters** menu, select **Fabric Assign Boot LUN**, and then press ENTER.
7. To enable this feature, type **1**.

Using a Fabric-Assigned Boot LUN with Noninteractive QConvergeConsole CLI

In noninteractive QConvergeConsole CLI, set a target as the fabric-assigned boot LUN for a specific adapter by issuing the following command:

```
# qaucli -pr fc -e  <hba instance>|<hba wwpn> enable|disable
FabricAssignBootLUN|fb
```

Where:

`hba instance` = Adapter number (use the `-g` command to find)

`hba wwpn` = Worldwide port name of the adapter

`FabricAssignBootLUN` or `fb` = Configure enable fabric-assigned boot LUN setting

Using a Fabric-Assigned Boot LUN with QConvergeConsole Plug-ins

For information about using a fabric-assigned boot LUN for QConvergeConsole VMware vSphere Web Client Plug-in, see the section “Configure Fibre Channel Port Boot Parameters” in the *User’s Guide—QConvergeConsole Plug-ins for vSphere* (part number TD-000965).

To configure a 2600 Series Adapter boot device with QConvergeConsole VMware vCenter Server Plug-in:

1. In the QConvergeConsole VMware vCenter Server Plug-in, click the **QConvergeConsole** tab.
2. In the tree pane on the left, select a 2600 Series Adapter port.
3. Click the **Boot** tab.

4. On the Boot Device Selection page, select the **Enable Fabric Assigned Boot LUN** check box. Figure 3-23 shows an example.

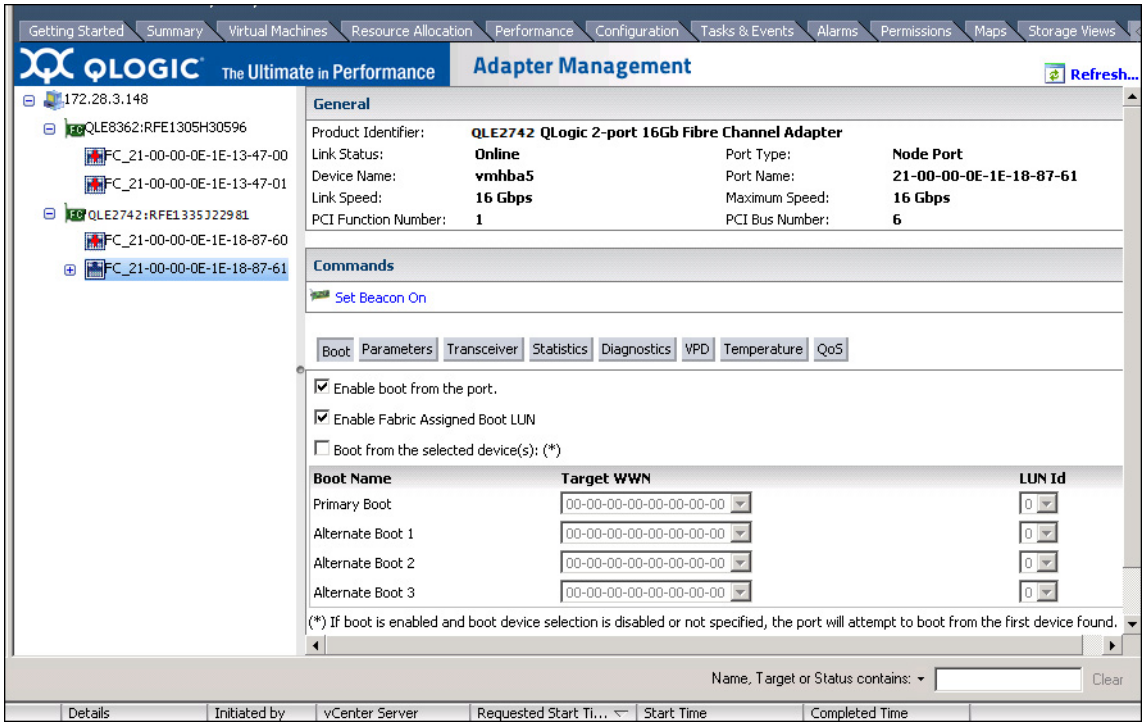

Figure 3-23. Fabric-Assigned Boot LUN in QConvergeConsole vSphere Client Plug-in

For details, refer to the “Boot Device Selection” topic in *QConvergeConsole Help*.

5. Click **Save**. Any previously saved configuration for the current 2600 Series Adapters is overwritten.

Running Diagnostics—Fibre Channel Ping and Trace Route

This section provides the following information on how to run Fibre Channel ping and trace route diagnostics:

- Pinging and Viewing Trace of Routing Using Topology Map
- Running a Fibre Channel CT Ping Test

Pinging and Viewing Trace of Routing Using Topology Map

The QConvergeConsole GUI topology map uses the common transport (CT) layer to enable the adapter's initiator port to ping or run a trace route that will navigate across the environment through multiple ports of the switch or switches and to the target.

Table 3-8 compares the differences between the Ping Test, CT Ping Test, and the CT FTR Tests.

Table 3-8. Ping Test, CT Ping Test, and CT FTR Test

Ping Test	CT Ping Test	CT FTR Test
SCSI layer ping	CT or Fibre Channel ping	CT or Fibre Channel trace route
Uses SCSI inquiry command	Uses echo through an external link	Requires scan of physical topology
Can operate on any SCSI device	Can operate with all Fibre Channel devices that support CT or an external link	Can operate with all devices that are connected through a port that supports CT or an external link

Supported configurations for Fibre Channel ping and trace route include:

- Marvell QLogic 2600 Series Adapters
- Brocade Switches with Fabric OS® (FOS) v6.0.0a or later
- Cisco® Switches with NX-OS v8.3.1 (FC Trace Route not supported)

NOTE

The switch and the Fibre Channel end device must both support the Fibre Channel ping and trace route feature. See Table 3-11 on page 86 for the supported Fibre Channel software version.

Running a Fibre Channel CT Ping Test

You can run a Fibre Channel CT ping test from QConvergeConsole CLI or the QConvergeConsole VMware plug-ins, as described in this section.

Running a Fibre Channel CT Ping from QConvergeConsole CLI

To run a Fibre Channel CT ping test from QConvergeConsole CLI:

1. From the QConvergeConsole CLI **FC Adapter Diagnostics** menu, select **7: CT Ping Test**.
2. From the **HBA Model** list, select the port that you want to ping.
3. From the **CT Ping Test** menu, select **4: Start**.

The following shows example output from the CT ping test.

```
-----------------------------------------------
Diagnostics Test Configuration
-----------------------------------------------
Diagnostic Mode           : CT Ping
Number of tests (1-10000) : 10
Number of Pass            : 1
Test Increment (1-10000)  : 1
Abort On Error            : Ignore
Test Continuous           : OFF

-------------------------------------------------
    ID      Data         Link      Sync       Signal      Invalid    Diagnostic
Port/Loop  Miscompare   Failure   Loss       Loss        CRC        Status
---------  ----------   -------   --------   ---------   ---------  ----------
01-0B-01            0         0          0           0           0     Success
```

NOTE

The default setting will repeat 10 times with output similar to the preceding for each pass of the test.

Running a Fibre Channel CT Ping from QConvergeConsole VMware Plug-ins

For information about running Fibre Channel CT ping from the QConvergeConsole VMware vCenter Server Plug-in, see the section “Diagnostic Tests—Fibre Channel Port” in the *User’s Guide—QConvergeConsole Plug-ins for vSphere* (part number TD-000965).

Configuring CS_CTL QoS

NOTE

For information about **end-to-end** CS_CTL QoS, see “Configuring End-to-End CS_CTL QoS” on page 79.

This section provides the following information on class-specific control (CS_CTL) quality of service (QoS):

- CS_CTL QoS Features
- Enabling CS_CTL QoS Mode for the Initiator and Target Ports
- Verifying and Confirming CS_CTL Mode Setup for Each Port

CS_CTL QoS Features

The following highlights some of the features of the CS_CTL QoS:

- Traffic priority classification for QLogic Gen 5 Fibre Channel SAN. QoS at the packet level is realized by using the CS_CTL field, which is specified in the FC-FS-2 specification. QoS allows one to prioritize frames between a host and target, depending on the value of the CS_CTL field in the FC Frame header.
- Extending QoS from the fabric to the host
- Enables Marvell QLogic Fibre Channel users to rapidly scale server virtualization without compromising service-level agreements
- Benefits for both physical and virtual server deployments
- Supported configuration:
  - Marvell QLogic 2600 Series Adapters
  - CS_CTL-based QoS is end-to-end (requires target to support CS_CTL)

Enabling CS_CTL QoS Mode for the Initiator and Target Ports

Figure 3-24 shows how to enable CS_CTL QoS mode for the initiator and target ports.

```
portdisable 17
portcfgqos --enable 17 csctl_mode
Enabling CSCTL mode flows causes QoS zone flows to lose priority on such ports.
Do you want to proceed?(y/n):y
portenable 17

portdisable 0
portcfgqos --enable 0 csctl_mode
Enabling CSCTL mode flows causes QoS zone flows to lose priority on such ports.
Do you want to proceed?(y/n):y
portenable 0
```

Figure 3-24. Enabling CS_CTL QoS Mode

Verifying and Confirming CS_CTL Mode Setup for Each Port

Figure 3-25 shows how to verify the CS_CTL mode setup for each port.

```
portcfgshow
Ports of Slot 0          0   1   2   3     4   5   6   7     8   9  10  11    12  13  14  15
-----------------------+---+---+---+---+-----+---+---+---+-----+---+---+---+-----+---+---+---
Octet Speed Combo        1   1   1   1     1   1   1   1     1   1   1   1     1   1   1   1

<ONLY PARTIAL OUTPUT SHOWN BELOW FOR THE APPROPRIATE AREA TO CHECK ON PORT 0 >

CSCTL mode               ON  ..  ..  ..    ..  ..  ..  ..    ..  ..  ..  ..    ..  ..  ..  ..
```

Figure 3-25. CS_CTL Mode Setup for Each Port

Configuring End-to-End CS_CTL QoS

NOTE

For information about CS_CTL QoS that is not end-to-end, see “Configuring CS_CTL QoS” on page 77.

QLogic’s class specific control (CS_CTL) QoS implementation is an end-to-end solution, encompassing the initiator (2600 Series Adapters), the switches, and the target (storage). The 2600 Series Adapter stamps the priority value (high, medium, low) in the CS_CTL field of the Fibre Channel frame. All three entities (initiator, switch, and target) honor this priority value at their respective levels.

This feature is enabled by default in the 2600 Series Adapter Windows and VMware drivers. Therefore, when all three entities support QoS (and it is enabled on the switch and the target), a default priority value is populated in all FCP_CMND, FCP_DATA, FCP_XFER_RDY, FCP_RSP, and FCP_CONFIRM frames.

NOTE

QoS must be supported by both the switch and target devices.

The following sections describe how to configure end-to-end CS_CTL QoS and change the default priority levels:

- Configuring CS_CTL QoS on the Switch
- Configuring CS_CTL QoS on the Storage Device
- Changing the QoS Priority Levels

Configuring CS_CTL QoS on the Switch

To enable CS_STL QoS on the Brocade SW6510 switch:

1. From the switch, issue the following command:

   ```
   root> configurechassis
   Configure...
   ```

2. A list of attributes appear one at a time. After either accepting the default or entering a different value, the next attribute appears.

   - Press ENTER to accept the default for `cfgload attributes`, `Custom attributes`, and `system attributes`.
   - Type **`y`** for `fos attributes`.
   - Type **`1`** for `CSCTL QoS Mode`.

   For example:

   ```
   cfgload attributes (yes, y, no, n): [no]
   Custom attributes (yes, y, no, n): [no]
   system attributes (yes, y, no, n): [no]
   fos attributes (yes, y, no, n): [no] y
   CSCTL QoS Mode (0 = default; 1 = auto mode): (0..1) [0] 1
   ```

3. Reboot the switch.

4. To verify that CS_CTL QoS is enabled, issue the following command:

   ```
   root> configshow -all | grep fos.csctlMode
   ```

   This command should return the following response:

   ```
   fos.csctlMode:1
   ```

5. To enable QoS on a switch-port basis, issue the following command:

   ```
   portcfgqos --enable [slot/]port csctl_mode
   ```

Configuring CS_CTL QoS on the Storage Device

See your device's user guide for QoS configuration instructions.

Changing the QoS Priority Levels

The QoS priority levels are set in the 2600 Series Adapters vPorts. The Windows and VMware adapter drivers assign default vPort priority levels. The following sections describe how to change the default vPort priority levels in Windows and VMware ESXi systems.

Changing the Priority Levels in Windows

To change the vPort priority levels in Windows, you must first assign the vPorts, and then change the priority levels. Both procedures can be accomplished using either QConvergeConsole GUI or interactive QConvergeConsole CLI, as described in the following sections.

For more detailed information, see the *QConvergeConsole Help* and the *User's Guide—QConvergeConsole CLI* (part number SN0054667-00).

Creating a vPort

To create a vPort in interactive QConvergeConsole CLI:

1. On the **Main Menu**, select **Adapter Configuration**.
2. On the **Adapter Type Selection** menu, select **Virtual Ports (NPIV)**.
3. On the **Virtual Ports** menu, select **Adapter Port**
4. On the **Adapter Port** menu, select **Create vPorts**.
5. On the **Create vPorts** menu, select **Generate**.
6. To create the vPorts, select **Commit**.

Changing the vPort Priority Level

To change the vPort priority level using interactive QConvergeConsole CLI:

1. On the **Main Menu**, select **Adapter Configuration**.
2. On the **Adapter Type Selection** menu, select **QoS**.
3. On the **QoS** menu, select **Adapter Port**.
4. On the **Adapter Port** menu, select **QoS Priority**.
5. On the **QoS Priority** menu, select **Modify QoS Settings**.
6. Choose the priority value for the vPorts.
7. Select **Apply QoS**.

Changing the Priority Levels in VMware ESXi

To change the priority levels using QConvergeConsole VMware vCenter Server Plug-in:

1. In the tree pane on the left, select a host.
2. Expand the host in the left pane by clicking the plus sign (**+**).
3. Locate and expand the 2600 Series Adapter whose priority levels you want to change.
4. Select a port in the 2600 Series Adapter.
5. On the content pane on the right, click **QoS**.

6. Set the vPort priority as follows:

   ❑ If the vPort is listed under **QoS**:

      a. Select the vPort.

      b. Set the QoS priority value for the vPort by moving the corresponding slider.

      c. Click **Save**.

   Figure 3-26 shows an example.

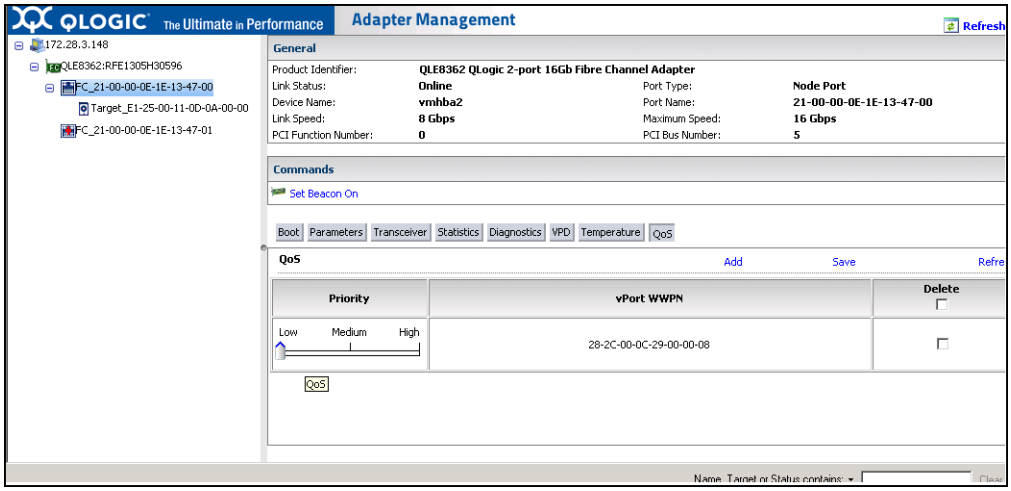

Figure 3-26. Changing vPort Priority Using QConvergeConsole Plug-in

   ❑ If vPort is *not* listed under **QoS**:

      a. Under **QoS**, click **Add**.

      b. Set the QoS priority value by moving the corresponding slider.

      c. Enter the vPort WWPN.

      d. Click **Save**.

7. Click **Refresh** to see the new QoS settings.

To remove a vPort priority:

1. Under **QoS**, locate the vPort whose priority you want to remove.
2. Do one of the following:

   ❑ Select the **Delete** check box of the vPort you want to remove.

   ❑ Select the **Delete** check box in the header to select or clear all of the vPorts.

3. Click **Save**.

Configuring FDMI

The fabric device management interface (FDMI) enables the centralized management of devices such as host adapters. Supported configuration for FDMI:

- Marvell QLogic 2600 Series Adapters
- Brocade switches with Fabric OS (FOS) v7.3 or later
- Cisco switches with NX-OS v8.3.1 or later

NOTE

Results can be viewed from Switch CLI. If the Fibre Channel switch does not support FDMI-2, it falls back to FDMI-1 support.

This section provides the following information on the FDMI:

Brocade Switch FOS CLI

Figure 3-27 shows the FOS CLI of a Brocade switch.

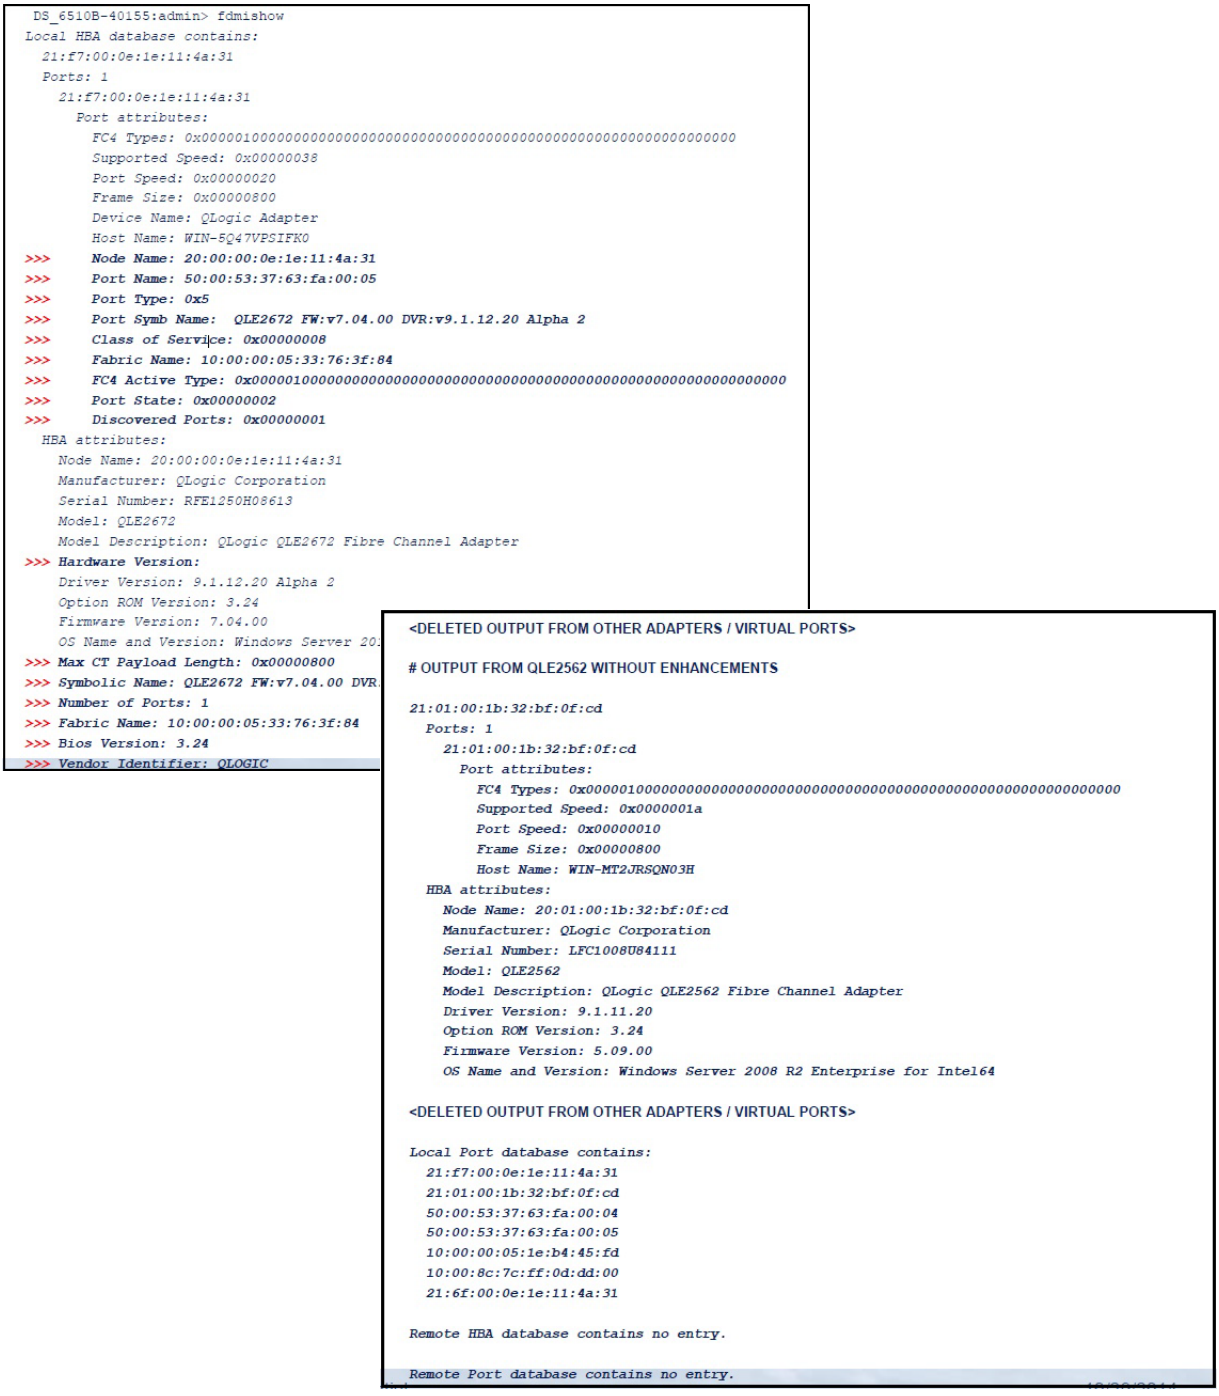

Figure 3-27. FOS CLI of Brocade Switch

Fabric Features

Table 3-9 shows the default configuration for Brocade Fibre Channel fabric features.

Table 3-9. Brocade Fibre Channel Fabric Features—Default Configuration

Feature	Description	Management Tools	Default Configuration
D_Port	Diagnostics port support	N/A	Diagnostic test; run from switch. Nothing to configure on adapter.
FA-WWPN	Fabric-assigned port world wide name	QConvergeConsole CLI QConvergeConsole VMware vCenter Server Plug-in	Disabled by default on adapter. Must be enabled in NVRAM.
F-BLD	Fabric-based boot LUN discovery	QConvergeConsole CLI QConvergeConsole VMware vCenter Server Plug-in	No default setting. Requires configuration on adapter.
FC Ping	Ping a Fibre Channel N_Port or end device	QConvergeConsole CLI QConvergeConsole VMware vCenter Server Plug-in	Diagnostic test; nothing to configure.
FC Trace Route	Obtains the path information between two F_Ports from the Fabric Configuration Server	QConvergeConsole CLI QConvergeConsole VMware vCenter Server Plug-in	Diagnostic test; nothing to configure.
QoS CS_CTL	Prioritize frames depending on the value of the CS_CTL field in the FC Frame header	QConvergeConsole CLI QConvergeConsole VMware vCenter Server Plug-in	By default, QoS of the physical port is set to **High**.
FDMI Enhancements	FDMI attributes	N/A	Nothing to configure on adapter.

Table 3-10 shows the default configuration for Cisco Fibre Channel fabric features.

Table 3-10. Cisco Fibre Channel Fabric Features—Default Configuration

Feature	Description	Management Tools	Default Configuration
D_Port	Diagnostics port support	N/A	Diagnostic test; run from switch. Nothing to configure on adapter.
FC Ping	Ping a Fibre Channel N_Port or end device	QConvergeConsole CLI QConvergeConsole VMware vCenter Server Plug-in	Diagnostic test; nothing to configure.
FDMI Enhancements	FDMI attributes	N/A	Nothing to configure on adapter.

Table 3-11 lists the minimum Dell software version of the Fibre Channel fabric features for both Brocade and Cisco.

Table 3-11. Fibre Channel Fabric Features, Software Version

Component	Dell Software Version
Firmware	9.07.00
Windows Driver	9.4.5.20A3
ESXi Driver	ESX 7.0 v4.1.38 ESX 8.0
Linux Driver	10.02.05.00-a7
QConvergeConsole CLI	2.4.00.21

FDMI Enhancements Support

FDMI enables the management of devices such as host adapters. Enabling additional attributes using FDMI provides you with detailed information about the adapters through a centralized management entity such as the Fibre Channel fabric switches.

To verify the additional adapter attributes registered on the Brocade switch, issue the `fdmishow` command on Brocade Fibre Channel switches to confirm that all the requested attributes are displayed in the switch. For details, see the Brocade switch documentation.

The following shows sample output of the `fdmishow` command.

Port attributes (Brocade switch):

```
"Device Name: QLogic Adapter"
"Node Name: 20:00:00:0e:1e:80:af:80"
"Port Name: 21:00:00:0e:1e:80:af:78"
"Port Type: 0x5"
"Port Symbolic Name: QLE8362 FW:v7.04.00 DVR:v9.1.12.20 Alpha 2"
"Class of Service: 0x00000008"
"Fabric Name: 10:00:00:05:33:8f:1a:49"
"FC4 Active Type: 0x0000010000000000000000000000000000000000000000000000000000000000"
"Port State: 0x00000002"
"Discovered Ports: 0x00000000"
```

HBA attributes:

```
"OS Name and Version: Windows Server 2012 R2 Datacenter Evaluation for Intel64"
"Max CT Payload Length: 0x00000800"
"Symbolic Name: QLE8362 FW:v7.04.00 DVR:v9.1.12.20 Alpha 2"
"Number of Ports: 1"
"Fabric Name: 10:00:00:05:33:8f:1a:43"
"Bios Version: 3.26"
"Vendor Identifier: QLOGIC"
```

The following example shows how to verify FDMI from a Cisco switch. For more details, see the Cisco switch documentation.

```
Cisco# sh fdmi database detail
Registered HBA List for VSAN 1

-------------------------------
HBA-ID: 21:00:00:24:ff:7f:29:39
-------------------------------
Node Name        :20:00:00:24:ff:7f:29:39
Manufacturer     :QLogic Corporation
Serial Num       :RFD1648N67032
Model            :QLE2692
Model Description:QLogic 16Gb 2-port FC to PCIe Gen3 x8 Adapter
Hardware Ver     :BK3210407-01 02
Driver Ver       :10.02.05.00.a3-k
ROM Ver          :3.62
Firmware Ver     :9.06.02 (d0d5)
--More--  Port-id: 21:00:00:24:ff:7f:29:39
```

```
Supported FC4 types:scsi-fcp 40
Supported Speed    :4G 8G 16G
Current Speed      :16G
Maximum Frame Size :2048
OS Device Name     :qla2xxx:host10
Host Name          :linux-ndyy
```

Enabling Marvell QLogic Forward Error Correction

QConvergeConsole supports the forward error correction (FEC) feature, which is a correction technique used for high-speed telecommunications devices to control errors in data transmission over unreliable or noisy communication channels.

NOTE

The FEC feature is supported only if the current adapter port is connected to a Fibre Channel switch with FEC support. See Table 3-11 on page 86 for the supported Fibre Channel software version.

FEC Process Overview

The following highlights how the FEC process works between the transmitting and receiving device:

1. Transmitter Device:
   a. Takes the incoming data stream.
   b. Inserts a redundant amount of parity bits of information on the data stream.
   c. Sends the data stream.
2. Receiving Device:
   a. Receives the data stream.
   b. Checks for missing or corrupt data:
      - If there are no corrupt or missing bits of information, the receiving device proceeds with normal process.
      - If there are corrupt or missing bits of information, the receiving device reconstructs the original data by using pieces of the original data stream (that are still intact) and combining those with the parity bits of information.

NOTE

Instead of having to request retransmission of the entire data frame, the receiving device can recreate or reconstruct the data stream itself. This process prevents multiple CRC errors from being triggered and causing application performance issues.

Enabling and Disabling Marvell QLogic FEC

Use the appropriate procedure to enable or disable FEC:

- Enabling and Disabling FEC on VMware ESXi
- Enabling and Disabling FEC on Linux and Windows

Enabling and Disabling FEC on VMware ESXi

To enable or disable FEC using the QConvergeConsole VMware vCenter Server Plug-in, see the *Marvell QLogic QConvergeConsole Plug-ins for vSphere*, (part number SN0054677-00).

To enable or disable FEC using the QConvergeConsole ESXCLI Plug-in for VMware vSphere (ESXi), see the *User's Guide—Marvell ESXCLI Plug-in for VMware vSphere* (part number BC0151101-00).

Enabling and Disabling FEC on Linux and Windows

To enable FEC on Linux or Windows using QConvergeConsole CLI, follow the instructions in the *User's Guide—QConvergeConsole CLI* (part number SN0054667-00).

To enable FEC on Linux or Windows using Windows PowerShell cmdlets, follow the instructions in the *User's Guide—Marvell QLogic PowerKit* (part number TD-001187).

To enable FEC on Windows using the QConvergeConsole extension for Windows Admin Center, follow the instructions in *User's Guide—Marvell QLogic Fibre Channel QConvergeConsole Extension for Windows Admin Center* ((part number TD-001728).

All guides are available on the Support pages at www.marvell.com

In Figure 3-28, the QConvergeConsole CLI interactive mode **Adapter Information** menu shows a port with an online FEC status.

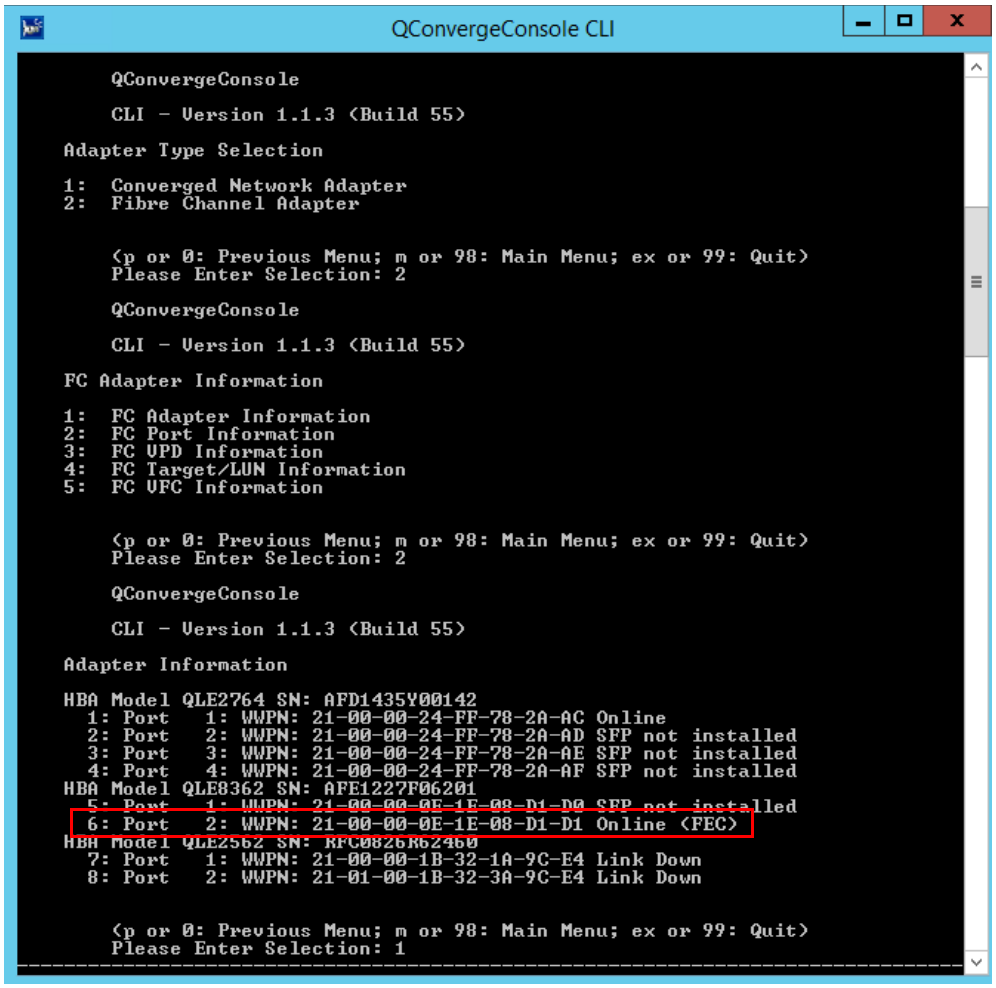

Figure 3-28. QConvergeConsole CLI Adapter Information Menu, FEC Status

Running Extended Link Service Commands

Commands for extended link service on the switch include the following:

- Link Cable Beacon Extended Link Service Command
- Read Diagnostic Parameters Extended Link Service Command

Link Cable Beacon Extended Link Service Command

NOTE

This feature is supported only on QLE2690-DEL and QLE2692-DEL Series Adapters. Run the LCB test only from a switch with either FOS version 7.4.*x* or later, or Cisco NX-OS 8.3.1 or later.

The link cable beacon (LCB) extended link service command locates the corresponding ends of a cable by flashing the specific LEDs on the corresponding ports of a link. The LCB command can also be sent through a Brocade switch using the following commands, which show examples of enabling, disabling, or verifying the LCB:

```
SW6505:admin> portpeerbeacon --enable 12
SW6505:admin> portpeerbeacon --show -all
portPeerBeacon is enabled on port :12
SW6505:admin> portpeerbeacon --disable 12
```

Read Diagnostic Parameters Extended Link Service Command

NOTE

This feature is only supported in QLE2690-DEL and QLE2692-DEL Series Adapters. You must run the RDP test from a switch with firmware as follows:

- Brocade switches: FOS version 7.4.*x* or later
- Cisco switches: NX-OS version 8.2.1 or later

The following example shows the read diagnostic parameters extended link service (RDP ELS) command output for Port 13 on a Brocade 16G switch:

```
SW6505:admin> sfpshow 13 -link -force
Identifier:  3    SFP
Connector:   7    LC
Transceiver: 7004404000000000 4,8,16_Gbps M5 sw Short_dist
Encoding:    6    64B66B
```

```
Baud Rate:   140  (units 100 megabaud)
Length 9u:   0    (units km)
Length 9u:   0    (units 100 meters)
Length 50u (OM2):  3    (units 10 meters)
Length 50u (OM3):  0    (units 10 meters)
Length 62.5u:0    (units 10 meters)
Length Cu:   0    (units 1 meter)
Vendor Name: HP-F     BROCADE
Vendor OUI:  00:05:1e
Vendor PN:   QK724A
Vendor Rev:  A
Wavelength:  850  (units nm)
Options:     003a Loss_of_Sig,Tx_Fault,Tx_Disable
BR Max:      0
BR Min:      0
Serial No:   HAF314370000B7F
Date Code:   140909
DD Type:     0x68
Enh Options: 0xfa
Status/Ctrl: 0x30
Pwr On Time: 0.30 years (2666 hours)
E-Wrap Control: 0
O-Wrap Control: 0
Alarm flags[0,1] = 0x0, 0x0
Warn Flags[0,1] = 0x0, 0x0
                                        Alarm                 Warn
                                  low          high       low          high
Temperature: 60      Centigrade      -5          85         0            75
Current:     8.430   mAmps           2.500       12.000     2.000        11.500
Voltage:     3322.2  mVolts          3000.0      3600.0     3100.0       3500.0
RX Power:    -3.1    dBm (495.3uW)   31.6   uW   1258.9 uW  31.6   uW    794.3  uW
TX Power:    -3.2    dBm (481.6 uW)  125.9  uW   1258.9 uW  251.2  uW    794.3  uW

State transitions: 5
Port Speed Capabilities    4Gbps 8Gbps 16Gbps

PEER Port Gbic Info
Laser Type:     Short Wave Laser
        SFP Type:       Optical Port Type
        Connecter Type: Other
```

```
Following SFP Parameters are Valid
        Temperature: 35      Centigrade [Range -128 - +128 C]
        Current:     8.176   mAmps     [Range 0 - 131 mAmps]
        Voltage:     3463.2  mVolts    [Range 0 - 3600 mVolts]
        Rx Power:    460.4   uW       [Range 0 - 6550 uW]
        Tx Power:    504.7   uW       [Range 0 - 6550 uW]
Signal Loss (Upstream)  :  -16.7   dBm (21.2  uW)
Signal Loss (Downstream):  -20.3   dBm (9.4   uW)
Port Speed Capabilities    4Gbps 8Gbps 16Gbps
Last poll time: 03-24-2015 UTC Tue 18:03:14
```

Universal SAN Congestion Mitigation (USCM)

NOTE

Universal SAN congestion mitigation (USCM) is supported on the QLE2690, QLE2692, QLE2694 and QLE2694L adapters. See “Prerequisites” on page 94 for more information.

- 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.

NOTE

Marvell USCM requires a Brocade or Cisco switch running a supported FOS or NX-OS version.

The USCM feature tracks fabric performance impact notifications (FPINs) and prevents both potential or actual congestion occurrences in the FC SAN environment. An FC port can become congested because of events such as link instability, credit stall, and over-subscription. When this feature is supported and enabled, the drivers and the application track FPINs and maintain statistics about the health of local and remote FC ports in NVMe and FCP environments. For a subset of FPINs that indicate a congestion occurrence in the fabric, depending on the cause of the congestion, the driver uses either one or a combination of multiple actions to mitigate congestion. The mechanisms/actions are described in the following sections.

By default, USCM is enabled on all ports of the adapter. USCM support can be configured (enabled/disabled) per port on all Marvell management tools.

NOTE

In the examples below the “0” denotes the first port detected by QCC CLI on the system.

For example, to enable USCM in QConvergeConsole CLI:

```
# qaucli -scm|-uscm 0 --enable
```

For example, to disable USCM in QConvergeConsole CLI:

```
# qaucli -scm|-uscm 0 --disable
```

- USCM allows you to monitor the following FPIN statistics:
  - Congestion Notification (FPIN-CN)
  - Peer Congestion Notification (FPIN-PN)
  - Link Integrity Notification (FPIN-LI)
  - Delivery Notification (FPIN-DN)
- You can request the following USCM information at any time:
  - A list of slow-drain device(s)
  - A list of normal-drain device(s)

For more information about using USCM, see the appropriate documents in “Related Materials” on page xi

Prerequisites

USCM is supported only on the following adapters:

- QLE2690-DEL
- QLE2690L-DEL
- QLE2690N-DEL
- QLE2690NL-DEL
- QLE2692-DEL
- QLE2692L-DEL
- QLE2692N-DEL
- QLE2692NL-DEL
- QME2692-DEL

The following drivers support USCM. You can download the drivers from the Marvell Web site.

- Linux: 10.02.04.00 or later
- Windows: 9.4.4.20 or later
- ESX: 3.1.64.0/4.1.34.0/5.3.1

Marvell USCM requires a Brocade or Cisco switch running a supported FOS or NX-OS version.

USCM is supported on the following Marvell management tools:

- QConvergeConsole CLI
- Marvell QLogic Fibre Channel QConvergeConsole Extension for Windows Admin Center
- QConvergeConsole Plug-in for VMware vCenter Server
- Marvell QLogic PowerKit
- Marvell ESXCLI Plug-in for VMware vSphere

To use these tools for USCM, see the associated user's guides in "Related Materials" on page xi.

USCM Profiles

USCM Profiles can be controlled:

- Through the Marvell QLogic applications listed in "Prerequisites" on page 94
- Using driver modules or registry parameters

USCM Profiles control the amount of *I/O throttling*, where the driver reduces the load on the fabric by *throttling down* outstanding requests to a target device. The Fabric Performance Impact Notifications (FPINs) provide details about the N_Port where the congestion originates (perpetrator port), along with the cause of the congestion (credit stall, over-subscription, and so on). Based on these details, the driver determines a throttling strategy.

Throttling appropriate requests reduces the load of the port at the congestion origin, which improves the performance of one or more *victim* ports that share downstream resources with the *perpetrator* port. However, reducing the load may reduce the performance of one or more sessions associated with the perpetrator port. While the main goal is improving the performance of victim ports, under certain circumstances, a performance drop at the perpetrator port may be undesirable.

The USCM Profiles feature provides controls by which an administrator can select an appropriate mitigation strategy using one of the Marvell applications listed in "Prerequisites" on page 94 or using a driver module/registry parameter.

Using Driver Modules or Registry Parameters for USCM Profiles

All of parameters in this section are module parameters used in both Linux and VMware ESXi; that is, both use the same name.

`scmrprofile`
 	
`ql2x_scmr_driver_profile`
	Windows/driver registry parameter
	
	Linux and VMware ESXi/driver module parameter
		Requests are throttled based on defined driver profiles. The default profile is Monitor Only.
		USCM Profiles are:
		0 = Monitor Only (default)
		1 = Conservative
		2 = Moderate
		3 = Aggressive

For more information, see the associated readme.txt file.

Configuring USCM Profiles

The USCM profile allows you to manage adapter congestion by selecting either the driver default profile settings, or customizing different profile settings for each adapter port on a case-by-case basis. Table 3-12 lists the available profiles.

Table 3-12. Adapter Port USCM Profiles

Profile	Description
`Monitor Only` (default)	Records adapter performance and congestion history for review. No actions are taken to resolve congestion.
`Conservative`	Maintains optimum throughput while gradually reducing congestion. Queue depth (outstanding I/Os) is reduced to half of the current value as part of the throttle down operation. Marvell recommends this setting for high-priority workloads.
`Moderate`	Queue depth (outstanding I/Os) is reduced to one-quarter of the current value as part of the throttle down operation.
`Aggressive`	Reduces congestion on priority while reducing throughput. Queue depth (outstanding I/Os) is reduced to one-eighth of the current value as part of the throttle down operation. Marvell recommends this setting for low-priority workloads.

Table 3-12. Adapter Port USCM Profiles (Continued)

Profile	Description
`Revert to Driver Settings`	Reverts the current, active profile to the driver profile default setting. The command takes effect immediately; a system reboot or driver reload is not required.

CLI

See the *User's Guide: QConvergeConsole CLI for instructions* on setting the USCM profiles using CLI.

Linux

To set the USCM profile in Linux, do one of the following:

- Set the `ql2x_scmr_profile` driver module parameter:

```
int ql2x_scmr_profile = 1;
module_param(ql2x_scmr_profile, int, 0600);
MODULE_PARM_DESC(ql2x_scmr_profile,
     " Defines SCM profile."
     "0 - Monitor only"
     "\t\t1 - Conservative(default)"
     "\t\t2 - Moderate "
     "\t\t3 - Aggressive \n");
```

- Edit the `/etc/modprobe.d/scsi.conf` file and add the following line:

```
options qla2xxx int ql2x_scmr_profile = 1
```

VMware ESXi

To set USCM profile in VMware ESXi, do one of the following:

- Set the `ql2x_scmr_driver_profile` driver module parameter. For example:

```
ql2x_scmr_driver_profile = 1;
VMK_MODPARAM(ql2x_scmr_driver_profile, int,
    " \n Driver assigned USCM profile - used only when QCC CLI
has not programmed a profile"
     "\t\t 0 = Monitor (Default)\n"
     "\t\t 1 = Conservative\n"
     "\t\t 2 = Moderate\n"
     "\t\t 3 = Aggressive\n");
```

- Set the `ql2x_scmr_driver_profile` parameter by issuing the following command:

  ```
  $ esxcfg-module -s "ql2x_scmr_driver_profile=1" qlnativefc
  ```

Windows

To set the USCM/SCM Profile in Windows, follow the instructions in the Windows driver readme.txt file and reference ‘scmprofile.’

FPIN Link Integrity Marginal Path Detection Support

The Device Mapper (DM) Multipath FPIN link integrity (FPIN-LI) marginal path detection is currently supported on RHEL 8.4, RHEL 8.5, RHEL 8.6, RHEL 9.0, SLES15 SP3 and later OSs with Brocade and Cisco fabrics. See the appropriate documentation for more information on the FPIN functionality provided by the Fibre Channel switch. The Fibre Channel Transport Services daemon (`fctxpd`) must be installed on the RHEL 8.*x* OS for marginal path detection to work. The transport daemon improves DM Multipath functionality to take corrective actions when link integrity events are received from switch. The steps are provided in “FC Transport daemon Installation for RHEL 8.4, 8.5, and 8.6” on page 98.

RHEL 9.0 and SLES 15 SP3 and later support marginal path detection natively. Therefore, you do not have to install the Fibre Channel transport services daemon on these OSs. However, changes are needed in the `multipath.config` file for marginal path detection to work correctly. The steps are provided in “Multipath Configuration File Changes for RHEL 9.0 and SLES 15 SP3 and Later” on page 100.

See the Brocade FOS and Cisco NX-OS documentation for more information on the fabric notification functionality provided by the switch.

FC Transport daemon Installation for RHEL 8.4, 8.5, and 8.6

The current Red Hat offering requires an out-of-box solution that requires the installation of an Extra Packages for Enterprise Linux (EPEL) component of the distribution. Following are the steps to install the FC transport daemon on the system when using RHEL 8.4 and 8.5.

1. Install the EPEL package:
   a. `dnf install https://dl.fedoraproject.org/pub/epel/epel-release-latest-8.noarch.rpm`
   b. `dnf update`
   c. `dnf install fctxpd.x86_64`

2. Start the FC Transport daemon.

   a. Run `systemctl status fctxpd`

      Following is a sample output.

      ```
      fctxpd.service - Fiber Channel Transport Daemon
      Loaded: loaded (/usr/lib/systemd/system/fctxpd.service;
      disabled; vendor preset: disabled)
      Active: inactive (dead)
      ```

      NOTE

      The fctxpd deamon is not started by default when installed and needs to be started manually.

   b. Run `systemctl enable fctxpd`

      NOTE

      This step is to ensure that the daemon is loaded automatically upon system reboot.

   c. Run `systemctl start fctxpd` to start the daemon.

3. Verify the FC Transport daemon status by running `systemctl status fctxpd`.

   Following is a sample output indicating that the daemon started successfully.

   ```
   systemctl status fctxpd

   â fctxpd.service - Fiber Channel Transport Daemon

   Loaded: loaded
   (/usr/lib/systemd/system/fctxpd.service; disabled;
   vendor preset: disabled)

   Active: active (running) since Mon 2020-07-27
   17:43:11 IST; 2s ago

   Main PID: 64691 (fctxpd)

   Tasks: 2 (limit: 403798)

   Memory: 688.0K
   ```

The marginal path detection package installation is complete.

Multipath Configuration File Changes for RHEL 9.0 and SLES 15 SP3 and Later

Edit the `multipath.conf` file as follows:

Enable the attribute `marginal_pathgroups` in the `/etc/multipath.conf` file. For example:

```
marginal_pathgroups fpinx3
```

If the `marginal_pathgroups` attribute is not enabled, marginal path detection will not work for the Fabric Notification event.

4 NVM Express over Fibre Channel

This chapter describes the installation and configuration of NVMe over Fibre Channel (FC-NVMe) on the Marvell QLogic adapters. It covers:

- Hardware and Software Requirements
- “Host Setup” on page 102
- “Boot From SAN Setup” on page 109
- “Limitations and Known Issues” on page 109
- “Unloading and Reloading the Linux Driver” on page 108

For assistance with the installation and configuration of FC-NVMe, contact Marvell Technical Support at this address:

https://www.marvell.com/support/contact/cust_support.do

Hardware and Software Requirements

The following sections list the required adapters, operating systems, and host software.

See the latest Readme and Release Notes for required versions of:

- Drivers
- Firmware
- QConvergeConsole CLI
- QConvergeConsole Plug-in for VMware vCenter Server
- Marvell QLogic Fibre Channel QConvergeConsole Extension for Windows Admin Center
- Marvell QLogic PowerKit
- Marvell ESXCLI Plug-in for VMware vSphere

Adapters

FC-NVMe requires one or more of the QLE269*x*-DEL or QME2692-DEL adapters.

Operating Systems

The following x86_64 operating systems are supported for FC-NVMe.

- Operating systems with inbox drivers:
  - Windows Server 2022 or later
  - SLES 15 SP2, or later
  - RHEL 7.8, or later
  - RHEL 8.0, or later
  - ESXi 7.0 U1, or later
- Operating systems with out-of-box drivers:
  - SLES 15 SP2, or later
  - RHEL 7.8, or later
  - RHEL 8.3, or later
  - Windows Server 2019, or later
  - ESXi 7.0, or later

Host Setup

Procedures for setting up the initiator host include the following:

- Installing the Drivers
- Finding Host NQN
- Programming the Host NQN onto the NVMe Storage Array
- Configuring System for NVMe Auto Discovery
- Verifying the Availability of FC-NVMe Devices
- Unloading and Reloading the Linux Driver
- Limitations and Known Issues

Installing the Drivers

The following instructions are for configuring Linux, Windows, and ESX Host for FC-NVMe. Installation of the OS driver is a prerequisite for configuring FC-NVMe. For instructions on installing the driver for each OS, see “Driver Installation and Configuration” on page 5.

Finding Host NQN

The Host NQN string is needed to configure and connect to NVMe Storage. The default Host NQN is displayed at the HII UEFI driver level. See “NVME options” on page 63.

Additionally, the Host NQN can also be displayed at the OS level using the methods described in the following sections. These methods use commands from either the operating system or QConvergeConsole CLI.

Linux

1. With the NVMe CLI already installed, issue the following command to verify that the hostnqn file exists:

   ```
   # ls /etc/nvme/hostnqn
   /etc/nvme/hostnqn
   ```

2. If the hostnqn file exists, issue the following command to view and note the hostnqn data to use when mapping the host initiator(s) to storage subsystems:

   ```
   # cat /etc/nvme/hostnqn
   nqn.2014-08.org.nvmexpress:uuid:1164eabb-e77f-4f15-ad16-c
   efb43d1fc9
   ```

3. If the hostnqn file does not exist, issue the following command to create it using the NVMe CLI:

   ```
   # nvme gen-hostnqn
   nqn.2014-08.org.nvmexpress:uuid:a0baf242-85fd-41e5-a524-e
   253d8531282
   ```

4. Take the output from the nvme gen-hostnqn command and save it to the `hostnqn` file. The edited file in the /etc/nvme/hostnqn directory should appear as follows.

   ```
   nqn.2014-08.org.nvmexpress:uuid:a0baf242-85fd-41e5-a524-e
   253d8531282
   ```

   If steps 3 and 4 are performed, a new dynamic Host NQN is created that may not match the default Host NQN displayed in the HII UEFI. Ensure the newly created Host NQN is configured into the NVMe storage, or you can edit the `hostnqn` file to match the Host NQN displayed as the default Host NQN at the HII UEFI.

Windows

The driver automatically generates a host NVMe qualified name (NQN). To view the host NQN, issue the `qaucli -i` command.

The following shows an example Host NQN output:

```
-------------------------------------------------------------------------------
Host Name : WIN-4KHVTBVO147
Host NQN : nqn.2018-08.com.marvell:nvme.host.sys.wwpn:21000024ff784720
HBA Instance : 2
HBA Model : QLE2692
```

NOTE

Host NQN is read-only and is generated by the driver. There is no option to modify/configure Host NQN from QConvergeConsole CLI.

ESXi

Issue the following command to list the current Host NQN:

```
# localcli nvme info get
Host NQN: nqn.2014-08.org.qlogic:nvme:localhost
```

Programming the Host NQN onto the NVMe Storage Array

Refer to the instructions provided by the vendor of your NVMe Storage Array to configure and map the Initiator Host NQN to NVMe namespaces.

Configuring System for NVMe Auto Discovery

Auto discovery is natively supported for SLES 15 SP2 or later, RHEL 8.3 or later, Windows Server 2019 or later, and ESXi 7.0 or later. No additional configuration is needed.

The following procedure is for configuring auto discovery on RHEL 7.8 or later only.

To set up a Linux host:

1. Verify that the FC-NVMe scripts are installed by issuing the following commands:

```
# ls -l /etc/udev/rules.d/99-nvme-fc.rules
-rw-r--r--. 1 root root 310 Dec 31 15:14
/etc/udev/rules.d/99-nvme-fc.rules
# ls -l /usr/sbin/qla2xxx/qla2xxx_nvme_boot_connections.sh
-rwxr-xr-x. 1 root root 859 Jan 31 10:56
/usr/sbin/qla2xxx/qla2xxx_nvme_boot_connections.sh
# ls -l
/usr/lib/systemd/system/qla2xxx-nvmefc-connect@.service
-rwxr-xr-x. 1 root root 859 Jan 31 10:56
/usr/lib/systemd/system/qla2xxx-nvmefc-connect@.service
# ls -l
/usr/lib/systemd/system/qla2xxx-nvmefc-connect@.service
```

```
-rw-r--r--. 1 root root 279 Jan 31 10:54
/usr/lib/systemd/system/qla2xxx-nvmefc-connect@.service
```

2. If the Linux out-of-box driver does not discover NVMe targets after a system reboot, perform the following steps to ensure the service namedqla2xxx-nvmefc-boot-connection.service is enabled.

   a. Open a Linux terminal window and issue the following command to check the state of the service.

   ```
   # systemctl list-unit-files | grep nvmefc-boot-connection
   qla2xxx-nvmefc-boot-connection.service disabled
   ```

   b. If the service is disabled, issue the following command to enable.

   ```
   # systemctl enable qla2xxx-nvmefc-boot-connection.service
   ```

   Following is an example response that should be seen:

   ```
   Created symlink
   /etc/systemd/system/default.target.wants/qla2xxx-nvmefc-b
   oot-connection.service
   →/usr/lib/systemd/system/qla2xxx-nvmefc-boot-connection.
   service.
   ```

   c. Issue the following command to verify that the service is enabled:

   ```
   # systemctl list-unit-files | grep nvmefc-boot-connection
   qla2xxx-nvmefc-boot-connection.service enabled
   ```

3. Reboot the system and verify that the NVMe targets are discovered.

Verifying the Availability of FC-NVMe Devices

Follow the procedures in this section to confirm FC-NVMe device availability on Linux and Windows systems.

Verifying in Linux

In Linux, verify that the target devices are available by issuing either of the following commands:

```
# lsblk
```

or

```
# nvme list
```

Example: `lsblk` Command

The following example of a portion of the `lsblk` output shows the FC-NVMe attached targets. The system discovered three target controllers (nvme0, nvme1, and nvme2). Each of these controllers has one FC-NVMe LUN behind it (n1).

```
NAME        MAJ:MIN RM   SIZE RO TYPE MOUNTPOINT

nvme0n1     259:2    0     1G  0 disk
nvme0n1p1   259:25   0  1023M  0 part
nvme0n2     259:8    0     1G  0 disk
nvme0n3     259:12   0     1G  0 disk
nvme0n4     259:16   0     1G  0 disk
nvme0n5     259:18   0     1G  0 disk
nvme0n6     259:20   0     1G  0 disk
nvme0n7     259:22   0     1G  0 disk
nvme0n8     259:24   0     1G  0 disk
```

Example: `nvme list` Command

The following example of a portion of the `nvme list` output shows the FC-NVMe attached targets.

```
Node     SN                   Model    Namespace Usage Format FW Rev
-------  -------------------- -------- ---- ---- ----------------------- --------- -----------
nvme0n1 00-200100110dd89300 SANBlaze VLUN P3T0 1 67.11  MB /  67.11  MB 512   B +  0 B   V8.x-dev
nvme1n1 00-200000110dd89200 SANBlaze VLUN P2T0 1 67.11  MB /  67.11  MB 512   B +  0 B   V8.x-dev
nvme2n1 00-200000110d4d1300 SANBlaze VLUN P0T0 1 67.11  MB /  67.11  MB 512   B +  0 B   V8.x-dev
```

Verifying in Windows

In Windows, verify that the target devices are available by issuing the `qaucli -t` command. For example:

`qaucli -t 0` (HBA instance 0)

```
Using config file: C:\Users\Administrator\qaucli.cfg
Installation directory: C:\Users\Administrator
Working dir: C:\Users\Administrator
------------------------------------------------------------------------------
HBA Instance 0: QLE2692 Port 1 WWPN 21:00:00:24:ff:8f:e3:9a PortID 01-21-00
Link: Online (FEC)
------------------------------------------------------------------------------
Path                            : 0
Target                          : 0
Device ID                       : 0x83
```

```
Product Vendor                  : NVMe
Product ID                      : SANBlaze VLUN P2
Product Revision                : -dev
Serial Number                   : 6001_10d0_0247_2600_0302_0000_929d_14c6.
Node Name                       : 20:02:00:11:0d:47:26:00
Port Name                       : 20:02:00:11:0d:47:26:00
Port ID                         : 01:1b:00
Product Type                    : NVME Disk (Ensure that this shows as an FC-NVMe target.)
LUN Count(s)                    : 51
Status                          : Online
------------------------------------------------------------
```

Figure 4-1 shows the devices in Windows Device Manager.

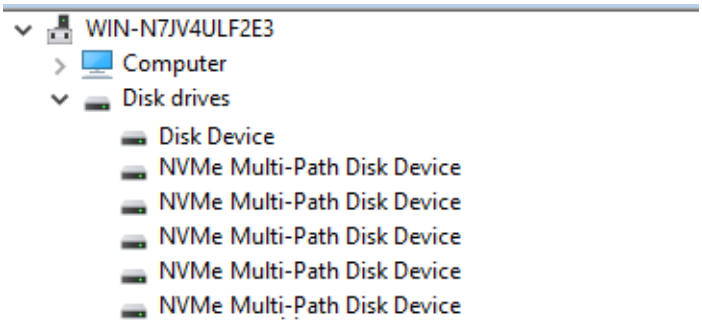

Figure 4-1. FC-NVMe Devices in Windows Device Manager

Verifying in ESX

To verify the controllers using ESX 7.0, issue the following command:

```
$ esxcli nvme controller list
```

Name	Controller Number	Adapter	Transport Type	Is Online
nqn.2014-08.org.sanblaze:virtualun.virtuallunnvme.2.0#vmhba67#200200110d27b200:200200110d27b200	260	vmhba67	FC	true
nqn.2014-08.com.sanblaze:virtualun.virtuallunnvme.2.2#vmhba67#200200110d27b202:200200110d27b202	267	vmhba67	FC	true
nqn.2014-08.com.sanblaze:virtualun.virtuallunnvme.2.4#vmhba67#200200110d27b204:200200110d27b204	268	vmhba67	FC	true
nqn.2014-08.com.sanblaze:virtualun.virtuallunnvme.2.6#vmhba67#200200110d27b206:200200110d27b206	269	vmhba67	FC	true
nqn.2014-08.org.sanblaze:virtualun.virtuallunnvme.2.0#vmhba68#200200110d27b200:200200110d27b200	274	vmhba68	FC	true
nqn.2014-08.com.sanblaze:virtualun.virtuallunnvme.2.2#vmhba68#200200110d27b202:200200110d27b202	281	vmhba68	FC	true
nqn.2014-08.com.sanblaze:virtualun.virtuallunnvme.2.4#vmhba68#200200110d27b204:200200110d27b204	282	vmhba68	FC	true
nqn.2014-08.com.sanblaze:virtualun.virtuallunnvme.2.6#vmhba68#200200110d27b206:200200110d27b206	283	vmhba68	FC	true

Figure 4-2. Listed NVMe Controllers

Display the NVMe namespace using the following command:

```
$ esxcli nvme namespace list
```

Name	Controller Number	Namespace ID	Block Size	Capacity in MB
eui.600110d00227b20003020000929d14c6	260	1	512	4096
eui.600110d00227b202030200022acfba0c	267	1	512	1
eui.600110d00227b20403020004b6297c9a	268	1	512	1
eui.600110d00227b20603020000641833f28	269	1	512	1
eui.600110d00227b20003020000929d14c6	274	1	512	4096
eui.600110d00227b202030200022acfba0c	281	1	512	1
eui.600110d00227b20403020004b6297c9a	282	1	512	1
eui.600110d00227b20603020000641833f28	283	1	512	1

Figure 4-3. NVMe Namespace List

NOTE

The ESX 7.0/8.0 `qlnativefc` driver may create and register two logical devices, one for NVMe and one for FCP. The vmhbas do not represent individual physical FC ports, but a logical device with which the operating system can interact. The WWPN/WWNN shows that it is the same physical FC port.

Unloading and Reloading the Linux Driver

This section applies to a locally booted Linux operating system.

Follow the instructions in this section to unload and reload the Linux driver after first disconnecting the NVMe devices.

To unload and reload the Linux driver:

1. List all the FC-NVMe connected target controllers by issuing the following command:

   ```
   # ls /dev/nvme* | grep -E nvme[0-9]+$
   ```

   The preceding command should list all connected `/dev/nvme[x]` devices. For example:

   ```
   /dev/nvme0
   /dev/nvme1
   ```

2. Disconnect all of the FC-NVMe target devices by issuing the following commands to each of the `/dev/nvme[x]` devices listed in Step 1:

   ```
   # nvme disconnect -d /dev/nvme0
   # nvme disconnect -d /dev/nvme1
   ```

3. Unload the current driver by issuing the following command:

```
# modprobe -r  qla2xxx
```

4. Reload the driver to auto-discover the FC-NVMe subsystems by issuing the following command:

```
# modprobe -v qla2xxx
```

Limitations and Known Issues

The following Linux and Windows limitations and known issues affect the installation and configuration of FC-NVMe.

Linux

For Linux operating systems, Marvell recommends that you disconnect all FC-NVMe devices prior to unloading the driver. Refer to the procedure in Unloading and Reloading the Linux Driver.

Windows

Windows FC-NVMe does not support point-to-point (also called N2N) direct attach to target/LUN configuration. It does support fabric configuration through FC Switch fabric.

Boot From SAN Setup

NVMe BFS is supported through UEFI and not supported through legacy BIOS.

NOTE

NVMe BFS is only supported on SLES 15 SP3 and later, Windows Server 2019, Windows Server 2022, and ESXi 7.0 or later. For instructions on configuring NVMe BFS refer to “Configuring the UEFI Driver” on page 52.

NOTE

Custom NQN configured in HII for NVMe BFS with SLES and Windows is not supported for SLES 15 SP3, SLES 15 SP4, Windows Server 2019, or Windows Server 2022. The default Host NQN string is required for NVMe BFS installation to succeed using the above listed operating systems.

Once the UEFI driver is configured using HII, and the NVMe storage has been configured with the Initiator Host NQN, there are no additional NVMe BFS steps required to detect the NVMe storage in order to install SLES 15 SP3, SLES 15 SP4, Windows Server 2019, or Windows Server 2022.

For ESX 7.0 and later, it is recommended to change the default Host NQN in HII to match the OS default Host NQN using the procedures in the following section.

Additional Setup for NVME Boot-from-SAN Using ESX 7.0 and Later

For VMware Host NQN is derived from the fully qualified domain name (FQDN) of the host. Following is the format for host NQN under VMware ESXi 7.0 and later:

`nqn.2014-08.<reverse-domain>:nvme:<hostname>`

For example, a host that has FQDN as R740-FRSCO.qlogic.org has the following NQN name:

`nqn.2014-08.org.qlogic:nvme:R740-FRSCO`

The UEFI driver assigns the host NQN based on a system UUID. The host NQN field can be changed on the HBA configuration page to match the format used by ESX. Marvell recommends that the UEFI and OS host NQN be the same, or configure the storage to accept both UEFI and OS NQN. See “NVME options” on page 63 for configuration options.

NOTE

To avoid FC-NVMe discovery related issues during the installation time, its important to make sure that the Host NQN in UEFI, storage host definition and ESXi installer match up.

If the fully qualified domain name is not already pushed out by the DNS server, it can be set during the installation by performing the following steps prior to ESXi installation. Once the VMware installer is fully loaded, access the console by pressing Alt+F1.

1. Enter `root` for the username and leave password field blank.
2. Issue the following command to list the current NQN:

   ```
   # localcli nvme info get
   Host NQN: nqn.2014-08.org.qlogic:nvme:localhost
   ```

   Use the host NQN and assign that string into the NVME storage Controller namespace. If necessary, modify the UEFI driver fields to use the host NQN and the storage NQN.

   Changing the host FQDN at a later time changes the NQN in the OS. To avoid this scenario, configure FQDN during the installation process using the following commands:

   ```
   # localcli system hostname set --fqdn=R740-FRSCO.qlogic.org
   ```

3. Check the system FQDN using the following command:

   # localcli system hostname get

4. Force a Host NQN change based on the FQDN that was just set using the following command:

   #localcli nvme info set --hostnqn=default

5. Query the NQN name in the system using the following command:

   # localcli nvme info get

6. Retrieve a list of vmhbas by issuing the following command:

   localcli storage san fc list

   Following is a partial sample output:

```
FcDevice:
   Adapter: vmhba65
   Port ID: 000000
   Node Name: 20:00:f4:c7:aa:0a:d7:27
   Port Name: 21:00:f4:c7:aa:0a:d7:27
   Speed: 0 Gbps
   Port Type: LPort
   Port State: LINK DOWN
   Model Description: QLogic QLE2872 Dual Port 64/32/16GFC
                      PCIe Gen4 x8 Adapter
   Hardware Version:
   OptionROM Version: 0.00
   Firmware Version: 9.10.01 (d0d5)
   Driver Name: qlnativefc
Error getting field DriverVersion
```

7. Reset the HBA and rescan using the following command:

   # localcli storage san fc reset -A vmhbaX

   # esxcfg-rescan vmhbaX

   A few attempts of reset and rescan may be required to force controller discovery.

8. List the namespaces using the following command:

   # localcli nvme namespace list

9. List the NVMe controllers using the following command:

   # localcli nvme controller list

NVMe namespaces should be available at this point to perform the installation. Press Alt+F2 to return to the installer GUI and continue with the installation process.

NOTE

Direct attached or node-to-node (N2N) NVMe BFS is not currently supported.

A Troubleshooting

This appendix provides information about Fibre Channel diagnostics and a Fibre Channel troubleshooting diagram, including:

Fibre Channel Diagnostics

Troubleshooting with Fibre Channel diagnostics include:

Fibre Channel Diagnostics Using Interactive QConvergeConsole CLI

To run the loopback test, read/write buffer test, or beacon test:

1. On the **Main Menu**, select **4: Adapter Diagnostics**.
2. On the **FC Diagnostics** menu, select the test type to run:

   Select **1: Loopback Test** to run the loopback test.

   Select **2: Read Write Buffer Test** to run the read/write buffer test.

   Select **3: Transceiver Diagnostics Monitoring Interface (DMI)** to view the transceiver diagnostics information.

   Select **4: Port Beacon** to run the beacon test.

For more information, refer to the section about the Fibre Channel interactive commands in the *User’s Guide—QConvergeConsole CLI* (part number SN0054667-00).

Fibre Channel Diagnostics Using Noninteractive QConvergeConsole CLI

For information, refer to the section about the Fibre Channel noninteractive commands in the *User's Guide—QConvergeConsole CLI* (part number SN0054667-00).

Fibre Channel Troubleshooting Diagram

Figure A-1 and Figure A-2 provide a two-part troubleshooting flowchart.

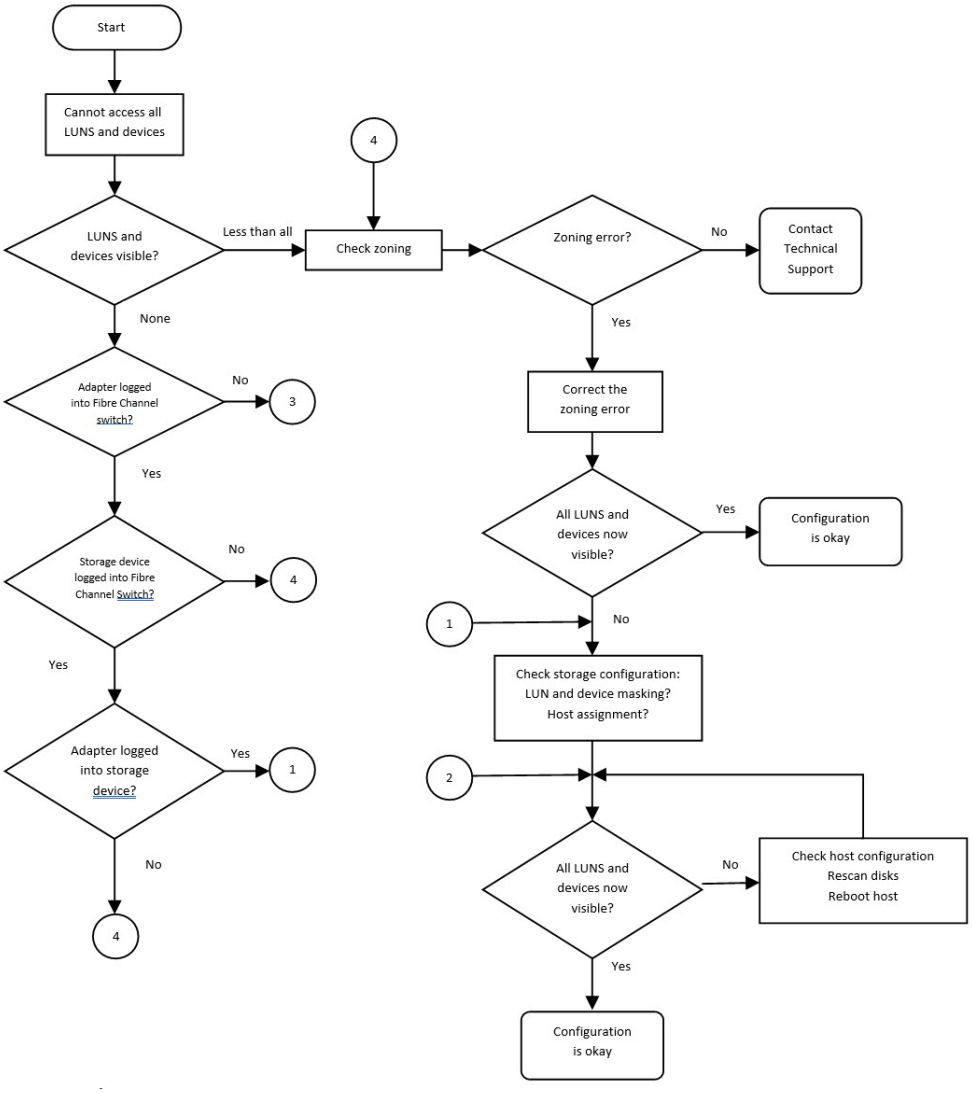

Figure A-1. Fibre Channel Troubleshooting Diagram (1 of 2)

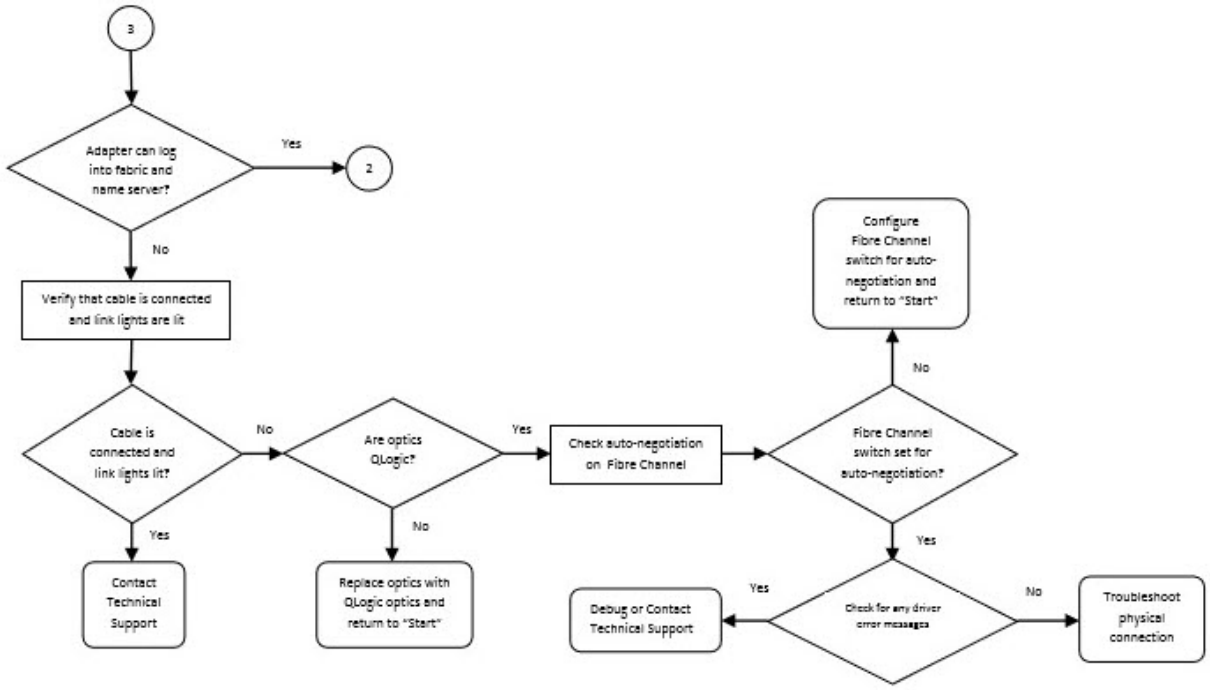

Figure A-2. Fibre Channel Troubleshooting Diagram (2 of 2)

Troubleshooting with a Diagnostic Port (D_Port)

Use the D_Port (diagnostic port) feature on the QLogic 2600 Series Adapters to identify and isolate link failures resulting from faulty modules (link, cable, or SFP) by running diagnostics over a Fibre Channel link.

- The D_Port feature is supported on Windows, Linux, and ESXi.
- The D_Port feature is not supported on the QME2692-DEL adapter
- D_Port can be run only in *dynamic* mode; it does not support *static* mode.
- D_Port requires a Fibre Channel Adapter that is linked to a Brocade or Cisco Fibre Channel 16G switch, and the supported Brocade or Cisco 16G Fibre Channel switch firmware.

Configure the diagnostic D_Port on the switch to dynamically allow a port to enter D_Port mode when a D_Port is configured on the switch end, and then return to operational mode when the test is completed. Also configure the switch to start and stop the diagnostic test on the adapter port. See the associated switch documentation for details.

NOTE

All Marvell QLogic-branded 16GFC SFP+ transceivers can run D_Port tests. If you have a supported non-QLogic branded transceiver, ensure that it supports the optical wrap (OWrap) and electrical wrap (EWrap) capability that is required for D_Port.

Before running diagnostics, verify that the FC switch supports diagnostics mode, and then enable the HBA Diagnostic option. The following example enables the HBA Diagnostic option, where `0` specifies the adapter instance number:

```
# qaucli -dport 0 --info
QLE2672
-----------
=======
Port 1:
=======
DPort  : Enable
```

When a switch port is enabled as a D_Port, the Marvell QLogic 2600 Series Adapter automatically runs the basic diagnostic test at both ends of the link. The Brocade switch then runs the following additional tests:

- Electric loopback test
- Optical loopback test

You should run diagnostic port tests on multipath configured boot from SAN or fabric-assigned boot LUN. Running diagnostic port tests on single path boot from SAN or fabric-assigned boot LUN may cause issues with OS stability in true boot from SAN environments.

The following sections describe how to configure and run D_Port mode, and how to verify D_Port mode from the QConvergeConsole utilities:

- Configuring D_Port on a Brocade Fibre Channel 16G Switch
- Configuring and Running Diagnostics Port on a Cisco Fibre Channel 16G Switch
- Checking D_Port Results from a Brocade Fibre Channel 16G Switch
- Verifying D_Port Mode with Interactive QConvergeConsole CLI
- Verifying D_Port Mode with Noninteractive QConvergeConsole CLI
- Verifying D_Port Mode with QConvergeConsole Plug-ins

Configuring D_Port on a Brocade Fibre Channel 16G Switch

The following example configures and starts a diagnostic test on port 17 of a Brocade Fibre Channel 16G switch (FOS 7.4.*x* or later). For more detail about the switch, see the Brocade switch documentation.

To enable D_Port diagnostic mode:

From the Brocade Fibre Channel 16G switch, issue the following commands:

```
# portdisable <port#>
# portcfgdport --enable <port#>
# portenable <port#>
```

The D_Port test starts automatically.

To view the diagnostics port test results from the Fibre Channel 16/32G switch:

1. To obtain the port number, issue the `switchshow` command as follows:

   ```
   # switchshow
   ```

   The following example output identifies port 17:

   ```
   17  17  011100  id  N16  Online  FC  D-Port Loopback->Port 17
   ```

2. Issue the `portdporttest` command as follows:

   ```
   # portdporttest --show <port#>
   ```

   For example:

```
# portdporttest --show 17
portdporttest --show 17
D-Port Information:
===================
Port:                           17
Remote WWPN:                    21:00:00:0e:1e:11:4a:31
Mode:                           Automatic
No. of test frames:             1 Million
Test frame size:                1024 Bytes
FEC (enabled/option/active):    Yes/No/No
CR (enabled/option/active):     Yes/No/No
Start time:                     Thu Aug 14 01:45:37 2014
End time:                       Thu Aug 14 01:47:40 2014
Status:                         PASSED
=============================================================================
Test                    Start time      Result          EST(HH:MM:SS)   Comments
```

```
================================================================================
Electrical loopback     01:46:08        PASSED          --------   ----------
Optical loopback        01:46:54        PASSED          --------   ----------
Link traffic test       --------        SKIPPED         --------   ----------
================================================================================
Roundtrip link latency:          157 nano-seconds
Estimated cable distance:        3 meters
Buffers required:                1 (for 2112 byte frames at 16Gbps speed)
```

To disable D_Port diagnostic mode:

From the Brocade Fibre Channel 16G switch, issue the following commands:

```
# portdisable <port_num>
# portcfgdport --disable <port_num>
# portenable <port_num>
```

The D_Port test stops.

NOTE

After the diagnostics port is disabled on the switch port, the adapter port automatically resumes normal operation as an N_Port.

Configuring and Running Diagnostics Port on a Cisco Fibre Channel 16G Switch

The following example configures and starts a diagnostic test on port 2 of a Cisco Fibre Channel 9396S or 97*xx* switch (NX-OS 8.2.1 or later). For more details about the switch, see the Cisco switch documentation.

```
Cisco# config t
Cisco# int fc 1/2
Cisco(config-if)# shut
Cisco(config-if)# switchport link-diag
Cisco(config-if)# no shut
Cisco(config-if)# diagnostic start interface fc 1/2 test link-diag
```

To check the diagnostic test status, type the following command:

```
Cisco(config-if)# show diagnostic result interface fc1/1 test
link-diag
WWPN of peer port: 21:00:00:24:ff:8f:c9:e0
Status: Supported (Reflector)
Reflector loopback capabilities:  Xcvr-optical Electrical
```

```
                Time of Test: Mon Feb  5 16:31:02 2018
                Total time taken: 20 seconds
==================|===========|===========|=========|============================|=============
                  |           |           | Discards |        Latency (ns)        |
Loopback Level    | Tx Frames | Rx Frames | IN | OUT |BAD WORDS|In-Switch|External|    Status
==================|===========|===========|====|=====|=========|=========|========|=============
Remote-Switched(R) |          0|          0|   0|    0|        0|        0|       0|       -NA-
Mac(R)             |          0|          0|   0|    0|        0|        0|       0|       -NA-
Xcvr-optical(R)    |    7225275|    7225275|   0|    0|        0|     2112|    1064|    Success
Electrical(R)      |      20000|      20000|        -NA-        |   -NA-  |  -NA-  |    Success
=============================================================================================

Overall Status                          : Success
Cable Length (approx. +/- 5 metres)     : 10.4 metres
```

To disable the diagnostic port, type the following commands:

```
Cisco#int fc 1/2
Cisco(config-if)# shut
Cisco(config-if)# no switchport link-diag
Cisco(config-if)# no shut
```

NOTE

After the diagnostics port is disabled on the switch port, the adapter port automatically resumes normal operation as an N_Port.

Checking D_Port Results from a Brocade Fibre Channel 16G Switch

To view the D_Port test results from a Brocade Fibre Channel 16G switch:

1. To get the port number, issue the `switchshow` command:

   ```
   # switchshow
   ```

   The following example output identifies port 17:

   ```
   17  17  011100  id  N16  Online  FC  D-Port Loopback->Port 17
   ```

2. Issue the `portdporttest` command as follows:

   ```
   # portdporttest --show <port_num>
   ```

For example:

```
# portdporttest --show 17
```

Following is an example output for this command.

```
portdporttest --show 17
D-Port Information:
===================
Port:                           17
Remote WWPN:                    21:00:00:0e:1e:11:4a:31
Mode:                           Automatic
No. of test frames:             1 Million
Test frame size:                1024 Bytes
FEC (enabled/option/active):    Yes/No/No
CR (enabled/option/active):     Yes/No/No
Start time:                     Thu Aug 14 01:45:37 2014
End time:                       Thu Aug 14 01:47:40 2014
Status:                         PASSED
============================================================================
Test                    Start time      Result          EST(HH:MM:SS)   Comments
============================================================================
Electrical loopback     01:46:08        PASSED          --------   ----------
Optical loopback        01:46:54        PASSED          --------   ----------
Link traffic test       --------        SKIPPED         --------   ----------
============================================================================
Roundtrip link latency:         157 nano-seconds
Estimated cable distance:       3 meters
Buffers required:               1 (for 2112 byte frames at 16Gbps speed)s
```

Checking Diagnostic Port Results from a Cisco Fibre Channel 16G Switch

The link diagnostic port result from a Cisco 16G switch can be viewed while the switch port is configured as a link diagnostic port:

```
CISCO32G1# show diagnostic result interface fc1/26 test link-diag
PWWN of peer port: 21:00:34:80:0d:73:05:83
Status: Supported (Reflector)
Reflector loopback capabilities:  Xcvr-optical Electrical
Time of Test: Wed Aug 24 15:17:00 2022
Total time taken: 42 seconds
```

```
==================|================|================|==========================|=============
=====|=================
                  |                |                |     Discards             |    Latency (ns)  |
Loopback Level     |    Tx Frames    |    Rx Frames    |  IN    |  OUT   |BAD
WORDS|In-Switch|External|    Status
==================|================|================|========|========|=========|=========|===
=====|=================
Remote-Switched(R) |               0|               0|       0|       0|        0|        0|
0|       -NA-
MAC(R)             |               0|               0|       0|       0|        0|        0|
0|       -NA-
Xcvr-optical(R)    |         48122705|         48122705|       0|       0|        0|     3644|
2684|    Success
Electrical(R)      |            20000|            20000|            -NA-           |     -NA-|
-NA-|    Success
=============================================================================================================
======================


Overall Status                          : Success
Cable Length (approx. +/- 5 meters)     : 10.5 meters
```

Verifying D_Port Mode with Interactive QConvergeConsole CLI

To verify D_Port mode in interactive QConvergeConsole CLI:

1. From the main menu, select the **Adapter Information** option, and then select the adapter type (**Fibre Channel Adapter**).
2. From the **FC Adapter Information** menu, select **2** (**FC Port Information**).

   The Adapter Information menu shows the affected port with a status of **Diag Port**.

```
QConvergeConsoleCLI
    CLI - Version 2.6.x <Build xx>

FC Adapater Information

1:  FC Adapter Informaton
2:  FC Port Information
3:  FC VPD Information
4:  FC Storage Device Information
5:  FC Hyper-V VFC Information
6:  FC Adapter FMB Information
```

```
        <p or 0: Previous Menu; m or 98: Main Menu; ex or 99: Quit>
        Please Enter Selection : 2

        QConvergeConsoleCLI
        CLI Version 2.6.x <Build xx>

    Adapter Information

    HBA Model QLE26xx SN: XXXXXXXXXXXXA
      1: Port   1: WWPN: xx:xx:xx:xx:xx:xx:xx:70 Diag Port
      2: Port   2: WWPN: xx:xx:xx:xx:xx:xx:xx:71 Online
    HBA Model QLE26xx SN: XXXXXXXXXXXXB
      3: Port   1: WWPN: xx:xx:xx:xx:xx:xx:xx:40 SFP not installed
      4: Port   2: WWPN: xx:xx:xx:xx:xx:xx:xx:41 SFP not installed
```

Verifying D_Port Mode with Noninteractive QConvergeConsole CLI

To verify D_Port mode in noninteractive QConvergeConsole CLI, issue the following command:

quacli -pr fc -g

All the 2600 Series Adapters are listed. The affected port shows a status of **Diag Port**.

```
C:\>qaucli -pr fc -g
-------------------------------------------------------------------------------------------------
Host Name                       : WIN-XXXXXXXXXXX
OS Type                         : Microsoft Windows Server 2008 Enterprise 32-bit x86
OS Version                      : Service Pack 2 (Build 6002)
SDM API Version                 : 1..28.0.116 QLSDM.DLL
-------------------------------------------------------------------------------------------------
HBA Model QLE26xx (SN: XXXXXXXXXXXXA)
  1: Port   1: WWPN: xx:xx:xx:xx:xx:xx:xx:cb (HBA instance 4) Online
  2: Port   2: WWPN: xx:xx:xx:xx:xx:xx:xx:cf (HBA instance 5) Online
HBA Model QLE26xx (SN: XXXXXXXXXXXXB)
  1: Port   1: WWPN: xx:xx:xx:xx:xx:xx:xx:70 (HBA instance 2) Diag Port
  2: Port   2: WWPN: xx:xx:xx:xx:xx:xx:xx:71 (HBA instance 3) Online
HBA Model QLE26xx (SN: XXXXXXXXXXXXC)
  1: Port   1: WWPN: xx:xx:xx:xx:xx:xx:xx:40 (HBA instance 0) SFP not installed
  2: Port   2: WWPN: xx:xx:xx:xx:xx:xx:xx:41 (HBA instance 1) SFP not installed
```

```
-------------------------------------------------------------------------------------------
Total QLogic HBA(s): 3
```

For more information, refer to the section about the Fibre Channel noninteractive commands in the *User's Guide—QConvergeConsole CLI* (part number SN0054667-00).

Verifying D_Port Mode with QConvergeConsole Plug-ins

For information about verifying D_Port mode using QConvergeConsole VMware vCenter Server Plug-in, see the "Managing Fibre Channel Adapters" section in the *User's Guide—QConvergeConsole Plug-ins for vSphere* (part number TD-000965).

B Specifications

This appendix provides specifications for the Q*x*E2600-DEL, QLE2690-DEL, and QLE2690-DEL Adapters.

Physical Characteristics

Table B-1 lists the physical characteristics for the Q*x*E2600-DEL, QLE2690-DEL, and QLE2690-DEL Adapters.

Table B-1. Physical Characteristics

Adapter	Type	Bracket Installed	Length	Width
QLE2660-DEL	Low profile PCIe card	Full height	6.6 inches	2.731 inches
QLE2662-DEL	Low profile PCIe card	Full height	6.6 inches	2.731 inches
QME2662-DEL	Dell PowerEdge mezzanine adapter	—	Conforms to the Dell mezzanine card form factor	
QLE2690-DEL	Low profile PCIe card	Full height	6.6 inches	2.731 inches
QLE2690L-DEL	Low profile PCIe card	Low profile	6.6 inches	2.731 inches
QLE2690N-DEL	Low profile PCIe card	Full height	6.6 inches	2.731 inches
QLE2690NL-DEL	Low profile PCIe card	Low profile	6.6 inches	2.731 inches
QLE2692-DEL	Low profile PCIe card	Full height	6.6 inches	2.731 inches
QLE2692L-DEL	Low profile PCIe card	Low profile	6.6 inches	2.731 inches
QLE2692N-DEL	Low profile PCIe card	Full height	6.6 inches	2.731 inches
QLE2692NL-DEL	Low profile PCIe card	Low profile	6.6 inches	2.731 inches
QME2692-DEL	Dell PowerEdge mezzanine adapter	—	Conforms to the Dell mezzanine card form factor	

Power Requirements

Table B-2 lists the power requirements for the Q*x*E2600-DEL, QLE2690-DEL, and QLE2690-DEL Adapters.

Table B-2. Power Requirements

Adapter	Nominal Power	Maximum Power
QLE2660-DEL	13W	15W
QLE2662-DEL	14W	16W
QME2662-DEL	12W	14W
QLE2690-DEL	8.6W	10.65W
QLE2690L-DEL	8.6W	10.65W
QLE2690N-DEL	8.6W	8.6W
QLE2690NL-DEL	8.6W	8.6W
QLE2692-DEL	9.3W	11.5W
QLE2692L-DEL	9.3W	11.5W
QLE2692N-DEL	9.3W	11.5W
QLE2692NL-DEL	9.3W	11.5W
QME2692-DEL	7W	8.3W

Standards Specifications

Standards specifications for the Q*x*E2600-DEL, QLE2690-DEL, and QLE2690-DEL Adapters include the following:

- Fibre Channel Tape (FC-TAPE) Profile
- PCI Express Base Specification, rev. 3.0
- PCI Express Card Electromechanical Specification, rev. 2.0
- PCI Bus Power Management Interface Specification, rev. 1.2
- SCSI-3 Fibre Channel Protocol (SCSI-FCP)
- SCSI Fibre Channel Protocol-2 (FCP-2)
- Second Generation FC Generic Services (FC-GS-2)
- Third Generation FC Generic Services (FC-GS-3)
- NVMe over Fibre Channel v2 (FC-NVMe-2)

Interface Specifications

Table B-3 list the interface specifications for the QxE2600-DEL, QLE2690-DEL, and QLE2690-DEL Adapters.

Table B-3. Interface Specifications

Port Type	Media
16Gb SFP+	SFF-8431 Specification for SFP+, Revision 3.2

Environmental Specifications

Table B-4 lists the environmental specifications for the QxE2600-DEL, QLE2690-DEL, and QLE2690-DEL Adapters.

Table B-4. Environmental Specifications

Characteristic		Value
Temperature		
Operational:	QME2662-DEL	0°C to 65°C (32°F to 149°F)
	All other models	0°C to 55°C (32°F to 131°F)
Storage:	All models	–20°C to 70°C (–4°F to 158°F)
Humidity		
Relative (noncondensing)		10% to 90%
Storage		5% to 95%

C Regulatory Information

This appendix provides warranty, regulatory, and compliance information for the Fibre Channel Adapters covered in this guide.

Warranty

For information about your Dell warranty, see your system documentation.

Regulatory and Compliance Information

Laser Safety, FDA Notice

This product complies with DHHS Rules 21CFR Chapter I, Subchapter J. This product has been designed and manufactured according to IEC60825-1 on the safety label of laser product.

CLASS I LASER

Class 1 Laser Product	**Caution**—Class 1 laser radiation when open Do not view directly with optical instruments
Appareil laser de classe 1	**Attention**—Radiation laser de classe 1 Ne pas regarder directement avec des instruments optiques
Produkt der Laser Klasse 1	**Vorsicht**—Laserstrahlung der Klasse 1 bei geöffneter Abdeckung Direktes Ansehen mit optischen Instrumenten vermeiden
Luokan 1 Laserlaite	**Varoitus**—Luokan 1 lasersäteilyä, kun laite on auki Älä katso suoraan laitteeseen käyttämällä optisia instrumenttej

Agency Certification

The following section contains a summary of EMI/EMC test specifications performed on model HD8310405 and CU0310421 to comply with radiated emission, radiated immunity, and product safety standards.

EMI and EMC Requirements

FCC Rules, CFR Title 47, Part 15, Subpart B: Class A

FCC compliance information statement: This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Industry Canada, ICES-003:2004: Class A

This Class A digital apparatus complies with Canadian ICES-003.
Cet appareil numérique de la classe A est conformé à la norme NMB-003 du Canada.

CE Mark 2004/108/EC EMC Directive Compliance

EN 55032:2012-13 CISPR 32:2015 Class A
EN55035:2017
EN61000-3-2:2014 EN61000-3-3: 2013
EN61000-3-3:

EN61000-4-2: ESD
EN61000-4-3: RF Electro Magnetic Field
EN61000-4-4: Fast Transient/Burst
EN61000-4-5: Fast Surge Common/Differential
EN61000-4-6: RF Conducted Susceptibility
EN61000-4-8: Power Frequency Magnetic Field
EN61000-4-11: Voltage Dips and Interrupt

VCCI: CISPR 32:2016; Class A

AS/NZS; CISPR 32:2015 Class A

Korea Certification

KC-RRA KN32 Class A, KN35

Product Name/Model—2600 Series Adapter
Certification Holder—QLogic Corporation
Manufactured Date—Refer to date code listed on product
Manufacturer/Country of Origin—QLogic Corporation/USA

A class equipment (Business purpose info/ telecommunications equipment)	As this equipment has undergone EMC registration for business purpose, the seller and/or the buyer is asked to beware of this point and in case a wrongful sale or purchase has been made, it is asked that a change to household use be made.

Korean Language Format—Class A

A급 기기 (업무용 정보통신기기)

이 기기는 업무용으로 전자파적합등록을 한 기기이오니 판매자 또는 사용자는 이 점을 주의하시기 바라며, 만약 잘못판매 또는 구입하였을 때에는 가정용으로 교환하시기 바랍니다.

Product Safety Compliance

UL, cUL Product Safety

UL60950-1 (2nd Edition), A1 + A2 2014-10-14
CSA C22.2 60950-1-07 (2nd Edition) A1 +A2 2014-10

Use only with listed ITE or equivalent.

Complies with 21 CFR 1040.10 and 1040.11.

Complies with 21 CFR 1040.10 and 1040.11.

For use with UL Listed Servers or equivalent.

Power off the computer and all attached devices such as monitors, printers, and external components before installing the accessory card or equivalent.

2006/95/EC Low Voltage Directive

TUV EN60950-1:2006+A11+A1+A12+A2 2nd Edition

IEC 60950-1 A1+A2

IEC 62368 2nd and 3rd Edition

TUV EN62368-1:2014+A11:2017

D Adapter LEDs

This appendix describes the LEDs on the 2600 Series Adapters.

Figure D-1 shows the QLE2660, 16Gb Adapter ports with three LEDs:

- The amber LED on the left is 16Gb.
- The green LED in the middle is 8Gb.
- The amber LED on the right is 4Gb.

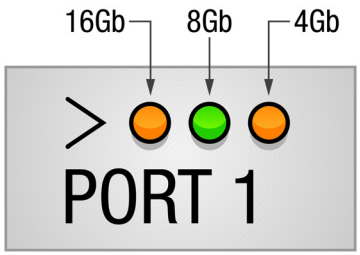

Figure D-1. QLE2660 16Gb Adapter LEDs

Table D-1 defines the LED behavior for the QLE2660, 16Gb Adapters.

Table D-1. QLE2660 16Gb Adapter LED Scheme

LED	Power Off	Power On (Before Firmware Initiali-zation)	Power On (After Firmware Initiali-zation)	Firmware Fault	4Gbps Link Up and Active	8Gbps Link Up and Active	16Gbps Link Up and Active	Beaconing
Amber LED (16Gbps)	Off	On	Flashing	Flashing in sequence	Off	Off	On or flashing	Flashing
Green LED (8Gbps)	Off	On	Flashing	Flashing in sequence	Off	On or flashing	Off	Off
Amber LED (4Gbps)	Off	On	Flashing	Flashing in sequence	On or flashing	Off	Off	Flashing

Figure D-2 shows the QLE2690, 16Gb Adapter ports with three LEDs:

- The green LED on the left is 16Gb.
- The green LED in the middle is 8Gb.
- The green LED on the right is 4Gb.

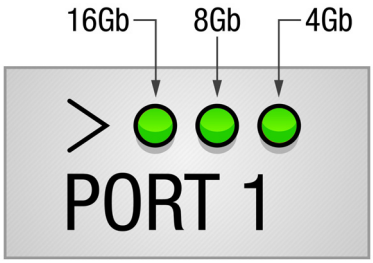

Figure D-2. QLE2690 16Gb Adapter LEDs

Table D-1 defines the LED behavior for the QLE2690, 16Gb Adapters.

Table D-2. QLE2690 16Gb Adapter LED Scheme

LED	Power Off	Power On (Before Firmware Initialization)	Power On (After Firmware Initialization)	Firmware Fault	4Gbps Link Up and Active	8Gbps Link Up and Active	16Gbps Link Up and Active	Beaconing
Green LED (16Gbps)	Off	On	Flashing	Flashing in sequence	Off	Off	On or flashing	Flashing
Green LED (8Gbps)	Off	On	Flashing	Flashing in sequence	Off	On or flashing	Off	Off
Green LED (4Gbps)	Off	On	Flashing	Flashing in sequence	On or flashing	Off	Off	Flashing

E Revision History

Document Revision History
Revision A, November 20, 2012
Revision B, April 8, 2013
Revision C, September 12, 2013
Revision D, June 9, 2014
Revision E, January 22, 2015
Revision F, July 1, 2015
Revision G, March 21, 2016
Revision H, April 19, 2016
Revision J, February 1, 2017
Revision K, August 24, 2017
Revision L, December 19, 2017
Revision M, April 13, 2018
Revision N, January 29, 2019
Revision P, June 14, 2019
Revision R, October 28, 2019
Revision T, April 3, 2020
Revision W, July 9, 2020
Revision X, January 21, 2021
Revision Y, May 14, 2021
Revision Z, July 20, 2021
Revision ZA, September 29, 2021
Revision ZB, December 15, 2021
Revision ZC, June 30, 2022
Revision ZD, October 21, 2022

Changes	Sections Affected
Removed support for Windows Server 2016, VMware 6.*x*. Added support for RHEL8.0, RHEL 9.0. Removed support for out of Order Frame Reassembly (OoOFR). Updated menu names, options, and output examples in sections containing instructions for QConvergeConsole CLI. Removed support for QConvergeConsole VMware vSphere Web Client Plug-in.	All
Added sub-bullets under USCM bullet.	"Key Features" on page xiv
Corrected the description to "The Restore Defaults option restores the **NVRAM** settings to the defaults."	"Restore Default Settings" on page 39
Added a second paragraph indicating that "This section is for use with `qla2xxx-kmp-default-<driver-version_kernel-version>-<release>.rpm` package." In Step 1, added instructions for what to do when a warning appears.	"Importing a Public Key for Secure Boot" on page 13
Updated section title and operating systems that apply (first paragraph).	"Building the Driver for RHEL 8.0 and 9.0" on page 14
Updated section title and operating systems that apply (first paragraph).	"Building the Driver for SLES 15" on page 16
In the first paragraph, third sentence, removed the Brocade switch version number.	"Configuring VM-ID" on page 29
Updated instructions to add UEFI boot devices as well as (legacy) BIOS boot devices.	"Configuring Boot Devices with Noninteractive QConvergeConsole CLI" on page 44
Added a NOTE that this procedure does not work in VMware ESXi 8.0.	"Configuring Boot Devices with QConvergeConsole Plug-in for VMware vSphere HTML5 Client" on page 46
Updated instructions to specify creating a virtual port with either an automatic or specific WWPN.	"Configuring NPIV with Noninteractive QConvergeConsole CLI" on page 46

In the first paragraph, clarified the second sentence to "The QLogic QoS solution is based on assigning QoS levels to virtual ports (NPIV ports), **in contrast to end-to-end QoS, which assigns QoS levels from initiator to target across the fabric**."	"Configuring NPIV Quality of Service" on page 48
In the description of the `para alias` parameter, added valid values.	"Configuring Fibre Channel Driver Parameters with Noninteractive QConvergeConsole CLI" on page 51
In Step 10, in the NOTE, added a new paragraph about restoring defaults in HII.	"Configuring the UEFI Driver" on page 52
In the second paragraph, added a new first sentence, "This feature is enabled by default in the 2600 Series Adapter Windows and VMware drivers."	"Configuring End-to-End CS_CTL QoS" on page 79
Updated the example outputs of the `fdmi show` command.	"FDMI Enhancements Support" on page 86
In the NOTE, updated the Brocade and Cisco switch OS versions.	"Read Diagnostic Parameters Extended Link Service Command" on page 91
Added explanation of using FPINs with USCM. Added a second paragraph to indicate that "USCM support can be controlled (enabled/disabled) per port on all Marvell QLogic management tools." Clarified that the examples are for QConvergeConsole CLI. Added a paragraph referencing the Related Materials section for documents that provide USCM information.	"Universal SAN Congestion Mitigation (USCM)" on page 93
In the list of supported drivers, updated ESX driver versions. Added a list of management tools that support USCM.	"Prerequisites" on page 94
Added instructions for creating I/O throttling profiles.	"Universal SAN Congestion Mitigation (USCM)" on page 93, "USCM Profiles" on page 95, "Using Driver Modules or Registry Parameters for USCM Profiles" on page 96
Expanded FPIN support for additional switches and OSs. Added new second, third paragraphs.	"FPIN Link Integrity Marginal Path Detection Support" on page 98

Changed title to reflect all RHEL supported OSs.	"FC Transport daemon Installation for RHEL 8.4, 8.5, and 8.6" on page 98
Added new section for multipath configuration file changes for RHEL and SLES OSs.	"Multipath Configuration File Changes for RHEL 9.0 and SLES 15 SP3 and Later" on page 100
Added a paragraph directing the reader to see the latest Readme and Release Notes for required versions of drivers, firmware, and management tools.	"Hardware and Software Requirements" on page 101
In the NOTE, updated the Brocade and Cisco switch versions.	"Read Diagnostic Parameters Extended Link Service Command" on page 91
In the NOTE, updated the first sentence to "The ESX7.0**/8.0** `qlnativefc` driver…"	"Verifying in ESX" on page 107
In the first NOTE, clarified the first sentence to "NVMe BFS is only supported on SLES15 SP3 **and later**…" In the second paragraph, added SLES 15 SP4.	"Boot From SAN Setup" on page 109
Changed title to include all versions after ESX 7.0. Added new Step 6.	"Additional Setup for NVME Boot-from-SAN Using ESX 7.0 and Later" on page 110
In Figure A-1, added a No option to "Storage device logged into Fibre Channel Switch?" In Figure A-2, added a Yes option to "Fibre Channel switch set for auto-negotiation?"	"Fibre Channel Troubleshooting Diagram" on page 114
Added introductory paragraph. Added instructions for viewing D_Port test results. At the end of the section, added a NOTE indicating that the adapter port switches to an N_Port after the diagnostics port is disabled on the switch port.	"Configuring D_Port on a Brocade Fibre Channel 16G Switch" on page 117,
At the end of the section, added a NOTE indicating that the adapter port switches to an N_Port after the diagnostics port is disabled on the switch port.	"Configuring and Running Diagnostics Port on a Cisco Fibre Channel 16G Switch" on page 118
Added a section for checking diagnostic results from a Cisco Fibre Channel Switch.	"Checking Diagnostic Port Results from a Cisco Fibre Channel 16G Switch" on page 120

Removed the figure and replaced with an updated output example; updated corresponding text.	"Verifying D_Port Mode with Interactive QConvergeConsole CLI" on page 121, "Verifying D_Port Mode with Noninteractive QConvergeConsole CLI" on page 122

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 www.marvell.com.

Doc. No. CU0354604-00 Rev. ZD Revised: October 21, 2022